
Molecular Contrastive Pretraining with Collaborative Featurizations

Yanqiao Zhu^{1*} Dingshuo Chen^{2,3*} Yuanqi Du^{4*} Yingze Wang⁵ Qiang Liu^{2,3} Shu Wu^{2,3†}

¹Department of Computer Science, University of California, Los Angeles

²Center for Research on Intelligent Perception and Computing
Institute of Automation, Chinese Academy of Sciences

³School of Artificial Intelligence, University of Chinese Academy of Sciences

⁴Department of Computer Science, Cornell University

⁵College of Chemistry and Molecular Engineering, Peking University

yzhu@cs.ucla.edu dingshuo.chen@cripac.ia.ac.cn ericwangyz@pku.edu.cn
yd392@cornell.edu {qiang.liu, shu.wu}@nlpr.ia.ac.cn

Abstract

Molecular pretraining, which learns molecular representations over massive unlabeled data, has become a prominent paradigm to solve a variety of tasks in computational chemistry and drug discovery. Recently, prosperous progress has been made in molecular pretraining with different molecular featurizations, including 1D SMILES strings, 2D graphs, and 3D geometries. However, the role of molecular featurizations with their corresponding neural architectures in molecular pretraining remains largely unexamined. In this paper, through two case studies—chirality classification and aromatic ring counting—we first demonstrate that different featurization techniques convey chemical information differently. In light of this observation, we propose a simple and effective MOlecular pretraining framework with COllaborative featurizations (MOCO). MOCO comprehensively leverages multiple featurizations that complement each other and outperforms existing state-of-the-art models that solely relies on one or two featurizations on a wide range of molecular property prediction tasks.

1 Introduction

Molecular representation learning, which automates the process of feature learning for molecules, is fast driving the development of computational chemistry and drug discovery. It has been recognized as crucial for a variety of downstream tasks, spanning from molecular property prediction to molecule design [1, 2]. Deep neural models, on the other hand, rely on a substantial amount of labeled data, which require expensive wet lab experiments in chemical domains. With insufficient annotated data, deep models easily overfit to such small training data and tend to learn spurious correlations [3].

In recent years, self-supervised pretraining has emerged as a promising strategy to alleviate the label scarcity problem and improve model robustness [4]. A typical framework pretrains the encoder model with training objectives over large-scale unlabeled datasets and then fine-tunes the learned model on labeled downstream tasks. Motivated by its success, many molecular pretraining models have been developed [5–12]. To capture chemical semantics of molecules, these models design several pretraining strategies based on different *molecular featurizations*, which translate chemical information into representations that can be recognized by machine learning algorithms. For example, early models [5, 6] propose to leverage masked language modeling [13] to pretrain Simplified Molecular-Input Line-Entry System (SMILES) strings [14], while others study contrastive learning on 2D graphs [7–9] or 3D conformations [10]. Some recent studies further propose to enrich 2D-topology-based pretraining with 3D geometry information [11, 12].

*These authors made equal contribution to this research.

†To whom correspondence should be addressed.




Figure 1: The proposed MOCO model. MOCO obtains four molecule featurizations with appropriate encoders. After that, an attention network is employed to aggregate each view embedding and compute a final embedding. The model is trained using a contrastive objective that maximizes the consistency between view embeddings and the final embedding.

Despite encouraging progress, prior studies tend to emphasize on pretraining on molecular graphs and overlook the impact of other molecular featurizations with their corresponding neural encoders, which represent chemical information in different ways. Consider SMILES strings as an example. It explicitly represents informative structures in special characters such as branches, rings, and chirality [15], which are difficult to learn in graph-based representations [16]. Moreover, the utility of different featurizations may vary across downstream tasks. Therefore, most previous models relying on only one or two featurizations might achieve sub-optimal performance across various downstream tasks. For example, 2D topology is important for many drug-related properties such as toxicity, while 3D geometry arguably determines properties related to quantum mechanics, such as single-point energy, atomic forces, or dipole moments [17, 18]. Therefore, it is natural to ask whether we can enjoy the benefits from multiple molecular featurizations and take the relative utilities of different featurizations into consideration during fine-tuning on downstream tasks.

In this work, we first revisit four commonly used featurizations techniques: (a) 2D topology graphs, (b) 3D geometry graphs, (c) Morgan fingerprints, and (d) SMILES strings. We leverage four accompanying neural encoders with proper inductive bias and conduct two case studies, classifying tetrahedral chiral centers and counting aromatic rings, both of which are informative chemical descriptors, on representations obtained on different featurization techniques. The results show there is no one single featurization that dominates the others, indicating that different featurizations encode chemical semantics of molecules in different ways.

In light of this observation, we then propose a simple and effective Molecular pretraining framework with Collaborative featurizations to comprehensively leverage every featurization during both pretraining and fine-tuning, which we term MOCO for brevity. Its graphical illustration is shown in Figure 1. The core idea of MOCO is to dynamically adjust the contribution of each featurization through an attention network, which *selectively* extracts information from each collaborative “view” of the raw molecular data. Besides, we design a novel multiview contrastive pretraining strategy, which trains the model by maximizing the consistency among different views in a self-supervised manner. Contrary to previous studies [11, 12] that only consider 2D graph structures during fine-tuning, our MOCO utilizes multiple featurizations in *both* pretraining and fine-tuning stages and further allows interpretation analysis of different downstream tasks for domain scientists. Note that our proposed MOCO framework is generic, allowing for seamless integration of off-the-shelf neural architectures. To the best of our knowledge, this is the first work that studies how various featurization techniques should be utilized for molecular pretraining and downstream tasks.

We evaluate the effectiveness of our MOCO model on widely-used benchmark datasets including MoleculeNet [19] and QM9 [20] that cover a wide range of molecular property prediction tasks. The results reveal that MOCO consistently improves non-pretraining baselines without negative transfer and outperforms existing state-of-the-art molecular pretraining models, achieving a 1.1% absolute improvement in terms of average ROC-AUC. Furthermore, the learned model weights of molecular featurizations for different end tasks are well aligned with prior chemical knowledge. We also suggest a series of guidelines on choosing effective featurization techniques for molecular representations.

The main contributions of this work are three-fold:

- We explore the featurization spaces of molecules with appropriate neural encoders and highlight the importance of incorporating different featurizations for molecular pretraining.
- We propose a novel molecular contrastive pretraining framework that adaptively integrates information from multiple collaborative featurizations during both pretraining and fine-tuning stages and provides interpretability for downstream molecular property prediction tasks.

- Extensive experiments conducted on public benchmark datasets validate the effectiveness of our proposed model. MOCO is able to achieve the state-of-the-art across various downstream datasets without negative transfer.

2 Preliminaries

2.1 A Brief Recapitulation of Molecular Featurization Techniques

Molecular featurizations translate chemical information of molecules into representations that can be understood by machine learning algorithms. Concretely, we consider the following molecular featurizations covering string-, graph-, scalar-, and vector-based representations for 1D/2D molecules and 3D structures, which are popular in literature [21, 22]:

- **2D topology graphs** model atoms and bonds as nodes and edges respectively. It is arguably a common technique, especially for capturing substructure information by means of graph topology.
- **3D geometry graphs** incorporate atomic coordinates (conformations) in their representations and are able to depict how atoms are positioned relative to each other in the 3D space. We consider conformers in an equilibrium state, corresponding to the minima in a potential energy surface.
- **Morgan fingerprints** [23, 24] encode molecules in fixed-length binary strings, with bits indicating presence or absence of specific substructures. They represent each atom according to a set of atomic invariants and iteratively update these features among neighboring atoms using a hash function.
- **SMILES strings** are a concise technique that represents chemical structures in a linear notation using ASCII characters, with explicitly depicting information about atoms, bonds, rings, connectivity, aromaticity, and stereochemistry.

2.2 Learning Representations with Different Featurizations

Next, we introduce four encoders with different inductive bias to capture the intrinsic information with each featurization. Here we only discuss the high-level design of each encoder; please refer to Supporting Information: Section 1 for detailed implementations of each encoder.

