

Energy-Motivated Equivariant Pretraining for 3D Molecular Graphs

Rui Jiao^{1,2}, Jiaqi Han^{1,2}, Wenbing Huang^{4,5,*}, Yu Rong⁶, Yang Liu^{1,2,3,*}

¹ Beijing National Research Center for Information Science and Technology (BNRist),

Department of Computer Science and Technology, Tsinghua University

² Institute for AI Industry Research (AIR), Tsinghua University

³ Beijing Academy of Artificial Intelligence

⁴ Gaoling School of Artificial Intelligence, Renmin University of China

⁵ Beijing Key Laboratory of Big Data Management and Analysis Methods

⁶ Tencent AI Lab

{jiaor21,hanjq21}@mails.tsinghua.edu.cn, hwenbing@126.com,

yu.rong@hotmail.com, liuyang2011@tsinghua.edu.cn

Abstract

Pretraining molecular representation models without labels is fundamental to various applications. Conventional methods mainly process 2D molecular graphs and focus solely on 2D tasks, making their pretrained models incapable of characterizing 3D geometry and thus defective for downstream 3D tasks. In this work, we tackle 3D molecular pretraining in a complete and novel sense. In particular, we first propose to adopt an equivariant energy-based model as the backbone for pretraining, which enjoys the merits of fulfilling the symmetry of 3D space. Then we develop a node-level pretraining loss for force prediction, where we further exploit the Riemann-Gaussian distribution to ensure the loss to be $E(3)$ -invariant, enabling more robustness. Moreover, a graph-level noise scale prediction task is also leveraged to further promote the eventual performance. We evaluate our model pretrained from a large-scale 3D dataset GEOM-QM9 on two challenging 3D benchmarks: MD17 and QM9. Experimental results demonstrate the efficacy of our method against current state-of-the-art pretraining approaches, and verify the validity of our design for each proposed component. Code is available at <https://github.com/jiaor17/3D-EMGP>.

Introduction

Learning informative molecular representation is a fundamental step for various downstream applications, including molecular property prediction (Gilmer et al. 2017; Kearnes et al. 2016), virtual screening (Wallach, Dzamba, and Heifets 2015; Zheng, Fan, and Mu 2019), and Molecular Dynamics (MD) simulation (Chmiela et al. 2017). Recent methods exploit Graph Neural Networks (GNNs) (Gilmer et al. 2017; Xu et al. 2019) for their power in capturing the topology of molecules, which yet is label-hungry and thus powerless for real scenarios when molecular annotations are unavailable. Therefore, the research attention has been paid to the *self-supervised* pretraining paradigm, to construct the surrogate task by exploring the intrinsic structure within unlabeled molecules. A variety of self-supervised works have

been proposed, ranging from generative-based models (Kipf and Welling 2016; Hu et al. 2020a,b) to contrastive learning (Sun et al. 2020; Velickovic et al. 2019; You et al. 2020, 2021a; Sun et al. 2021).

In many applications, using and analyzing 3D geometry is crucial and even indispensable; for instance, we need to process 3D coordinates for energy prediction in MD simulation or protein-ligand docking. Owing to the fast development in data acquisition, it is now convenient to access large-scale unlabeled molecules with rich 3D conformations (Axelrod and Gomez-Bombarelli 2022). It would be quite exciting if we could develop techniques to obtain pretrained models from these unlabeled molecules for 3D tasks with limited data. Nevertheless, existing self-supervised methods (Hu et al. 2020a; You et al. 2020; Rong et al. 2020) are weak in leveraging the 3D geometry information. First, from the input side, the backbone models they pretrain can only process the input of 2D molecules without the consideration of 3D coordinates. As demonstrated by Schütt et al. (2017), certain molecular properties (*e.g.* potential energy) are closely related to the 3D structure with which they can be better predicted. Second, for the output side, their pretraining tasks are not 3D-aware, making the knowledge they discover less generalizable in 3D space. Recently, the study by Liu et al. (2021) proposes to impose the 3D information for pretraining; however, its goal is still limited to enhancing 2D models for 2D tasks.

In this paper, we investigate 3D molecular pretraining in a complete and novel sense: using 3D backbone models, designing 3D-aware pretraining tasks, and targeting 3D downstream evaluations. However, this is not trivial by any means. The challenges mainly stem from how to maintain the symmetry of our biological world—rotating/translating the 3D conformation of a molecule does not change the law of its behavior. Mathematically, we should make the backbone $E(3)$ -equivariant, and the pretraining loss $E(3)$ -invariant, where the group $E(3)$ collects the transformations of rotations, reflections, and translations (Satorras, Hoogeboom, and Welling 2021). Unfortunately, typical GNNs (Gilmer et al. 2017; Xu et al. 2019) and 3D losses based on Euclidean distance (Luo and Hu 2020) do not satisfy such constraints.

*Wenbing Huang and Yang Liu are the corresponding authors.
Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

To address the above challenges, this paper makes the following contributions: **1.** We propose an energy-based representation model that predicts E(3)-equivariant force for each atom in the input 3D molecule, by leveraging recently-proposed equivariant GNNs (Satorras, Hoogeboom, and Welling 2021; Schütt et al. 2017; Thölke and De Fabritiis 2021) as the building block. **2.** To pretrain this model, we formulate a physics-inspired node-level force prediction task, which is further translated to a 3D position denoising loss in an equivalent way. More importantly, we develop a novel denoising scheme with the aid of the proposed Riemann-Gaussian distribution, to ensure the E(3)-invariance of the pretraining task. **3.** We additionally design a graph-level surrogate task on 3D molecules, in line with the observation from traditional 2D methods (Rong et al. 2020) that performing node-level and graph-level tasks jointly is able to promote the eventual performance. For this purpose, we teach the model to identify the noise scale of the input tuple consisting of a clean sample and a noisy one. The above ingredients are unified in a general pretraining framework: energy-motivated 3D Equivariant Molecular Graph Pretraining (3D-EMGP).

We pretrain our model on a large-scale dataset with 3D conformations: GEOM-QM9 (Axelrod and Gomez-Bombarelli 2022), and then evaluate its performance on the two popular 3D tasks: MD17 (Chmiela et al. 2017) and QM9 (Ramakrishnan et al. 2014). Extensive experiments demonstrate that our model outperforms state-of-the-art 2D approaches, even though their inputs are augmented with 3D coordinates for fair comparisons. We also inspect how the performance changes if we replace the components of our architecture with other implementations. The results do support the optimal choice of our design.