Notations. Each molecule can be represented as an undirected graph, where nodes are atoms and edges describe inter-atomic bonds. Formally, each graph is denoted as $\mathcal{G} = (\mathbf{A}, \mathbf{R}, \mathbf{X}, \mathbf{E})$, where $\mathbf{A} \in \{0, 1\}^{N \times N}$ is the adjacency matrix of N nodes, $\mathbf{R} \in \mathbb{R}^{N \times 3}$ is the 3D position matrix, $\mathbf{X} \in \mathbb{R}^{N \times K}$ is the matrix of atom attributes of K dimension, and $\mathbf{E} \in \mathbb{R}^{N \times N \times E}$ is the tensor for bond attributes of E dimension. Additionally, each molecule is attached with a binary fingerprint vector $\mathbf{f} \in \{0, 1\}^F$ of length F and a SMILES string $\mathbf{S} = [s_j]_{j=1}^S$ of length S . In what follows, the subscript i is used to index the i -th molecule.

Embedding 2D graphs. To capture the 2D topological information, we employ a widely-used Graph Isomorphism Network (GIN) model [25] denoted by f_{2D} , which receives as input the graph adjacency matrix and attributes of atoms and bonds, and produces the embedding vector $\mathbf{z}_i^{2D} \in \mathbb{R}^D$:

$$\mathbf{z}_i^{2D} = f_{2D}(\mathbf{X}_i, \mathbf{E}_i, \mathbf{A}_i). \quad (1)$$

Embedding 3D graphs. To model additional spatial coordinates associated with atoms, we leverage SchNet [26] as the backbone, which models message passing as continuous-filter convolutions and is able to preserve rotational invariance for energy predictions. We denote its encoding function as f_{3D} which takes atom features and positions as input and produces the 3D embedding $\mathbf{z}_i^{3D} \in \mathbb{R}^D$:

$$\mathbf{z}_i^{3D} = f_{3D}(\mathbf{X}_i, \mathbf{R}_i). \quad (2)$$

Embedding molecular fingerprints. Since there is a lack of proper neural encoders for fingerprints, we propose an attention-based network to model interactions of feature fields in fingerprint vectors, which considers the discrete and extremely sparse nature of fingerprints. Specifically, we first transform all F feature fields into a dense embedding matrix $\mathbf{F}_i \in \mathbb{R}^{F \times D_F}$ via embedding lookup.

Table 1: Results of two case studies with different featurizations: chirality classification and aromatic ring count regression.

Target	2D	3D	SM	FP
Chirality (AP, \uparrow)	0.4952	0.4959	0.5505	0.5246
#Rings (MAE, \downarrow)	0.1949	0.2021	0.3077	0.2590

Then, we use a multihead self-attention network f_{FP} [27] to model the interaction among those feature fields, resulting in an embedding matrix $\widehat{\mathbf{Z}}_i^{FP} \in \mathbb{R}^{F \times D_F}$. Following that, we perform sum pooling and use a linear model f_{LIN} to obtain the final fingerprint embedding $\mathbf{z}_i^{FP} \in \mathbb{R}^D$:

$$\widehat{\mathbf{Z}}_i^{FP} = f_{FP}(\mathbf{F}_i), \quad \mathbf{z}_i^{FP} = f_{LIN} \left(\sum_{d=1}^{D_F} \widehat{\mathbf{Z}}_{i,d}^{FP} \right). \quad (3)$$

Embedding SMILES strings. To encode SMILES strings, we use a pretrained RoBERTa [28] as the backbone model. As SMILES strings do not possess consecutive relationships, the RoBERTa model is pretrained using the masked language model as the only objective, unlike conventional natural language models [29]. After that, in order to reduce the computational burden, we freeze the RoBERTa encoder (denoted by f_{SM}) in our model and employ an additional learnable MultiLayer Perceptron (MLP) on the representation $\mathbf{s}_i \in \mathbb{R}^{D_S}$ to get the final embedding $\mathbf{z}_i^{SM} \in \mathbb{R}^D$:

$$\mathbf{s}_i = f_{SM}(\mathbf{S}_i), \quad \mathbf{z}_i^{SM} = f_{MLP}(\mathbf{s}_i). \quad (4)$$

2.3 Case Studies

We present two case studies—chirality classification and aromatic ring counting—to demonstrate that the representation ability of each featurization with the corresponding neural encoder is different. For chirality classification, we randomly select 10K molecules with one chirality center from GEOM-Drugs [30] and test whether the representations obtained using the four featurizations can classify tetrahedral chiral centers as R/S. For aromatic ring counting, we randomly draw another 10K molecules and test whether these models can recognize the number of aromatic rings of each molecule. Note that both chirality properties and ring counts are informative chemical descriptors [31] and can be easily computed with existing implementations such as RDKit [32].


We report classification and regression performance in Average Precision (AP) and Mean Absolute Error (MAE) respectively. The results are summarized in Table 1. It is seen from the table that no single featurization performs the best on all targets and four representations contain complementary information to each other, suggesting us to leverage multiple featurizations for molecular pretraining.

3 Molecular Pretraining with Collaborative Featurizations

As with generic self-supervised learning pipelines, the MOCO framework is divided into two stages, pretraining and fine-tuning. In the first stage, given an unlabeled dataset, we train an encoding function that learns representations with the four featurization techniques. In the subsequent fine-tuning phase, we take the weights of the encoders from the pretrained model and tune the model on molecules with annotations of particular properties in a supervised fashion.

We next introduce the MOCO pretraining framework in detail. We first use obtain four “view” representations based on the aforementioned four featurizations. Then, we integrate these four embeddings to compute a final representation for each molecule through an attention network. Finally, we pretrain the whole model using a contrastive objective.

Figure 2: (a) Chirality: even if two graphs are isomorphic, they can have two distinct stereochemistry structures. (b) The aromatic ring is an important functional group.



3.1 Representation Aggregation from Multiple Featurizations

Since each featurization technique reflects the molecule from one certain aspect, we take weighted average of every view embedding to obtain a comprehensive final representation:

$$\mathbf{z}_i = \sum_{m \in \mathcal{M}} \alpha^m \mathbf{z}_i^m, \quad (5)$$

where $\mathcal{M} = \{2D, 3D, FP, SM\}$ is the set of all views. We leverage an attention network [33] that learns to adjust the contribution of each view. Formally, the attention coefficient α^m denoting the contribution of the m -th view is computed by:

$$\alpha^m = \frac{\exp(w^m)}{\sum_{m' \in \mathcal{M}} \exp(w^{m'})}, \quad w^m = \frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \mathbf{q}^\top \cdot \tanh \left(\mathbf{W} \frac{\mathbf{z}_i^m}{\|\mathbf{z}_i^m\|_2} + \mathbf{b} \right), \quad (6)$$

where $\mathbf{q}, \mathbf{b} \in \mathbb{R}^D$, $\mathbf{W} \in \mathbb{R}^{D \times D}$ are trainable parameters in the attention network, and \mathcal{B} denotes the set of molecules in the current training batch. Note that we perform ℓ_2 normalization on all embeddings to regularize the scale across different views when computing the attention scores.

3.2 Contrastive Objectives for Pretraining

Finally, we train the model using a contrastive objective by aligning the aggregated embedding with all view-specific embeddings. Particularly, for one molecule i , we designate its four view embeddings \mathbf{z}_i^m as the anchors and the aggregated embeddings \mathbf{z}_i as the positive instance. Other aggregated embeddings $\{\mathbf{z}_j\}_{i \neq j}$ in the same batch are then chosen as the negative samples. Following prior studies [8, 34–38], we leverage the Information Noise Contrastive Estimation (InfoNCE) objective, which can be formally written as:

$$\mathcal{L} = \frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \left[\frac{1}{|\mathcal{M}|} \sum_{m \in \mathcal{M}} -\log \frac{\exp(\theta(\mathbf{z}_i^m, \mathbf{z}_i)/\tau)}{\sum_{j \in \mathcal{B}} \exp(\theta(\mathbf{z}_i^m, \mathbf{z}_j)/\tau)} \right], \quad (7)$$

where the critic function θ computes the likelihood scores of contrastive pairs and the hyperparameter τ adjusts the dynamic range of the likelihood scores of contrastive pairs. Specifically, the critic function θ performs non-linear transformation via an MLP function g [34] and then measures their cosine similarity:

$$\theta(\mathbf{x}, \mathbf{y}) = \frac{g(\mathbf{x})^\top g(\mathbf{y})}{\|g(\mathbf{x})\|_2 \|g(\mathbf{y})\|_2}. \quad (8)$$

After pretraining the model with the self-supervised objective function \mathcal{L} , we fine-tune the model weights of view encoders along with the attentive representation aggregation module with the supervision of downstream tasks at a smaller learning rate.