Related Works

Self-supervised molecular pretraining Self-supervised learning has been well developed in the field of molecular graph representation learning. Many pretraining tasks have been proposed to extract information from large-scale molecular dataset, mainly divided into three categories: *contrastive*, *generative* and *predictive*. Contrastive methods aim to maximize the mutual information between different views of the same graph (Sun et al. 2020; Velickovic et al. 2019; You et al. 2020, 2021a; Sun et al. 2021), while generative methods focus on reconstructing the information from different levels of the 2D topological structure (Kipf and Welling 2016; Hu et al. 2020a,b). As for predictive methods, they learn the molecule representation by predicting the pseudo-labels created from the input graphs. For example, GROVER (Rong et al. 2020) proposes to classify the subgraph structure and predict the existence of specific motifs, which leverages domain knowledge into molecule pretraining. However, all the above methods mainly focus on pretraining 2D GNNs without 3D information. More recently, several methods propose to tackle 3D graphs, including 3D Infomax (Stärk et al. 2022) that maximizes the mutual information between the representations encoded from a 2D and a 3D model, and GraphMVP (Liu et al. 2021) which uses contrastive and generative methods to incorporate 3D informa-

tion into the 2D model. However, the motivation of these two methods remains to benefit 2D models with 3D information. On the contrary, this paper attempts to pretrain 3D models via 3D objectives with the usage for 3D downstream tasks. Besides, ChemRL-GEM (Fang et al. 2022) predicts bond length, bond angle prediction, and atom distance for 3D pretraining. While they only employ pairwise or triplet invariant tasks, we formulate both invariant and equivariant pretraining objectives in accordance with our proposed energy-based molecular representation model. Zhu et al. (2022) formulates the pretraining task as translation between 2D and 3D views, distinct from our goal of realizing pretraining mainly based on 3D conformations. Zhou et al. (2023) applies masked position denoising, similar to the implemented baseline PosPred in our experiments.

Equivariant graph neural networks To better analyze the physical and chemical properties of molecules, many researchers regard the molecules as geometric graphs, which additionally assign 3D coordinates on each atom apart from the 2D topological information. Geometric graphs present rotational, translational, and/or reflectional symmetry, as the properties are invariant and the dynamic processes are equivariant to the E(3) or SE(3) transformations in 3D space. To introduce this inductive bias, geometrically equivariant graph neural networks have been proposed to model the geometric graphs. According to Han et al. (2022), current 3D GNNs achieve equivariance mainly in three ways: extracting irreducible representations (Thomas et al. 2018; Fuchs et al. 2020), utilizing group regular representations (Finzi et al. 2020; Hutchinson et al. 2021) or transforming the 3D vectors into invariant scalars (Schütt et al. 2017; Satorras, Hoogeboom, and Welling 2021; Thölke and De Fabritiis 2021). Previous works showcase the superiority of equivariant models on several 3D molecular tasks (Thölke and De Fabritiis 2021; Liu et al. 2022), and our goal is to further improve the performance of 3D models via equivariant pretraining on large-scale 3D datasets.

Method

Energy-based Molecular Modeling

In general, a molecule consisting of N atoms can be modeled as a *molecular graph* $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of nodes (atoms) and \mathcal{E} is the set of edges modeling the connectivity, e.g., bond connection or spatial proximity. Each atom is assigned a node feature $h_i, 1 \leq i \leq N$, representing the properties of the atom, such as atom type. In this work, we elaborate on the 3D information of a molecule, that is, apart from the node attribute $h_i \in \mathbb{R}^m$ for atom i , we extra consider the 3D coordinate $\mathbf{x}_i \in \mathbb{R}^3$. We denote the configurations of all nodes as $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N] \in \mathbb{R}^{3 \times N}$, and similarly the node features as $\mathbf{H} = [h_1, h_2, \dots, h_N] \in \mathbb{R}^{m \times N}$. Our goal is to pretrain a capable GNN model φ that can be generally applied to different downstream 3D tasks, which is depicted as $\varphi(\mathbf{X}, \mathbf{H}, \mathcal{E})$. We hereafter omit the input of \mathcal{E} for conciseness, unless otherwise specified.

Unlike conventional 2D molecular graph pretraining, we are now provided with vital 3D information, making it possible to leverage the rich geometric context in a unified frame-

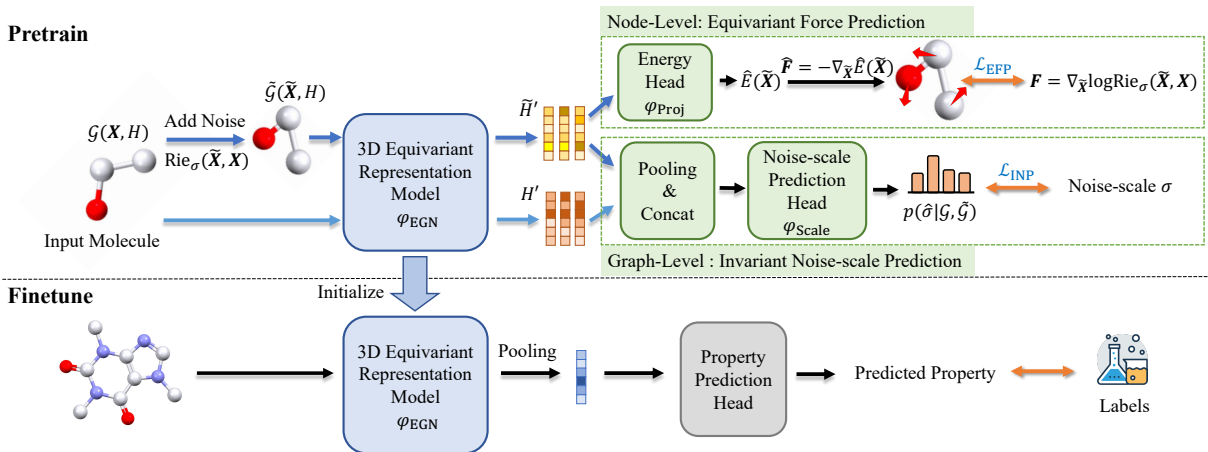


Figure 1: An overview of our 3D-EMGP. It consists of two tasks: node-level equivariant force prediction and graph-level invariant noise scale prediction. \mathbf{X} , $\tilde{\mathbf{X}}$ are the original and perturbed coordinates. H is the input node feature and H' , \tilde{H}' are the output features of the original and perturbed graph. $\text{Rie}_\sigma(\tilde{\mathbf{X}} | \mathbf{X})$ is the proposed Riemann-Gaussian distribution in Eq. (9).

work. However, involving 3D conformation is not free of difficulty, and one constraint we should enforce is to make φ equivariant for vector outputs, and invariant for scalar outputs. The notion of equivariance/invariance is defined below.

Definition 1 A GNN model φ is called $E(3)$ -equivariant, if for any transformation $g \in E(3)$, $\varphi(g \cdot \mathbf{X}, H) = g \cdot \varphi(\mathbf{X}, H)$; and it is called $E(3)$ -invariant if $\varphi(g \cdot \mathbf{X}, H) = \varphi(\mathbf{X}, H)$.