4 Experiments

In this section, we present empirical evaluation of our proposed work. Specifically, the experiments aim to investigate the following three key questions.

- **RQ1 (Overall performance).** Is the proposed MOCO able to improve non-pretraining baselines and outperform state-of-the-arts on molecular property prediction tasks?
- **RQ2 (Interpretation).** Are the learned attention weights of molecular featurizations on different downstream tasks consistent with chemical knowledge?
- **RQ3 (Ablation studies).** How do the representation aggregation module and the fine-tuning strategy affect the model performance?

In the following, we first summarize experimental setup and proceed to results and analysis.

4.1 Experimental Configurations

Datasets. We closely follow the experimental setup of GraphMVP [12] for fair comparison. Specifically, we pretrain the model using the GEOM-Drugs dataset [30] containing both 2D and 3D information. For fine-tuning, we choose a variety datasets extracted from MoleculeNet [19], ChEMBL [39],

and CEP [40], that cover a wide range of applications, including physiological, biological, and pharmaceutical tasks, and QM9 [20] that focuses on quantum property prediction. These downstream tasks include 8 binary classification and 12 regression tasks. For those datasets for fine-tuning, we follow OGB [41] that uses scaffolds to split training/test/validation subsets with a split ratio of 80%/10%/10%. For detailed description, we refer readers of interest to Supporting Information: Section 2.

Baselines. For comprehensive comparison, we select the following two groups of SSL methods as primary baselines in our experiments.

- Generic graph SSL models: GraphSAGE [42], InfoGraph [43], GPT-GNN [44], AttrMask, ContextPred [7], GraphLoG [9], GraphCL [8], JOAO [45], and GraphMAE [46].
- Molecular SSL models: GROVER-Contextual (GROVER-C), GROVER-Motif (GROVER-M) [47], UnifiedMol [48], GraphMVP³ [12] and BET [49].

To ensure a fair comparison, all the above SSL baselines are pre-trained on GEOM-Drugs dataset with identical atomic and bond features and evaluated based on the same settings following Hu et al. [7]. However, some hand-crafted features are closely related to the design of UnifiedMol [48], e.g. subgraph isomorphism, thus we keep these features during reproduction. We also report performance with a randomly initialized model as the non-pretraining baseline.

Implementation details. In the GEOM-Drugs dataset, since the original full set is too large (containing 317K molecules with over 9M conformations), we randomly select 50K molecules as the pretraining dataset. For each molecule, we select to use its top-5 conformers of the lowest energy in virtue of their sufficient geometry information. Since molecules in the fine-tuning datasets do not have 3D information available, we use ETKDG [50] in RDkit [32] to compute molecular conformations. For both pretraining and fine-tuning datasets, we use RDkit to generate 1024-bit molecular fingerprints with radius $R = 2$, which is roughly equivalent to the ECFP4 scheme [51]. We would like to emphasize that all dataset preprocessing and graph encoder architectures are kept in line with GraphMVP [12] to ensure fair comparison. Readers of interest may refer to Supporting Information: Section 3 for implementation details regarding software/hardware platforms, model training, and hyperparameter specifications.

Evaluation protocols. For classification tasks, we report the performance in terms of the Area Under the ROC-Curve (ROC-AUC), where higher values indicate better performance. For quantum property and other non-quantum regression tasks, we measure the performance in Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) respectively, where lower values are better. We repeat every experiment on three seeds with scaffold splitting and report the averaged performance with standard deviation, following previous work [12].

4.2 Main Results on Molecular Property Prediction

The performance of molecular property prediction tasks is summarized in Table 2. It can be found that our MOCO shows strong empirical performance across all eight low-data downstream datasets, delivering seven out of eight state-of-the-art results and acquiring a 1.1% absolute improvement on average. The outstanding results validate the superiority of our proposed model.

We make other observations as follows. Firstly, MOCO obtains more accurate and stabler predictions compared to the randomly initialized baseline, indicating that our pretraining framework can transfer the knowledge from large, unannotated datasets to smaller downstream datasets without negative transfer. Secondly, previous work has already achieved pretty high performance. For example, the current state-of-the-art UnifiedMol [48] only obtains a 0.05% absolute improvement over its best baseline BET [49] in terms of average ROC-AUC. Our work pushes that boundary without extensive hyperparameter tuning, with an absolute improvement of up to 1.4% over UnifiedMol in terms of average ROC-AUC. Lastly, it is worth mentioning that, the non-pretraining baseline even achieves better performance than some graph-based pretraining models. On some challenging datasets (e.g., Tox21, MUV, and ToxCast), it even achieves the second to best performance. This once more demonstrates the effectiveness of leveraging multiple featurization techniques.

³In our experiments, we do not include its two variants GraphMVP-G and GraphMVP-C since they are essentially two ensemble models that combine AttrMask and ContextPred [7] respectively.

Table 2: Results for eight molecule property prediction tasks in terms of ROC-AUC (%). We highlight the best- and the second-best performing results in **boldface** and underlined, respectively.

Pretraining	BBBP	Tox21	ToxCast	SIDER	ClinTox	MUV	HIV	BACE	Avg.
—	71.0 \pm 0.5	<u>75.9\pm0.3</u>	<u>64.7\pm2.3</u>	57.7 \pm 3.1	71.5 \pm 5.3	<u>77.7\pm1.0</u>	75.9 \pm 0.7	71.5 \pm 2.7	70.63
GraphSAGE [42]	64.5 \pm 3.1	74.5 \pm 0.4	60.8 \pm 0.5	56.7 \pm 0.1	55.8 \pm 6.2	73.3 \pm 1.6	75.1 \pm 0.8	64.6 \pm 4.7	65.64
AttrMask [7]	70.2 \pm 0.5	74.2 \pm 0.8	62.5 \pm 0.4	60.4 \pm 0.6	68.6 \pm 9.6	73.9 \pm 1.3	74.3 \pm 1.3	77.2 \pm 1.4	70.16
GPT-GNN [44]	64.5 \pm 1.1	75.3 \pm 0.5	62.2 \pm 0.1	57.5 \pm 4.2	57.8 \pm 3.1	76.1 \pm 2.3	75.1 \pm 0.2	77.6 \pm 0.5	68.27
InfoGraph [43]	69.2 \pm 0.8	73.0 \pm 0.7	62.0 \pm 0.3	59.2 \pm 0.2	75.1 \pm 5.0	74.0 \pm 1.5	74.5 \pm 1.8	73.9 \pm 2.5	70.10
ContextPred [7]	<u>71.2\pm0.9</u>	73.3 \pm 0.5	62.8 \pm 0.3	59.3 \pm 1.4	73.7 \pm 4.0	72.5 \pm 2.2	75.8 \pm 1.1	78.6 \pm 1.4	70.89
GraphLoG [9]	67.8 \pm 1.7	73.0 \pm 0.3	62.2 \pm 0.4	57.4 \pm 2.3	62.0 \pm 1.8	73.1 \pm 1.7	73.4 \pm 0.6	78.8 \pm 0.7	68.47
GROVER-C [47]	70.3 \pm 1.6	75.2 \pm 0.3	62.6 \pm 0.3	58.4 \pm 0.6	59.9 \pm 8.2	72.3 \pm 0.9	75.9 \pm 0.9	79.2 \pm 0.3	69.21
GROVER-M [47]	66.4 \pm 3.4	73.2 \pm 0.8	62.6 \pm 0.5	60.6 \pm 1.1	77.8 \pm 2.0	73.3 \pm 2.0	73.8 \pm 1.4	73.4 \pm 4.0	70.14
GraphCL [8]	67.5 \pm 3.3	75.0 \pm 0.3	62.8 \pm 0.2	60.1 \pm 1.3	78.9 \pm 4.2	77.1 \pm 1.0	75.0 \pm 0.4	68.7 \pm 7.8	70.64
JOAO [45]	66.0 \pm 0.6	74.4 \pm 0.7	62.7 \pm 0.6	60.7 \pm 1.0	66.3 \pm 3.9	77.0 \pm 2.2	76.6 \pm 0.5	72.9 \pm 2.0	69.57
GraphMVP [12]	68.5 \pm 0.2	74.5 \pm 0.4	62.7 \pm 0.1	62.3\pm1.6	79.0 \pm 2.5	75.0 \pm 1.4	74.8 \pm 1.4	76.8 \pm 1.1	71.69
GraphMAE [46]	70.9 \pm 0.9	75.0 \pm 0.4	64.1 \pm 0.1	59.9 \pm 0.5	<u>81.5\pm2.8</u>	76.9 \pm 2.6	76.7 \pm 0.9	<u>81.4\pm1.4</u>	73.31
UnifiedMol [48]	70.4 \pm 2.2	75.2 \pm 0.4	63.3 \pm 0.3	62.2 \pm 1.1	81.3 \pm 4.9	76.9 \pm 1.5	77.5 \pm 1.1	80.3 \pm 2.2	73.38
BET [49]	68.3 \pm 1.2	76.1 \pm 0.3	62.2 \pm 0.4	<u>63.4\pm0.9</u>	81.1 \pm 6.3	76.2 \pm 1.1	78.5\pm0.8	80.9 \pm 2.4	73.34
MOCO	71.6\pm1.0	76.7\pm0.4	64.9\pm0.8	61.2 \pm 0.6	81.6\pm3.7	78.5\pm1.4	<u>78.3\pm0.4</u>	82.6\pm0.3	74.41

4.3 Interpretation and Analysis

In order to analyze the correlation between tasks and featurization techniques, we visualize the attention weights α learned on different downstream tasks in Figure 3. Note that most of the datasets in MoleculeNet [19] are ADMET property prediction tasks: chemical Absorption (A), Distribution (D), Metabolism (M), Excretion (E), and Toxicity (T), and we thus group the eight end tasks according to their prediction targets in the following analysis.

In general, we can interpret from the visualization that *2D-based features are more significant than 3D-based features in the studied tasks*, which is well aligned with chemical knowledge. We provide detailed analysis as follows:

- In Tox21, ClinTox, SIDER, and ToxCast, we find that 2D graphs play the most important role. These four datasets are related to toxicity (or side effects). Although it is a very complex biological issue to explain, such properties can still be partially deduced from certain functional groups patterns contained in 2D graphs. Actually, medicinal chemists have developed such a database to provide them with necessary alerts of potential side effects in drug design [52].
- BBBP, which measures blood-brain barrier permeability, is mostly dominated by the following properties: liposolubility/water-solubility, molecular weight, and interaction between molecules and transporter proteins. Similarly, these properties can also be inferred from 2D topology, such as molecules with too many hydrogen bond acceptors/donors are unlikely to break the blood-brain barrier due to poor liposolubility [53].
- On BACE and MUV we see 2D graphs and SMILES strings contribute most. These two datasets are about predicting protein-ligand binding activities, which are theoretically relevant to 3D conformations. However, it is still an open question that whether the conformation sampling methods can produce conformations that resemble bioactive conformations, which provide the key information for protein-ligand binding. Nevertheless, in each of these tasks, the target protein is fixed so that bioactivity can be partially deduced from 2D structures, which is supported by the success of fragment-based Quantitative Structure-Activity Relationship (QSAR) models [54].
- Due to the complicated pathogenetic mechanisms, it is hard to draw an explanation to why attention weights of fingerprints outweigh the other three features in the HIV task. Given that the HIV dataset is the largest one (over 40,000 molecules per task), one possible explanation of this phenomenon is that we use a high-dimensional fingerprint representations (1024 bits).

Concerning the difference between three 2D-based features (namely 2D topological graphs, fingerprints, and SMILES strings), we make the following findings, which we hope could serve as guidelines for future research on molecular representation learning:




Figure 3: Visualizing the learned attention weights on eight molecular property prediction datasets.

Table 3: Results of two case studies: chirality classification and ring counting. We highlight the best and second-best performing results in **boldface** and underline, respectively.

	2D	3D	SM	FP	AttrMask	GraphMAE [46]	GraphMVP [12]	UnifiedMol [48]	MOCO
Chirality (AP, \uparrow)	0.4952	0.4959	0.5505	0.5246	0.5068	0.5196	0.5152	0.5207	<u>0.5469</u>
#Rings (MAE, \downarrow)	0.1949	0.2021	0.3077	0.2590	0.1785	0.1658	0.1694	<u>0.1625</u>	0.1533

- 2D graph representations can encode local information explicitly by resembling chemical structures. Besides, graph-based neural networks can capture long-range local chemical environment through message passing. For example, with molecular graphs, it is more convenient to identify which part of the molecule serves as a scaffold.
- In principle, SMILES strings contain all 2D information of certain molecules, but with atoms and bonds represented in ASCII characters, neural networks may have difficulty in distilling semantic meanings of chemical structures in a numerical way.
- Fingerprint representations are based on local structures and thus such features may be less effective in circumstances where long-range effects induced by topologically distant functional groups predominate, which accounts for relatively small attention weights of fingerprints in Figure 3.

4.4 Additional Experiments

4.4.1 Case Studies: Chirality Classification and Aromatic Ring Counting

In this section, we evaluate MOCO and some state-of-the-art baselines on the two case studies described in Section 2.3 — chirality classification and aromatic ring counting — under the same experimental configuration as in Section 4.1. The results are summarized in Table 3. It can be observed that the performance of MOCO surpasses all of the other methods, validating the effectiveness of our approach. However, it is important to note that while our model achieves the best performance, it does not outperform that of a single modality in the task of chirality classification. Such phenomenon is not uncommon in large vision-language models, e.g., CLIP [55], which are less capable of single-modality models on specific tasks. Our method can identify the most sensible modality-complementary information but not necessarily guarantee superior performance since we encourage alignment across multiple featurizations.


4.4.2 Molecular Property Regression

To demonstrate that the conformations generated by RDKit are helpful, we further conduct an experiment on quantum property regression on the QM9 dataset [20], where 3D conformations generated by RDKit are used for the fine-tuning datasets. This task is known to be closely related to 3D structures. Table 4 presents the performance comparison of MEMO with two non-pretraining (supervised) baselines SchNet and MOCO (denoted by SchNet-NP and MOCO-NP) and two state-of-the-art pretraining baselines GraphMVP [12] and 3D Infomax [11].

It is seen that our MOCO model achieves the best performance on all datasets. GraphMVP that consider only 2D structures during fine-tuning even result in negative transfer on some datasets. Our MOCO, on the contrary, achieves better performance than the supervised baseline, underscoring the value of leveraging 3D structures (as well as other sources of 2D information) during fine-tuning.

Table 4: Results for eight molecule quantum property regression tasks in terms of Mean Absolute Error (MAE, \downarrow). The highest performance is highlighted in **bold**.

Target Unit	μ D	α Bohr ³	ϵ_{HOMO} meV	ϵ_{LUMO} meV	ϵ_{gap} meV	U_0 meV	U meV	$\langle R^2 \rangle$ Bohr ³
SchNet-NP [26]	0.4604	0.3251	95.9740	78.5870	136.4720	98.1240	100.1650	24.3277
MOCO-NP	0.3767	0.2439	73.0625	69.8780	102.2332	77.4708	92.8562	17.5842
GraphMVP [12]	0.3726	0.4390	75.3750	72.3820	104.8370	278.8900	325.8021	22.6433
3D Infomax [11]	0.3644	0.4190	72.0558	67.6203	99.4032	207.2148	219.5415	20.3934
MOCO	0.3618	0.2236	71.5120	58.5890	97.7440	64.3550	66.3958	15.5571

**Figure 4:** Ablation studies on representation aggregation and the fine-tuning strategy.

We also perform experiments on non-quantum property regression tasks. Our proposed MOCO also obtains promising improvements compared to the current state-of-the-art baselines. Please refer to Supporting Information: Section 4.1 for performance comparison and analysis.