In Definition 1, the group action $g \cdot \mathbf{X}$ is implemented as matrix multiplication $\mathbf{O}\mathbf{X}$ for orthogonal transformation $\mathbf{O} \in \mathbb{R}^{3 \times 3}$ and addition $\mathbf{X} + \mathbf{t}$ for translation $\mathbf{t} \in \mathbb{R}^3$. Basically, for an equivariant function, the output will translate/rotate/reflect in the same way as the input, while for the invariant case, the output stays unchanged whatever group action conducted on the input. Equivariance/invariance is so essential that it characterizes the symmetry of the 3D biochemistry: rotating or translating a molecule will never change its potential energy. There is a variety of $E(3)$ -equivariant GNNs (Han et al. 2022) that can be utilized as our backbone. By choosing an $E(3)$ -equivariant GNN, φ is instantiated as φ_{EGN} . Notably, φ_{EGN} is also permutation equivariant regarding the order of the columns in \mathbf{X} and H .

We now introduce our idea of how to tackle 3D molecular pretraining in a more domain-knowledge-reliable manner. As well studied in biochemistry, the interaction between atoms in 3D space is captured by the forces and potentials, depending on the positions of the atoms, *i.e.*, the molecular conformation. This connection inspires us to incorporate the concept of energy and force into our representation model, making room for designing unsupervised pretraining objectives with 3D geometric information. In light of this, we introduce a representation model that jointly takes into account both energy and force. We denote the energy of a molecule as $E \in \mathbb{R}$ and the resultant interaction force exerting on atom i as $\mathbf{f}_i \in \mathbb{R}^3$, $1 \leq i \leq N$. The forces over all atoms are collected in the matrix $\mathbf{F} \in \mathbb{R}^{3 \times N}$. Clearly, E is an invariant graph-level scalar, while \mathbf{F} consists of equivariant node-level vectors, in terms of the input transformation.

To derive E and \mathbf{F} by the equivariant model φ_{EGN} , we first obtain a node-level representation in the latent space:

$$H' = \varphi_{\text{EGN}}(\mathbf{X}, H, \mathcal{E}), \quad (1)$$

where $H' \in \mathbb{R}^{k \times N}$ is an invariant representation. Let \hat{E} , $\hat{\mathbf{F}}$ denote the predicted energy and force. We yield the graph-level energy of the molecule via a graph pooling operation:

$$\hat{E}(\mathbf{X}) = \varphi_{\text{Proj}}\left(\sum_{i=1}^N h'_i\right), \quad (2)$$

where h'_i is the i -th column of H' , $\varphi_{\text{Proj}}: \mathbb{R}^k \mapsto \mathbb{R}$ is the projection head, realized by a Multi-Layer Perceptron (MLP). Essentially, force corresponds to the direction that causes potential energy to decrease, which implies

$$\hat{\mathbf{F}}(\mathbf{X}) = -\lim_{\Delta \mathbf{X} \rightarrow 0} \frac{\Delta \hat{E}}{\Delta \mathbf{X}} = -\nabla_{\mathbf{X}} \hat{E}(\mathbf{X}), \quad (3)$$

where $\nabla_{\mathbf{X}}$ denotes the gradient *w.r.t.* \mathbf{X} . It is easy to verify that \hat{E} is invariant and $\hat{\mathbf{F}}$ is equivariant¹.

We attempt to design the first proxy task by leveraging the predicted force $\hat{\mathbf{F}}$ to fit the force implied in the molecule. However, there is usually no force label provided in the pre-training dataset. Fortunately, we can fulfill this purpose from the lens of node-level denoising—we first add noise to each node’s coordinate and then estimate the virtual force that pulls the noisy coordinate back to the original one. Upon the denoising process in the first pretraining task, we further construct a graph-level pretraining objective in an intuitive sense: a desirable model should be able to tell how much noise is added to its input. The details of the two pretraining tasks are presented in the subsequent contents.

¹More precisely, $\hat{\mathbf{F}}$ is orthogonality-equivariant but translation-invariant.

Node-Level: Equivariant Force Prediction

We start by designing a pretraining objective that well characterizes the 3D geometric information. To fulfill this goal, we resort to the force $\hat{\mathbf{F}}$ produced by our energy-based molecular model. Yet, it is challenging and non-straightforward to provide a reasonable instantiation, since there is usually no available ground truth force labels in large-scale pretraining datasets. Interestingly, we find a way through by establishing a connection between \mathbf{F} and the distribution of the conformations \mathbf{X} , and manage to provide the predicted $\hat{\mathbf{F}}$ with pseudo supervision. The connection is identified by first assuming a Boltzmann energy distribution (Boltzmann 1868) over the training conformers \mathbb{G} :

$$p(\mathbf{X}) = \frac{1}{Z} \exp\left(-\frac{E(\mathbf{X})}{kT}\right), \quad (4)$$

where E denotes the assumed energy, k is the Boltzmann constant, T is the temperature, and Z is the normalization. By taking the logarithm and computing the gradient over \mathbf{X} , we acquire

$$\nabla_{\mathbf{X}} \log p(\mathbf{X}) \propto -\nabla_{\mathbf{X}} E(\mathbf{X}) := \mathbf{F}. \quad (5)$$

It is thus applicable to approach $\hat{\mathbf{F}}$ by the first term in Eq. (5), serving as a pseudo force. In light of this, we formulate an equivariant (pseudo) force prediction (EFP) objective over training data \mathbb{G} :

$$\mathcal{L}_{\text{EFP}} = \mathbb{E}_{\mathbf{G} \sim \mathbb{G}} [\|\hat{\mathbf{F}}(\mathbf{X}) - \nabla_{\mathbf{X}} \log p(\mathbf{X})\|_F^2], \quad (6)$$

where $\hat{\mathbf{F}}$ is produced by the model φ_{EGN} (Eq. (1-3)), $\|\cdot\|_F$ computes the Frobenius norm.