4.5 Ablation Studies

Finally, we conduct ablation studies on the representation aggregation module and the fine-tuning strategy. We consider the following model variants for further inspection. Except the modifications in specific modules, other implementations remain the same as previously described.

- **MOCO-Max** removes the attention network in the representation aggregation module in Equation (5) and simply uses max pooling to combine view embeddings.
- **MOCO-Mean** modifies representation aggregation by taking average over view embeddings.
- **MOCO-Freeze** does not fine-tune the representation aggregation module but instead uses the frozen weights of the pretrained model.

We report the performance of model variants in Figure 4. It is seen that all three variants achieve downgraded performance, which empirically rationalizes the design choice of our molecular pre-training framework with collaborative featurizations. Specifically, the performance of MOCO-Max and MOCO-Mean without attention aggregation mechanisms of multiple featurizations is inferior to that of MOCO, demonstrating the necessity of adaptively combining information from multiple featurizations. In addition, MOCO-Freeze occasionally obtains better performance than the two other variants, which indicates that our proposed attention network is able to select information from different views. It does not, however, fine-tune the contribution of featurizations with downstream datasets, where the optimal combination might differ, resulting in performance deterioration.

Moreover, we conduct ablation studies on models that include only three view representations, where the results can be found in Supporting Information: Section 4.2. Results demonstrate the necessity of comprehensively leveraging four views in the proposed MOCO model.

5 Related Work

Molecular representation learning. Traditional methods [51, 56, 57] represent molecular structures with fingerprints. Some prior studies [19, 58, 59] employ tree-based machine learning models such as random forests [60] and XGBoost [61] on fingerprints to predict the properties of molecules. With

Table 5: Comparing MOCO with representative self-supervised methods on molecular pretraining.

Method	Pretraining				Fine-tuning			
	2D	3D	Fingerprint	SMILES	2D	3D	Fingerprint	SMILES
SMILES-BERT [5]				✓				✓
ChemBERTa [6]				✓				✓
AttrMask, ContextPred [7]	✓				✓			
GraphCL [8]	✓				✓			
GraphLoG [9]	✓				✓			
GROVER [47]	✓				✓			
GEM [10]		✓				✓		
3D Infomax [11]	✓	✓			✓			
GraphMVP [12]	✓	✓			✓			
UnifiedMol [48]	✓	✓			✓			
AGPT [80]	✓			✓	✓			✓
MOCO (Ours)	✓	✓	✓	✓	✓	✓	✓	✓

the development of deep learning, neural approaches have been dominating the field given their strong representation ability. One line of work [5, 6, 49] leverages language modeling techniques such as BERT [29] to learn molecular representations based on SMILES strings [14]. However, some argue that sequence-based representations cannot fully capture substructure information and propose to leverage Graph Neural Networks (GNNs), which model molecules as graphs with atoms as nodes and bonds as edges [62–65]. Despite the prosperous progress, they only model 2D topological structures of molecules, without considering the 3D coordinates of atoms that are known to determine certain chemical and physical functionalities of molecules. To address this deficiency, recent work further explicitly considers such 3D geometry and designs equivariant networks to obtain the representations [26, 66–75].

Molecular pre-training. Even though molecular representation learning techniques have been extensively investigated, there are very few labeled datasets available for studying the molecular properties of interest (e.g., drug-likeness or quantum properties). On the other hand, there are abundant unannotated molecules available, which motivates researchers to study pretraining techniques that learn the model weights in a self-supervised manner and transfer the knowledge to downstream datasets with limited annotations via fine-tuning. A series of pretraining frameworks on 2D molecular graph representations have been developed so far [7, 47, 76–78]. Recent work GEM [10] studies large-scale pretraining for 3D geometry representations. Additionally, researchers also study to supplement 2D-graph-based pretraining with 3D conformation information [11, 12, 48, 79]. Note that apart from directly applying the pretrained model to downstream tasks, the molecular representations obtained from pretrained model can also be combined with traditional machine learning methods. For example, blending various representations using random forest is one possible avenue. [80]

Other related work. The application of language models has been gaining prominence in this field. Their usage can be categorized into three main approaches: (1) direct prompting pre-trained language models, such as GPT4 [81], for tasks like property prediction for molecules [82], (2) training or fine-tuning language models specifically tailored for molecular discovery tasks [83], and (3) developing self-supervised transformer models on molecular string representations [6, 84]. The latter approach, which utilizes a variety of molecular representations including string, structural, topological, and geometrical information, aligns closely with our work. Generative machine learning, particularly diffusion models [85, 86], has emerged as a powerful tool for the inverse design of molecules. These models use principles from nonequilibrium thermodynamics to simulate a diffusion process from complex to prior distributions. Learning to reverse this process enables the iterative generation of new molecular structures [87]. Their capacity for pre-training and representation learning, along with their ability to construct new molecular structures, is exemplified in recent works [88, 89].

A succinct comparison of our work with other representative methods is provided in Table 5. Compared to the above studies, our proposed MOCO is the only model that can *adaptively* leverage multiple featurizations for both pretraining and fine-tuning stages.

6 Conclusion

We investigate the neural scaling behavior of molecular representation learning to explore how quantity and quality of molecular data affect the performance. Our research confirms that the performance of MRL follows a power-law relationship with data quantity. Specifically, the graph modality and uniform data distribution exhibit higher learning efficiency in the studied datasets, and the optimal model capacity is highly correlated with task requirements. Moreover, with increasing downstream data, the positive transfer stemming from pre-training diminishes and even lead to negative transfer in the high-data regime. To challenge these scaling laws, we further adapt seven data pruning strategies to molecular data and benchmark their performance. Surprisingly, none of them with simple adaptations can beat the random pruning baseline in MRL. Based on our experimental findings, we raise several key considerations for molecular representation learning, particularly from a data-centric perspective, providing valuable insights for future research endeavors.

Future Work. The flexibility of our method allows its application to a broad range of target properties, extending beyond the specific models and tasks discussed above. For instance, in the domain of protein property prediction, modeling mechanical and dynamical properties [90] could be very important, as exemplified in normal mode frequency prediction [91]. Our approach can also aid in predictions by integrating various protein features like amino acid sequences and protein structures, which have complementary inductive bias [92]. Moreover, the concept of integrating multiple featurizations holds potential for application in chemical reactions. In certain cases, reaction coordinates might be more discernible and easier to learn using one type of featurization over others. For instance, in predicting the forces within different conformations, 3D conformational featurization would be more important than the other featurizations [93].

Data Availability and Acknowledgements

The pretraining data and molecular property prediction benchmarks used in this work, and the source code of our MOCO implementation can be accessed at <https://github.com/Data-reindeer/MOCO>. This work is supported by National Key Program for Technological Innovation 2030 of China (SQ2023AAA010048) and National Natural Science Foundation of China (62141608, 62206291, and 62372454).