Nevertheless, we still have no idea of what the exact form of the data density $p(\mathbf{X})$ looks like. Hence it is infeasible to directly apply the loss Eq. (6). Fortunately, the work by Vincent (2011) draws a promising conclusion that Eq. (6) can be equivalently translated to a denoising problem which is tractable to solve (see Appendix for details). In a nutshell, we instead sample a noisy sample $\tilde{\mathbf{X}}$ from \mathbf{X} according to a certain conditional distribution, *i.e.*, $\tilde{\mathbf{X}} \sim p(\tilde{\mathbf{X}} | \mathbf{X})$. Then we substitute the noisy sample into the model φ_{EGN} and perform the replacement of Eq. (6) by

$$\mathcal{L}_{\text{EFP-DN}} = \mathbb{E}_{\mathbf{G} \sim \mathbb{G}, \tilde{\mathbf{X}} \sim p(\tilde{\mathbf{X}} | \mathbf{X})} [\|\hat{\mathbf{F}}(\tilde{\mathbf{X}}) - \nabla_{\tilde{\mathbf{X}}} \log p(\tilde{\mathbf{X}} | \mathbf{X})\|_F^2]. \quad (7)$$

We now discuss the formulation of the conditional probability $p(\tilde{\mathbf{X}} | \mathbf{X})$. Different from the traditional denoising process on images or other Euclidean data (Song and Ermon 2020; Shi et al. 2021; Luo et al. 2021), in our case when considering the 3D geometry, the noise we add should be geometry-aware other than conformation-aware. In other words, $p(\tilde{\mathbf{X}} | \mathbf{X})$ should be *doubly E(3)-invariant*, namely,

$$p(g_1 \cdot \tilde{\mathbf{X}} | g_2 \cdot \mathbf{X}) = p(\tilde{\mathbf{X}} | \mathbf{X}), \forall g_1, g_2 \in E(3), \quad (8)$$

with the illustrations provided in Fig. 2. This is consistent with our understanding: the behavior of molecules with the same geometry should be independent of different conformations. For example, when we rotate the sample $\tilde{\mathbf{X}}$, the




Figure 2: Illustration of different distributions. For typical Gaussian, a data point (in dashed circle) is a specific conformation \mathbf{X} , while for Riemann Gaussian, it is a set of conformations with the same geometry $[\mathbf{X}] := \{g \cdot \mathbf{X} | g \in E(3)\}$.

property of $p(\tilde{\mathbf{X}} | \mathbf{X})$ by Eq. (8) ensures the loss in Eq. (7) to be unchanged, which is what we desire; similarly, conducting rotation on \mathbf{X} should also obey the same rule. A conventional choice of $p(\tilde{\mathbf{X}} | \mathbf{X})$ is utilizing the standard Gaussian with noise scale σ : $p(\tilde{\mathbf{X}} | \mathbf{X}) = \mathcal{N}(\mathbf{X}, \sigma^2 \mathbf{I})$. This naive form fails to meet the doubly E(3)-invariant property in Eq. (8), which could cause mistaken supervision in Eq. (7). To show this, we derive $\nabla_{\tilde{\mathbf{X}}} \log p(\tilde{\mathbf{X}} | \mathbf{X}) = -\frac{\tilde{\mathbf{X}} - \mathbf{X}}{\sigma^2}$ as the force target; if we set $\tilde{\mathbf{X}} = \mathbf{R}\mathbf{X}$ for some rotation matrix $\mathbf{R} \neq \mathbf{I}$, then we have $\nabla_{\tilde{\mathbf{X}}} \log p(\tilde{\mathbf{X}} | \mathbf{X}) = -\frac{1}{\sigma^2}(\mathbf{R} - \mathbf{I})\mathbf{X} \neq 0$, which, however, does not align with the true fact that the force between $\tilde{\mathbf{X}} = \mathbf{R}\mathbf{X}$ and \mathbf{X} should be zero since they share the same geometry.

To devise the form with the symmetry in Eq. (8), we instead resort to Riemann-Gaussian distribution (Said et al. 2017) defined as follows:

$$p_{\sigma}(\tilde{\mathbf{X}} | \mathbf{X}) = \text{Rie}_{\sigma}(\tilde{\mathbf{X}} | \mathbf{X}) := \frac{1}{Z(\sigma)} \exp\left(-\frac{d^2(\tilde{\mathbf{X}}, \mathbf{X})}{4\sigma^2}\right), \quad (9)$$

where $Z(\sigma)$ is the normalization term, and d is the metric that calculates the difference between $\tilde{\mathbf{X}}$ and \mathbf{X} . Riemann-Gaussian is a generalization version of typical Gaussian, by choosing various distances d beyond the Euclidean metric. Here, to pursue the constraint in Eq. (8), we propose to use

$$d(\mathbf{X}_1, \mathbf{X}_2) = \|\mathbf{Y}_1^{\top} \mathbf{Y}_1 - \mathbf{Y}_2^{\top} \mathbf{Y}_2\|_F, \quad (10)$$

where $\mathbf{Y} = \mathbf{X} - \boldsymbol{\mu}(\mathbf{X})$ re-positions \mathbf{X} towards zero mean ($\boldsymbol{\mu}(\mathbf{X})$ denotes the mean of the columns in \mathbf{X}). One clear benefit is that the distance function d defined in Eq. (10) satisfies the doubly E(3)-invariance constraint in Eq. (8). Note that d is also permutation invariant with regard to the order of the columns of $\tilde{\mathbf{X}}$ and \mathbf{X} . We summarize the above discussion as a formal proposition as follows.

Proposition 1 For Riemann-Gaussian $\text{Rie}_{\sigma}(\tilde{\mathbf{X}} | \mathbf{X})$ defined in Eq. (9), it is doubly E(3)-invariant as per Eq. (8).

The gradient of Riemann-Gaussian is calculated as follows, with the detailed derivations in Appendix A.2:

$$\nabla_{\tilde{\mathbf{X}}} \log p_{\sigma}(\tilde{\mathbf{X}} | \mathbf{X}) = -\frac{1}{\sigma^2} [(\tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^{\top}) \tilde{\mathbf{Y}} - (\tilde{\mathbf{Y}} \mathbf{Y}^{\top}) \mathbf{Y}]. \quad (11)$$

Meanwhile, as proved in Appendix, the calculation in Eq. (11) is of the complexity $\mathcal{O}(N)$, making it computationally efficient for even large-scale molecules.

The last remaining recipe is how to sample $\tilde{\mathbf{X}}$ from \mathbf{X} according to the Riemann-Gaussian distribution to provide the input to Eq. (7). It is non-straightforward to accomplish this goal, since the normalization term $Z(\sigma)$ of Riemann-Gaussian is unknown. Here we resort to Langevin dynamics (Schlick 2010) which is widely used for approximated sampling when only non-normalized probability density is given. We provide the details in Appendix A.6. Furthermore, to better explore the conformation space, we employ a sampling scheme with multiple levels of noise (Song and Ermon 2020). Particularly, let $\{\sigma_l\}_{l=1}^L$ be a series of noises with different scales. The final EFP loss is given by

$$\mathcal{L}_{\text{EFP-Final}} = \mathbb{E}_{\mathcal{G} \sim \mathbb{G}, l \sim U(1, L), \tilde{\mathbf{X}} \sim p_{\sigma_l}(\tilde{\mathbf{X}}|\mathbf{X})} \left[\sigma_l^2 \left\| \frac{1}{\sigma_l} \hat{\mathbf{F}}(\tilde{\mathbf{X}}) - \frac{1}{\alpha} \nabla_{\tilde{\mathbf{X}}} \log p_{\sigma_l}(\tilde{\mathbf{X}}|\mathbf{X}) \right\|_F^2 \right], \quad (12)$$

where $U(1, L)$ is the discrete uniform distribution, and $\nabla_{\tilde{\mathbf{X}}} \log p_{\sigma_l}(\tilde{\mathbf{X}}|\mathbf{X})$ is provided by Eq. (11). We apply a weighting coefficient σ_l^2 for different noise scales and scale the predicted forces by $1/\sigma_l$ as suggested by Song and Ermon (2020); Shi et al. (2021). We also add α as a normalization for numerical stability of the inner product; its value is given by $\alpha = (\|\tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^\top\|_F + \|\tilde{\mathbf{Y}}\mathbf{Y}^\top\|_F)/2$ in our experiments. It is proved in Appendix A.5 that the normalization term α also satisfies the doubly E(3)-invariant property.

Graph-Level: Invariant Noise-scale Prediction

In the last subsection, we have constructed a node-level pretraining objective for local force prediction. To further discover global patterns within the input data, this subsection presents how to design a graph-level self-supervised task. Previous studies (Hu et al. 2020a; Rong et al. 2020) have revealed for 2D molecules that the node- and graph-level tasks are able to promote each other. Here, we investigate this idea on 3D geometric graphs.

Recalling that $\tilde{\mathbf{X}}$ is distributed by $p_{\sigma_l}(\tilde{\mathbf{X}}|\mathbf{X})$, it is expected that a well-behaved model should identify how much the perturbed sample deviates from the original data. Such intuition inspires us to set up a classification problem as noise scale prediction. Specifically, our φ_{EGN} shares the same EGN backbone as in Eq. (1), yielding exactly the same invariant node- and graph-level embeddings. For the input \mathbf{X} and $\tilde{\mathbf{X}}$, we first obtain their graph-level embedding \mathbf{u} and $\tilde{\mathbf{u}}$ via φ_{EGN} , respectively. Instead of using the scalar projection head φ_{Proj} for energy computation, we employ a classification head φ_{Scale} that takes as input a concatenation of the graph-level embeddings of the original conformation \mathbf{u} and the perturbed conformation $\tilde{\mathbf{u}}$. The output of φ_{Scale} is the logits $\mathbf{p} \in \mathbb{R}^L = \varphi_{\text{Scale}}(\mathbf{u} \parallel \tilde{\mathbf{u}})$, where L is the number of noise levels. Finally, a cross-entropy loss is computed between the logits and the label, which is the sampled noise level for the current input. The objective of the invariant noise-scale prediction task is thus given by

$$\mathcal{L}_{\text{INP}} = \mathbb{E}_{\mathcal{G} \sim \mathbb{G}, l \sim U(1, L), \tilde{\mathbf{X}} \sim p_{\sigma_l}(\tilde{\mathbf{X}}|\mathbf{X})} [\mathcal{L}_{\text{CE}}(\mathbb{I}[l], \mathbf{p})], \quad (13)$$

where \mathcal{L}_{CE} is the cross-entropy loss and $\mathbb{I}[l]$ is the one-hot encoding of l .

Our overall training objective, as illustrated in Fig. 1, is a combination of both node-level equivariant force prediction loss and graph-level invariant noise scale prediction loss:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{\text{EFP-Final}} + \lambda_2 \mathcal{L}_{\text{INP}}, \quad (14)$$

where λ_1 and λ_2 are the balancing coefficients.

Experiments

Experimental Setup

Pretraining dataset We leverage a large-scale molecular dataset GEOM-QM9 (Axelrod and Gomez-Bombarelli 2022) with corresponding 3D conformations as our pretraining dataset. Specifically, we select the conformations with top-10 Boltzmann weight² for each molecule, and filter out the conformations that overlap with the testing molecules in downstream tasks, leading to 100k conformations in total.

Downstream tasks To thoroughly evaluate our proposed pretraining framework, we employ the two widely-adopted 3D molecular property prediction datasets: MD17 (Chmiela et al. 2017) and QM9 (Ramakrishnan et al. 2014), as the downstream tasks. In detail, MD17 contains the simulated dynamical trajectories of 8 small organic molecules, with the recorded energy and force at each frame. We select 9,500/500 frames as the training/validation set of each molecule. We jointly optimize the energy and force predictions by firstly obtaining the energy and deriving the force by $\mathbf{F} = -\nabla_{\mathbf{X}} E$. QM9 labels 12 chemical properties of small molecules with stable 3D structures. We follow the data split in Anderson, Hy, and Kondor (2019) and Satorras, Hoogeboom, and Welling (2021), where the sizes of training, validation, and test sets are 100k, 18k, and 13k, respectively.

Baselines The baseline without any pretraining is termed as *Base*. Several widely-used 2D pretraining tasks are evaluated, including the generative methods *AttrMask* (Hu et al. 2020a), *EdgePred* (Hamilton, Ying, and Leskovec 2017), *GPT-GNN* (Hu et al. 2020b), and the contrastive methods *InfoGraph* (Sun et al. 2020), *GCC* (Qiu et al. 2020), *GraphCL* (You et al. 2020), *JOAO* and its improved version *JOAOv2* (You et al. 2021b). In addition, we also compare with *GraphMVP* (Liu et al. 2021) and *3D Infomax* (Stärk et al. 2022) which simultaneously train 2D- and 3D-GNN models. Notably, different from the original setting in GraphMVP and 3D Infomax, which evaluates the pretrained 2D GNN, we preserve the 3D model for our 3D tasks in the experiments. We further involve *GEM* (Fang et al. 2022) which applies bond length prediction, bond angle prediction, and atom distance prediction as 3D pretraining tasks. We also propose *PosPred*, an extension of 2D AttrMask to 3D, as a competitive 3D baseline which masks the positions of a random subset of atoms with the center of each input molecule, and then reconstructs the masked positions. *For all above model-agnostic methods, we adapt exactly the*

²Boltzmann weight is the statistic weight for each conformer determined by its energy.