References

- [1] Kevin Yang, Kyle Swanson, Wengong Jin, Connor Coley, Philipp Eiden, Hua Gao, Angel Guzman-Perez, Timothy Hopper, Brian Kelley, Miriam Mathea, Andrew Palmer, Volker Settels, Tommi Jaakkola, Klavs Jensen, and Regina Barzilay. Analyzing Learned Molecular Representations for Property Prediction. *J. Chem. Inf. Model.*, 59(8):3370–3388, 2019. 1
- [2] Yuanqi Du, Tianfan Fu, Jimeng Sun, and Shengchao Liu. MolGenSurvey: A Systematic Survey in Machine Learning Models for Molecule Design. *arXiv.org*, 2022. 1
- [3] Shiori Sagawa, Aditi Raghunathan, Pang Wei Koh, and Percy Liang. An Investigation of Why Overparameterization Exacerbates Spurious Correlations. In *ICML*, pages 8346–8356, 2020. 1
- [4] Longlong Jing and Yingli Tian. Self-supervised Visual Feature Learning with Deep Neural Networks: A Survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43(1):4037–4058, 2021. 1, 28
- [5] Sheng Wang, Yuzhi Guo, Yuhong Wang, Hongmao Sun, and Junzhou Huang. SMILES-BERT: Large Scale Unsupervised Pre-Training for Molecular Property Prediction. In *BCB*, pages 429–436, 2019. 1, 10
- [6] Seyone Chithrananda, Gabriel Grand, and Bharath Ramsundar. ChemBERTa: Large-Scale Self-Supervised Pretraining for Molecular Property Prediction. *arXiv.org*, 2020. 1, 10
- [7] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. Strategies for Pre-training Graph Neural Networks. In *ICLR*, 2020. 1, 6, 7, 10, 21, 28

- [8] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. Graph Contrastive Learning with Augmentations. In *NeurIPS*, pages 5812–5823, 2020. 1, 5, 6, 7, 10, 21, 28
- [9] Minghao Xu, Hang Wang, Bingbing Ni, Hongyu Guo, and Jian Tang. Self-supervised Graph-level Representation Learning with Local and Global Structure. In *ICML*, pages 11548–11558, 2021. 1, 6, 7, 10, 28
- [10] Xiaomin Fang, Lihang Liu, Jieqiong Lei, Donglong He, Shanzhuo Zhang, Jingbo Zhou, Fan Wang, Hua Wu, and Haifeng Wang. Geometry Enhanced Molecular Representation Learning for Property Prediction. *Nat. Mach. Intell.*, 4:127–134, 2022. 1, 10
- [11] Hannes Stärk, Dominique Beaini, Gabriele Corso, Prudencio Tossou, Christian Dallago, Stephan Günnemann, and Pietro Liò. 3D Infomax improves GNNs for Molecular Property Prediction. *arXiv.org*, 2021. 1, 2, 8, 9, 10, 23
- [12] Shengchao Liu, Hanchen Wang, Weiyang Liu, Joan Lasenby, Hongyu Guo, and Jian Tang. Pre-training Molecular Graph Representation with 3D Geometry. In *ICLR*, 2022. 1, 2, 5, 6, 7, 8, 9, 10, 21, 23, 25
- [13] Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Janvin. A Neural Probabilistic Language Model. *J. Mach. Learn. Res.*, 3:1137–1155, 2003. 1, 28
- [14] David Weininger. SMILES, A Chemical Language and Information System. *J. Chem. Inf. Comput. Sci.*, 28(1):31–36, 1988. 1, 10
- [15] Jerret Ross, Brian Belgodere, Vijil Chenthamarakshan, Inkit Padhi, Youssef Mroueh, and Payel Das. Molformer: Large Scale Chemical Language Representations Capture Molecular Structure and Properties. *Research Square*, 2022. 2
- [16] Zhengdao Chen, Lei Chen, Soledad Villar, and Joan Bruna. Can Graph Neural Networks Count Substructures? In *NeurIPS*, pages 10383–10395, 2020. 2
- [17] Linfeng Zhang, Jiequn Han, Han Wang, Roberto Car, and Weinan E. Deep Potential Molecular Dynamics: A Scalable Model with the Accuracy of Quantum Mechanics. *Phys. Rev. Lett.*, 120:143001, 2018. 2
- [18] J. S. Smith, O. Isayev, and A. E. Roitberg. ANI-1: An Extensible Neural Network Potential with DFT Accuracy at Force Field Computational Cost. *Chem. Sci.*, 8:3192–3203, 2017. 2
- [19] Zhenqin Wu, Bharath Ramsundar, Evan N. Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S. Pappu, Karl Leswing, and Vijay Pande. MoleculeNet: A Benchmark for Molecular Machine Learning. *Chem. Sci.*, 9:513–530, 2018. 2, 5, 7, 9, 23
- [20] Raghunathan Ramakrishnan, Pavlo O. Dral, Matthias Rupp, and O. Anatole von Lilienfeld. Quantum Chemistry Structures and Properties of 134 kilo Molecules. *Sci. Data*, 1(1):1–7, 2014. 2, 6, 8, 24
- [21] Bharath Ramsundar, Peter Eastman, Patrick Walters, and Vijay Pande. *Deep Learning for the Life Sciences: Applying Deep Learning to Genomics, Microscopy, Drug Discovery, and More*. O’Reilly Media, 2019. 3
- [22] Kenneth Atz, Francesca Grisoni, and Gisbert Schneider. Geometric Deep Learning on Molecular Representations. *Nat. Mach. Intell.*, 3(12):1023–1032, 2021. 3
- [23] H. L. Morgan. The Generation of a Unique Machine Description for Chemical Structures — A Technique Developed at Chemical Abstracts Service. *J. Chem. Doc.*, 5(2):107–113, 1965. 3, 24
- [24] Robert C. Glem, Andreas Bender, Catrin H. Arnby, Lars Carlsson, Scott Boyer, and James Smith. Circular Fingerprints: Flexible Molecular Descriptors with Applications from Physical Chemistry to ADME. *IDrugs*, 9(3):199–204, 2006. 3, 24
- [25] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How Powerful are Graph Neural Networks? In *ICLR*, 2019. 3, 21

- [26] Kristof Schütt, Pieter-Jan Kindermans, Huziel Enoc Saucedo Felix, Stefan Chmiela, Alexandre Tkatchenko, and Klaus-Robert Müller. SchNet: A Continuous-Filter Convolutional Neural Network for Modeling Quantum Interactions. In *NIPS*, pages 991–1001, 2017. 3, 9, 10, 21
- [27] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Uszkoreit Kaiser, and Illia Polosukhin. Attention is All You Need. In *NIPS*, pages 5998–6008, 2017. 4, 22
- [28] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv.org*, 2019. 4, 22, 28
- [29] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT*, pages 4171–4186, 2019. 4, 10, 28
- [30] Simon Axelrod and Rafael Gómez-Bombarelli. GEOM, Energy-Annotated Molecular Conformations for Property Prediction and Molecular Generation. *Sci. Data*, 9(1):185, 2022. 4, 5, 23
- [31] Timothy J. Ritchie and Simon J.F. Macdonald. The Impact of Aromatic Ring Count on Compound Developability – Are Too Many Aromatic Rings a Liability in Drug Design? *Drug Discov. Today*, 14(21):1011–1020, 2009. 4
- [32] Rdkit: Open-source cheminformatics. <http://www.rdkit.org>. 4, 6, 24
- [33] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural Machine Translation by Jointly Learning to Align and Translate. In *ICLR*, 2015. 5
- [34] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A Simple Framework for Contrastive Learning of Visual Representations. In *ICML*, pages 1597–1607, 2020. 5, 28
- [35] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum Contrast for Unsupervised Visual Representation Learning. In *CVPR*, pages 9726–9735, 2020. 5, 28
- [36] Philip Bachman, R. Devon Hjelm, and William Buchwalter. Learning Representations by Maximizing Mutual Information Across Views. In *NeurIPS*, pages 15509–15519, 2019. 5, 28
- [37] Yanqiao Zhu, Yichen Xu, Feng Yu, Qiang Liu, Shu Wu, and Liang Wang. Deep Graph Contrastive Representation Learning. In *GRL+@ICML*, 2020. 5, 28
- [38] Yanqiao Zhu, Yichen Xu, Qiang Liu, and Shu Wu. An Empirical Study of Graph Contrastive Learning. In *NeurIPS Datasets and Benchmarks*, 2021. 5, 28
- [39] Anna Gaulton, Louisa J. Bellis, A. Patricia Bento, Jon Chambers, Mark Davies, Anne Hersey, Yvonne Light, Shaun McGlinchey, David Michalovich, Bissan Al-Lazikani, and John P. Overington. ChEMBL: A Large-Scale Bioactivity Database for Drug Discovery. *Nucleic Acids Res.*, 40(D1):D1100–D1107, 2011. 5, 23
- [40] Johannes Hachmann, Roberto Olivares-Amaya, Sule Atahan-Evrenk, Carlos Amador-Bedolla, Roel S. Sánchez-Carrera, Aryeh Gold-Parker, Leslie Vogt, Anna M. Brockway, and Alán Aspuru-Guzik. The Harvard Clean Energy Project: Large-Scale Computational Screening and Design of Organic Photovoltaics on the World Community Grid. *J. Phys. Chem. Lett.*, 2(17): 2241–2251, 2011. 6, 23
- [41] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open Graph Benchmark: Datasets for Machine Learning on Graphs. In *NeurIPS*, pages 22118–22133, 2020. 6, 24
- [42] William L. Hamilton, Zhitao Ying, and Jure Leskovec. Inductive Representation Learning on Large Graphs. In *NIPS*, pages 1024–1034, 2017. 6, 7, 28