Force	Aspirin	Benzene	Ethanol	Malon.	Naph.	Salicylic	Toluene	Uracil	Average
Base	0.3885	0.1861	0.0599	0.1464	0.3310	0.2683	0.1563	0.1323	0.2086
AttrMask	0.3643	0.2277	0.0567	0.1456	0.1773	0.3890	0.1093	0.1560	0.2032
EdgePred	0.4707	0.2036	0.0743	0.1268	0.2310	0.3400	0.1854	0.1933	0.2281
GPT-GNN	0.4278	0.2492	0.0703	0.1484	0.2080	0.3609	0.1541	0.2219	0.2301
InfoGraph	0.6578	0.2743	0.1257	0.2647	0.2860	0.5793	0.3821	0.4238	0.3742
GCC	0.3996	0.2346	0.0662	0.1484	0.2798	0.4263	0.3378	0.2369	0.2662
GraphCL	<u>0.2333</u>	<u>0.1845</u>	<u>0.0503</u>	0.0852	<u>0.0966</u>	<u>0.1587</u>	<u>0.0725</u>	<u>0.1167</u>	<u>0.1247</u>
JOAO	<u>0.3646</u>	<u>0.2331</u>	<u>0.0642</u>	0.1029	<u>0.2017</u>	<u>0.3020</u>	<u>0.1322</u>	0.1683	0.1961
JOAOv2	0.3447	0.2198	0.0568	0.0981	0.1889	0.2753	0.1001	0.1850	0.1836
GraphMVP	0.3198	0.2800	0.0629	<u>0.0788</u>	0.2350	0.2641	0.0903	0.1339	0.1831
3D Infomax	0.4592	0.1914	0.0705	<u>0.1263</u>	0.2642	0.3401	0.2032	0.1836	0.2298
GEM	0.3994	0.2105	0.0871	0.1161	0.1489	0.2344	0.1193	0.1827	0.1873
PosPred	0.3050	0.2023	0.0519	0.0937	0.0971	0.2481	0.0945	0.1270	0.1525
3D-EMGP	0.1560	0.1648	0.0389	0.0737	0.0829	0.1187	0.0619	0.0773	0.0968

Table 1: MAE (lower is better) on MD17 force prediction. All methods share the same backbone as Base.

	α	Δ_{ϵ}	ϵ_{HOMO}	ϵ_{LUMO}	μ	C_{ν}	G	H	R^2	U	U_0	ZPVE
Base	0.070	49.9	28.0	24.3	0.031	0.031	<u>10.1</u>	10.9	0.067	9.7	<u>9.3</u>	1.51
AttrMask	0.072	50.0	31.3	37.8	<u>0.020</u>	0.062	11.2	11.4	0.423	10.8	10.7	1.90
EdgePred	0.086	58.2	37.4	31.9	0.039	0.038	14.5	14.8	0.112	14.2	14.7	1.81
GPT-GNN	0.103	54.1	35.7	28.8	0.039	0.032	12.2	14.8	0.158	24.8	12.0	1.75
InfoGraph	0.099	72.2	48.1	38.1	0.041	0.030	16.5	14.5	0.114	14.9	16.4	1.69
GCC	0.085	57.7	37.7	32.3	0.041	0.034	12.8	14.5	0.104	13.2	13.1	1.66
GraphCL	<u>0.066</u>	45.5	26.8	22.9	0.027	<u>0.028</u>	10.2	9.6	0.095	<u>9.7</u>	9.6	<u>1.42</u>
JOAO	<u>0.068</u>	46.0	28.2	22.8	0.028	<u>0.030</u>	10.5	10.0	0.076	9.9	10.1	1.48
JOAOv2	<u>0.066</u>	45.0	27.8	22.2	0.027	<u>0.028</u>	9.9	<u>9.2</u>	0.087	9.8	9.5	1.43
GraphMVP	0.070	46.9	28.5	26.3	0.031	0.033	11.2	10.4	0.082	10.3	10.2	1.63
3D Infomax	0.075	48.8	29.8	25.7	0.034	0.033	13.0	12.4	0.122	12.5	12.7	1.67
GEM	0.081	52.1	33.8	27.7	0.034	0.035	13.2	13.3	0.089	12.6	13.4	1.73
PosPred	0.067	<u>40.6</u>	<u>25.1</u>	<u>20.9</u>	0.024	0.035	10.9	10.2	0.115	10.3	10.2	1.46
3D-EMGP	0.057	37.1	21.3	18.2	0.020	0.026	9.3	8.7	0.092	8.6	8.6	1.38

Table 2: MAE (lower is better) on QM9. All methods share the same backbone as Base.

same 3D backbone as our method, ensuring fairness. Particularly, we leverage EGNN (Satorras, Hoogeboom, and Welling 2021), a widely adopted equivariant GNN, as our backbone. Details are deferred to Appendix B.2.

Main Results

Table 1 and 2 document the results of all pretraining methods on MD17 and QM9, respectively, where the underlined numbers indicate the previous SOTAs on that task, and the numbers in bold are the best results. We interpret the results by answering the questions as follows. 1. *How does our 3D-EMGP perform in general?* It is observed from both Table 1 and Table 2 that 3D-EMGP achieves the best performance in most cases, and its general effectiveness is better justified by checking the average MAE of the last column in Table 1. Particularly for force prediction, the superiority of 3D-EMGP to other methods is more remarkable (3D-EMGP achieves 0.0969, while the second best GraphCL is 0.1247), probably because the design of our node-level force prediction during pretraining is generalizable to the real force distribution after finetuning. 2. *Are the 3D-aware pretraining*

tasks helpful? Compared with Base, 3D-EMGP consistently delivers meaningful improvement on MD17, and gains better performance on QM9 except for the evaluation on R^2 . We conjecture that the quantity R^2 assessing Electronic spatial extent is hardly recovered by the pretraining dataset, hence incurring negative transfer for all pretraining methods. Interestingly, PosPred usually behaves promisingly on MD17, although its 3D prediction objective is simple. 3. *How do the traditional 2D methods perform on 3D tasks?* Most 2D methods struggle, especially for force prediction on MD17 and tasks on QM9. Taking MD17 as an example, we observe a serious negative transfer phenomenon on several 2D-pretraining baselines like EdgePred, GPT-GNN, and InfoGraph. This may be due to the dissimilarity between the pretraining 2D task and downstream 3D targets (Rosenstein et al. 2005). Meanwhile, GraphCL and GraphMVP achieve better results, because GraphMVP has considered the 3D information, while GraphCL employs the contrastive learning technique to capture the invariant information which may exist in both 2D and 3D areas. Our method 3D-EMGP achieves the best performance in most cases, which verifies

	Proposed Components				Average MAE	
	EFP	INP	RG	Energy	Energy	Force
Base					0.1191	0.2086
Ours	✓	✓	✓	✓	0.0876	0.0968
INP only		✓	✓	✓	0.0974	0.1350
EFP only	✓		✓	✓	0.0905	0.1193
Gaussian	✓	✓		✓	0.0912	0.1060
Distance	✓ ¹	✓			0.0931	0.1292
Direct	✓	✓	✓		0.0914	0.1267

Table 3: Ablation studies on MD17. ¹Denoising on distance.

the effectiveness of our 3D pretraining task.

Ablation Studies

Contribution of each component We provide extensive ablation results on MD17 to show how each component in our model contributes in Table 3. In detail, we study the following aspects. **1.** We inspect the contributions of the node-level task (*i.e.* EFP) and the graph-level task (*i.e.* INP) by comparing our method with its variants without EFP or INP. It is shown that both EFP and INP improve the performance individually, and their combination leads to more precise predictions. **2.** To evaluate the importance of the proposed Riemann-Gaussian (RG) distribution, we relax the distribution in Eq. (9) as the Gaussian distribution $p(\tilde{\mathbf{X}} | \mathbf{X}) = \mathcal{N}(\mathbf{X}, \sigma^2 \mathbf{I})$, violating the doubly E(3)-invariance in Eq. (8). The results suggest that such relaxation causes certain performance detriment. We also compare with a variant which alternatively applies denoising on the E(3)-invariant distance matrix. This surrogate does not fit in the energy framework, and the performance also drops by a margin. This verifies the empirical significance of leveraging Riemann-Gaussian distribution. **3.** We analyze the necessity of the proposed energy-based modeling. Instead of deriving the force as the gradient of the energy model, it is also possible to straightly apply the equivariant output from EGNN as the predicted force signal in the EFP loss in Eq. (7). We name this variant as Direct. Results in Table 3 report that this variant suffers higher MAEs. From an algorithmic point of view, the energy-based strategy is able to better capture the global patterns and therefore lead to preferable performance, by first pooling the embeddings of all nodes as the energy and then computing the gradient of energy as the force.

Performance with different backbones We further apply our method to another two 3D backbones, SchNet (Schütt et al. 2017) and TorchMD-ET (Thölke and De Fabritiis 2021) to evaluate the generalization of our proposed self-supervised tasks. The averaged MAEs of the MD17 force prediction task are reduced by 36.1% and 6.9% for SchNet and TorchMD-ET, respectively. The compelling improvement verifies that our pretraining method is widely applicable and generalizes well to a broad family of 3D backbones consistently.

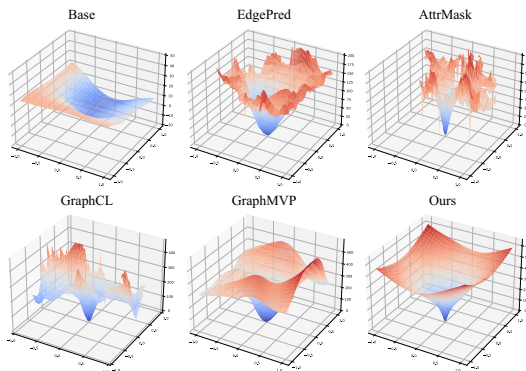


Figure 3: Energy landscape of different pretrained models.

Visualization

To probe the representation space of different pretrained models, we visualize the local energy landscape around a given conformation. To do so, we first fix the pretrained representation model and finetune an energy projection head on MD17 to fit ground-truth energy labels, in order to project the pretrained representations onto the energy surface. Note that there is initially no energy projection head for other methods, and we manually add an MLP on top of their backbone models similar to Eq. (2). After training the energy head, we select a random aspirin conformation \mathbf{X} from MD17 and randomly generate two directions $\mathbf{D}_1, \mathbf{D}_2 \in \mathbb{R}^{3 \times N}$ according to Gaussian distribution. We construct a 2-dimension conformation plane as $\{\tilde{\mathbf{X}}(i, j) | \tilde{\mathbf{X}}(i, j) = \mathbf{X} + i\mathbf{D}_1 + j\mathbf{D}_2\}$. For each point by varying the values of i and j , we calculate its output energy by $E_{i,j} = E(\varphi_{\text{EGN}}(\tilde{\mathbf{X}}(i, j)))$, where E, φ_{EGN} denote the energy projection head and the pretrained model, respectively. Fig. 3 plots the energy landscape $(i, j, E_{i,j})$ for several compared approaches and our 3D-EMGP. We interestingly find that the landscape by our method converges towards the original conformation smoothly and decreasingly, which implies the observed conformation corresponds to a metastable state with locally-lowest energy on the projected conformation plane. However, the 2D-based pretrained models such as EdgePred, AttrMask, and GraphCL deliver rugged landscapes. We speculate the reason is that their knowledge acquired from the pretraining process does not comply with the underlying energy distribution. The Base method outputs a flat surface, as it is less knowledgeable by solely learning from the small data.

Conclusion

In this work, we propose a general self-supervised pretraining framework for 3D tasks on molecules. It consists of a node-level Equivariant Force Prediction (EFP) and a graph-level Invariant Noise-scale Prediction (INP) task to jointly extract the geometric information from a large-scale 3D molecular dataset. Experiments on MD17 and QM9 showcase the superiority of our method to conventional 2D counterparts. Necessary ablations, visualizations, and analyses are also provided to support the validity of our design.

Acknowledgments

This work is jointly supported by the National Natural Science Foundation of China (No.61925601, No.62006137); Guoqiang Research Institute General Project, Tsinghua University (No. 2021GQG1012); Beijing Academy of Artificial Intelligence; Huawei Noah's Ark Lab; Beijing Outstanding Young Scientist Program (No. BJJWZYJH012019100020098); Tencent AI Lab. Rhino-Bird Visiting Scholars Program (VS2022TEG001).