- [43] Fan-Yun Sun, Jordan Hoffmann, Vikas Verma, and Jian Tang. InfoGraph: Unsupervised and Semi-supervised Graph-Level Representation Learning via Mutual Information Maximization. In *ICLR*, 2020. 6, 7, 28
- [44] Ziniu Hu, Yuxiao Dong, Kuansan Wang, Kai-Wei Chang, and Yizhou Sun. GPT-GNN: Generative Pre-Training of Graph Neural Networks. In *KDD*, pages 1857–1867, 2020. 6, 7, 28
- [45] Yuning You, Tianlong Chen, Yang Shen, and Zhangyang Wang. Graph Contrastive Learning Automated. In *ICML*, pages 12121–12132, 2021. 6, 7, 21, 28
- [46] Zhenyu Hou, Xiao Liu, Yukuo Cen, Yuxiao Dong, Hongxia Yang, Chunjie Wang, and Jie Tang. GraphMAE: Self-Supervised Masked Graph Autoencoders. In *KDD*, pages 594–604, 2022. 6, 7, 8
- [47] Yu Rong, Yatao Bian, Tingyang Xu, Weiyang Xie, Ying Wei, Wenbing Huang, and Junzhou Huang. Self-Supervised Graph Transformer on Large-Scale Molecular Data. In *NeurIPS*, pages 12559–12571, 2020. 6, 7, 10
- [48] Jinhua Zhu, Yingce Xia, Lijun Wu, Shufang Xie, Tao Qin, Wengang Zhou, Houqiang Li, and Tie-Yan Liu. Unified 2d and 3d pre-training of molecular representations. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 2626–2636, 2022. 6, 7, 8, 10
- [49] Dong Chen, Jiaxin Zheng, Guo-Wei Wei, and Feng Pan. Extracting predictive representations from hundreds of millions of molecules. *The journal of physical chemistry letters*, 12(44): 10793–10801, 2021. 6, 7, 10
- [50] Sereina Riniker and Gregory A. Landrum. Better Informed Distance Geometry: Using What We Know To Improve Conformation Generation. *J. Chem. Inf. Model.*, 55(12):2562–2574, 2015. 6, 24
- [51] David Rogers and Mathew Hahn. Extended-Connectivity Fingerprints. *J. Chem. Inf. Model.*, 50(5):742–754, 2010. 6, 9, 24
- [52] Jonathan B. Baell and Georgina A. Holloway. New substructure filters for removal of pan assay interference compounds (pains) from screening libraries and for their exclusion in bioassays. *J. Med. Chem.*, 53(7):2719–2740, 2010. 7
- [53] Anthony J. Suckling, M.G. Rumsby, and Michael William Blackburn Bradbury. *Blood-Brain Barrier in Health and Disease*. Ellis Horwood Health Science Series. Ellis Horwood, 1986. 7
- [54] Prabu Manoharan, R.S.K. Vijayan, and Nanda Ghoshal. Rationalizing Fragment Based Drug Discovery for BACE1: Insights From FB-QSAR, FB-QSSR, Multi Objective (MO-QSPR) and MIF Studies. *J. Comput. Aided Mol. Des.*, 24(10):843–864, 2010. 7
- [55] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 8
- [56] Raymond E. Carhart, Dennis H. Smith, and R. Venkataraghavan. Atom Pairs as Molecular Features in Structure-Activity Studies: Definition and Applications. *J. Chem. Inf. Comput. Sci.*, 25(2):64–73, 1985. 9
- [57] Ramaswamy Nilakantan, Norman Bauman, J. Scott Dixon, and R. Venkataraghavan. Topological Torsion: A New Molecular Descriptor for SAR Applications. Comparison with Other Descriptors. *J. Chem. Inf. Comput. Sci.*, 27(2):82–85, 1987. 9
- [58] Vladimir Svetnik, Andy Liaw, Christopher Tong, and Ting Wang. Application of Breiman’s Random Forest to Modeling Structure-Activity Relationships of Pharmaceutical Molecules. In *MCS*, pages 334–343, 2004. 9
- [59] Jesse G. Meyer, Shengchao Liu, Ian J. Miller, Joshua J. Coon, and Anthony Gitter. Learning Drug Functions from Chemical Structures with Convolutional Neural Networks and Random Forests. *J. Chem. Inf. Model.*, 59(10):4438–4449, 2019. 9

- [60] Leo Breiman. Random Forests. *Mach. Learn.*, 45(1):5–32, 2001. 9
- [61] Ricky T. Q. Chen and Carlos Guestrin. XGBoost: A Scalable Tree Boosting System. In *KDD*, pages 785–794, 2016. 9
- [62] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural Message Passing for Quantum Chemistry. In *ICML*, pages 1263–1272, 2017. 10
- [63] Shengchao Liu, Mehmet Furkan Demirel, and Yingyu Liang. N-Gram Graph: Simple Un-supervised Representation for Graphs, with Applications to Molecules. In *NeurIPS*, pages 8464–8476, 2019. 10
- [64] Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu. Do Transformers Really Perform Badly for Graph Representation? In *NeurIPS*, pages 28877–28888, 2021. 10
- [65] Dingshuo Chen, Yanqiao Zhu, Jieyu Zhang, Yuanqi Du, Zhixun Li, Qiang Liu, Shu Wu, and Liang Wang. Uncovering neural scaling laws in molecular representation learning. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. 10
- [66] Johannes Klicpera, Janek Groß, and Stephan Günnemann. Directional Message Passing for Molecular Graphs. In *ICLR*, 2020. 10
- [67] Victor Garcia Satorras, Emiel Hooeboom, and Max Welling. E(n) Equivariant Graph Neural Networks. In *ICML*, pages 9323–9332, 2021. 10
- [68] Fabian Fuchs, Daniel E. Worrall, Volker Fischer, and Max Welling. SE(3)-Transformers: 3D Roto-Translation Equivariant Attention Networks. In *NeurIPS*, pages 1970–1981, 2020. 10
- [69] Kristof Schütt, Oliver T. Unke, and Michael Gastegger. Equivariant Message Passing for the Prediction of Tensorial Properties and Molecular Spectra. In *ICML*, pages 9377–9388, 2021. 10
- [70] Weitao Du, He Zhang, Yuanqi Du, Qi Meng, Wei Chen, Bin Shao, and Tie-Yan Liu. Equivariant Vector Field Network for Many-Body System Modeling. *arXiv.org*, 2021. 10
- [71] Yi Liu, Limei Wang, Meng Liu, Xuan Zhang, Bora Oztekin, and Shuiwang Ji. Spherical Message Passing for 3D Graph Networks. *arXiv.org*, 2021. 10
- [72] Johannes Gasteiger, Florian Becker, and Stephan Günnemann. GemNet: Universal Directional Graph Neural Networks for Molecules. In *NeurIPS*, pages 6790–6802, 2021. 10
- [73] Simon L. Batzner, Tess E. Smidt, Lixin Sun, Jonathan P. Mailoa, Mordechai Kornbluth, Nicola Molinari, and Boris Kozinsky. SE(3)-Equivariant Graph Neural Networks for Data-Efficient and Accurate Interatomic Potentials. *arXiv.org*, 2021. 10
- [74] Johannes Brandstetter, Rob Hesselink, Elise van der Pol, Erik J. Bekkers, and Max Welling. Geometric and Physical Quantities Improve E(3) Equivariant Message Passing. In *ICLR*, 2022. 10
- [75] Minkai Xu, Shitong Luo, Yoshua Bengio, Jian Peng, and Jian Tang. Learning Neural Generative Dynamics for Molecular Conformation Generation. In *ICLR*, 2021. 10
- [76] Zaixi Zhang, Qi Liu, Hao Wang, Chengqiang Lu, and Chee-Kong Lee. Motif-based Graph Self-Supervised Learning for Molecular Property Prediction. In *NeurIPS*, pages 15870–15882, 2021. 10
- [77] Yuyang Wang, Jianren Wang, Zhonglin Cao, and Amir Barati Farimani. Molecular Contrastive Learning of Representations via Graph Neural Networks. *Nat. Mach. Intell.*, 4:279–287, 2022. 10
- [78] Pengyong Li, Jun Wang, Yixuan Qiao, Hao Chen, Yihuan Yu, Xiaojun Yao, Peng Gao, Guotong Xie, and Sen Song. Learn Molecular Representations From Large-Scale Unlabeled Molecules for Drug Discovery. *arXiv.org*, 2020. 10