References

- Anderson, B.; Hy, T. S.; and Kondor, R. 2019. Cormorant: Covariant molecular neural networks. *Advances in neural information processing systems*, 32.
- Axelrod, S.; and Gomez-Bombarelli, R. 2022. GEOM, energy-annotated molecular conformations for property prediction and molecular generation. *Scientific Data*, 9(1): 1–14.
- Boltzmann, L. 1868. Studien uber das Gleichgewicht der lebenden Kraft. *Wissenschaftliche Abhandlungen*, 1: 49–96.
- Chmiela, S.; Tkatchenko, A.; Sauceda, H. E.; Poltavsky, I.; Schütt, K. T.; and Müller, K.-R. 2017. Machine learning of accurate energy-conserving molecular force fields. *Science advances*, 3(5): e1603015.
- Fang, X.; Liu, L.; Lei, J.; He, D.; Zhang, S.; Zhou, J.; Wang, F.; Wu, H.; and Wang, H. 2022. Geometry-enhanced molecular representation learning for property prediction. *Nature Machine Intelligence*, 1–8.
- Finzi, M.; Stanton, S.; Izmailov, P.; and Wilson, A. G. 2020. Generalizing Convolutional Neural Networks for Equivariance to Lie Groups on Arbitrary Continuous Data. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, 3165–3176. PMLR.
- Fuchs, F.; Worrall, D. E.; Fischer, V.; and Welling, M. 2020. SE(3)-Transformers: 3D Roto-Translation Equivariant Attention Networks. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural message passing for quantum chemistry. In *International conference on machine learning*, 1263–1272. PMLR.
- Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30.
- Han, J.; Rong, Y.; Xu, T.; and Huang, W. 2022. Geometrically Equivariant Graph Neural Networks: A Survey. *arXiv preprint arXiv:2202.07230*.
- Hu, W.; Liu, B.; Gomes, J.; Zitnik, M.; Liang, P.; Pande, V. S.; and Leskovec, J. 2020a. Strategies for Pre-training Graph Neural Networks. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Hu, Z.; Dong, Y.; Wang, K.; Chang, K.; and Sun, Y. 2020b. GPT-GNN: Generative Pre-Training of Graph Neural Networks. In Gupta, R.; Liu, Y.; Tang, J.; and Prakash, B. A., eds., *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020*, 1857–1867. ACM.
- Hutchinson, M. J.; Le Lan, C.; Zaidi, S.; Dupont, E.; Teh, Y. W.; and Kim, H. 2021. Lietransformer: Equivariant self-attention for lie groups. In *International Conference on Machine Learning*, 4533–4543. PMLR.
- Kearnes, S.; McCloskey, K.; Berndl, M.; Pande, V.; and Riley, P. 2016. Molecular graph convolutions: moving beyond fingerprints. *Journal of computer-aided molecular design*, 30(8): 595–608.
- Kipf, T. N.; and Welling, M. 2016. Variational Graph Auto-Encoders. *NIPS Workshop on Bayesian Deep Learning*.
- Liu, S.; Wang, H.; Liu, W.; Lasenby, J.; Guo, H.; and Tang, J. 2021. Pre-training Molecular Graph Representation with 3D Geometry. In *International Conference on Learning Representations*.
- Liu, Y.; Wang, L.; Liu, M.; Lin, Y.; Zhang, X.; Oztekin, B.; and Ji, S. 2022. Spherical Message Passing for 3D Molecular Graphs. In *International Conference on Learning Representations*.
- Luo, S.; and Hu, W. 2020. Differentiable manifold reconstruction for point cloud denoising. In *Proceedings of the 28th ACM international conference on multimedia*, 1330–1338.
- Luo, S.; Shi, C.; Xu, M.; and Tang, J. 2021. Predicting Molecular Conformation via Dynamic Graph Score Matching. *Advances in Neural Information Processing Systems*, 34.
- Qiu, J.; Chen, Q.; Dong, Y.; Zhang, J.; Yang, H.; Ding, M.; Wang, K.; and Tang, J. 2020. Gcc: Graph contrastive coding for graph neural network pre-training. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1150–1160.
- Ramakrishnan, R.; Dral, P. O.; Rupp, M.; and Von Lilienfeld, O. A. 2014. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1(1): 1–7.
- Rong, Y.; Bian, Y.; Xu, T.; Xie, W.; Wei, Y.; Huang, W.; and Huang, J. 2020. Self-Supervised Graph Transformer on Large-Scale Molecular Data. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Rosenstein, M. T.; Marx, Z.; Kaelbling, L. P.; and Dietterich, T. G. 2005. To transfer or not to transfer. In *In NIPS'05 Workshop, Inductive Transfer: 10 Years Later*.
- Said, S.; Bombrun, L.; Berthoumieu, Y.; and Manton, J. H. 2017. Riemannian Gaussian distributions on the space of symmetric positive definite matrices. *IEEE Transactions on Information Theory*, 63(4): 2153–2170.
- Satorras, V. G.; Hoogeboom, E.; and Welling, M. 2021. E(n) equivariant graph neural networks. In *International Conference on Machine Learning*, 9323–9332. PMLR.

- Schlick, T. 2010. *Molecular modeling and simulation: an interdisciplinary guide*, volume 2. Springer.
- Schütt, K.; Kindermans, P.-J.; Sauceda Felix, H. E.; Chmiela, S.; Tkatchenko, A.; and Müller, K.-R. 2017. Schnet: A continuous-filter convolutional neural network for modeling quantum interactions. *Advances in neural information processing systems*, 30.
- Shi, C.; Luo, S.; Xu, M.; and Tang, J. 2021. Learning gradient fields for molecular conformation generation. In *International Conference on Machine Learning*, 9558–9568. PMLR.
- Song, Y.; and Ermon, S. 2020. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33: 12438–12448.
- Stärk, H.; Beaini, D.; Corso, G.; Tossou, P.; Dallago, C.; Günnemann, S.; and Liò, P. 2022. 3d infomax improves gnns for molecular property prediction. In *International Conference on Machine Learning*, 20479–20502. PMLR.
- Sun, F.; Hoffmann, J.; Verma, V.; and Tang, J. 2020. InfoGraph: Unsupervised and Semi-supervised Graph-Level Representation Learning via Mutual Information Maximization. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Sun, M.; Xing, J.; Wang, H.; Chen, B.; and Zhou, J. 2021. MoCL: Data-Driven Molecular Fingerprint via Knowledge-Aware Contrastive Learning from Molecular Graph. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, KDD '21*, 3585–3594. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383325.
- Thölke, P.; and De Fabritiis, G. 2021. Equivariant transformers for neural network based molecular potentials. In *International Conference on Learning Representations*.
- Thomas, N.; Smidt, T.; Kearnes, S.; Yang, L.; Li, L.; Kohlhoff, K.; and Riley, P. 2018. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219*.
- Velickovic, P.; Fedus, W.; Hamilton, W. L.; Liò, P.; Bengio, Y.; and Hjelm, R. D. 2019. Deep Graph Infomax. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Vincent, P. 2011. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7): 1661–1674.
- Wallach, I.; Dzamba, M.; and Heifets, A. 2015. AtomNet: a deep convolutional neural network for bioactivity prediction in structure-based drug discovery. *arXiv preprint arXiv:1510.02855*.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful are Graph Neural Networks? In *International Conference on Learning Representations*.
- You, Y.; Chen, T.; Shen, Y.; and Wang, Z. 2021a. Graph Contrastive Learning Automated. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, 12121–12132. PMLR.
- You, Y.; Chen, T.; Shen, Y.; and Wang, Z. 2021b. Graph contrastive learning automated. In *International Conference on Machine Learning*, 12121–12132. PMLR.
- You, Y.; Chen, T.; Sui, Y.; Chen, T.; Wang, Z.; and Shen, Y. 2020. Graph Contrastive Learning with Augmentations. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Zheng, L.; Fan, J.; and Mu, Y. 2019. Onionnet: a multiple-layer intermolecular-contact-based convolutional neural network for protein–ligand binding affinity prediction. *ACS omega*, 4(14): 15956–15965.
- Zhou, G.; Gao, Z.; Ding, Q.; Zheng, H.; Xu, H.; Wei, Z.; Zhang, L.; and Ke, G. 2023. Uni-Mol: A Universal 3D Molecular Representation Learning Framework. In *The Eleventh International Conference on Learning Representations*.
- Zhu, J.; Xia, Y.; Wu, L.; Xie, S.; Qin, T.; Zhou, W.; Li, H.; and Liu, T.-Y. 2022. Unified 2d and 3d pre-training of molecular representations. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2626–2636.