- [79] Shuwen Yang, Ziyao Li, Guojie Song, and Lingsheng Cai. Deep Molecular Representation Learning via Fusing Physical and Chemical Information. In *NeurIPS*, pages 16346–16357, 2021. 10
- [80] Dong Chen, Kaifu Gao, Duc Duy Nguyen, Xin Chen, Yi Jiang, Guo-Wei Wei, and Feng Pan. Algebraic graph-assisted bidirectional transformers for molecular property prediction. *Nature communications*, 12(1):3521, 2021. 10
- [81] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mo Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rameev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. GPT-4 Technical Report, 2023. 10
- [82] Shengchao Liu, Jiong Xiao Wang, Yijin Yang, Chengpeng Wang, Ling Liu, Hongyu Guo, and Chaowei Xiao. Chatgpt-powered conversational drug editing using retrieval and domain feedback, 2023. 10

- [83] Jie Pan. Large language model for molecular chemistry. *Nature Computational Science*, 3(1): 5–5, 2023. 10
- [84] Jerret Ross, Brian Belgodere, Vijil Chenthamarakshan, Inkit Padhi, Youssef Mroueh, and Payel Das. Large-scale chemical language representations capture molecular structure and properties. *Nature Machine Intelligence*, 4(12):1256–1264, 2022. 10
- [85] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 10
- [86] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019. 10
- [87] Emiel Hoogeboom, Victor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3d. In *International conference on machine learning*, pages 8867–8887. PMLR, 2022. 10
- [88] Weitao Du, Jiujiu Chen, Xuecang Zhang, Zhiming Ma, and Shengchao Liu. Molecule joint auto-encoding: Trajectory pretraining with 2d and 3d diffusion. *arXiv preprint arXiv:2312.03475*, 2023. 10
- [89] Shengchao Liu, Hongyu Guo, and Jian Tang. Molecular geometry pretraining with se (3)-invariant denoising distance matching. *arXiv preprint arXiv:2206.13602*, 2022. 10
- [90] Zuobai Zhang, Minghao Xu, Aurelie Lozano, Vijil Chenthamarakshan, Payel Das, and Jian Tang. Pre-training protein encoder via siamese sequence-structure diffusion trajectory prediction. In *Annual Conference on Neural Information Processing Systems*, 2023. 11
- [91] Kai Guo and Markus J Buehler. Rapid prediction of protein natural frequencies using graph neural networks. *Digital Discovery*, 1(3):277–285, 2022. 11
- [92] Yiwen Hu and Markus J Buehler. End-to-end protein normal mode frequency predictions using language and graph models and application to sonification. *ACS nano*, 16(12):20656–20670, 2022. 11
- [93] Markus Meuwly. Machine learning for chemical reactions. *Chemical Reviews*, 121(16): 10218–10239, 2021. 11
- [94] Boris Weisfeiler and Andrei Leman. A Reduction of a Graph to a Canonical Form and an Algebra Arising During This Reduction. *Nauchno-Technicheskaya Informatsia*, 2(9):12–16, 1968. 21
- [95] Philip Gage. A New Algorithm for Data Compression. *C Users J.*, 12(2):23–38, 1994. 22
- [96] Stefan Grimme. Exploration of Chemical Compound, Conformer, and Reaction Space with Meta-Dynamics Simulations Based on Tight-Binding Quantum Chemical Calculations. *J. Chem. Theory Comput.*, 15(5):2847–2862, 2019. 23
- [97] John S. Delaney. ESOL: Estimating Aqueous Solubility Directly from Molecular Structure. *J. Chem. Inf. Comput. Sci.*, 44(3):1000–1005, 2004. 23
- [98] AIDS Antiviral Screen Data. URL <https://wiki.nci.nih.gov/display/NCIDTPdata/AIDS+Antiviral+Screen+Data>. 23
- [99] Sebastian G. Rohrer and Knut Baumann. Maximum Unbiased Validation (MUV) Data Sets for Virtual Screening Based on PubChem Bioactivity Data. *J. Chem. Inf. Model.*, 49(2):169–184, 2009. 23
- [100] Francisco-Javier Gamero, Laura M. Sanz, Jaume Vidal, Cristina de Cozar, Emilio Alvarez, Jose-Luis Lavandera, Dana E. Vanderwall, Darren V. S. Green, Vinod Kumar, Samiul Hasan, James R. Brown, Catherine E. Peishoff, Lon R. Cardon, and Jose F. Garcia-Bustos. Thousands of Chemical Starting Points for Antimalarial Lead Identification. *Nature*, 465(7296):305–310, 2010. 23

- [101] Ines Filipa Martins, Ana L. Teixeira, Luis Pinheiro, and Andre O. Falcao. A Bayesian Approach to in Silico Blood-Brain Barrier Penetration Modeling. *J. Chem. Inf. Model.*, 52(6):1686–1697, 2012. 23
- [102] Tox21 Data Challenge 2014, 2014. URL <https://tripod.nih.gov/tox21/challenge/>. 23
- [103] Ann M. Richard, Richard S. Judson, Keith A. Houck, Christopher M. Grulke, Patra Volarath, Inthirany Thillainadarajah, Chihae Yang, James Rathman, Matthew T. Martin, John F. Wambaugh, Thomas B. Knudsen, Jayaram Kancherla, Kamel Mansouri, Grace Patlewicz, Antony J. Williams, Stephen B. Little, Kevin M. Crofton, and Russell S. Thomas. ToxCast Chemical Landscape: Paving the Road to 21st Century Toxicology. *Chem. Res. Toxicol.*, 29(8):1225–1251, 2016. 23
- [104] Kaitlyn M. Gayvert, Neel S. Madhukar, and Olivier Elemento. A Data-Driven Approach to Predicting Successes and Failures of Clinical Trials. *Cell Chem. Biol.*, 23(10):1294–1301, 2016. 23
- [105] Michael Kuhn, Ivica Letunic, Lars Juhl Jensen, and Peer Bork. The SIDER Database of Drugs and Side Effects. *Nucleic Acids Res.*, 44(D1):D1075–D1079, 2016. 23
- [106] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *NeurIPS*, pages 8024–8035, 2019. 24
- [107] Matthias Fey and Jan Eric Lenssen. Fast Graph Representation Learning with PyTorch Geometric. In *RLGM@ICLR*, 2019. 24
- [108] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. HuggingFace’s Transformers: State-of-the-Art Natural Language Processing. In *EMNLP (Demo)*, pages 38–45, 2020. 24
- [109] Xavier Glorot and Yoshua Bengio. Understanding the Difficulty of Training Deep Feedforward Neural Networks. In *AISTATS*, pages 249–256, 2010. 24
- [110] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *ICLR*, 2015. 24
- [111] Nitish Srivastava, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan R. Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks From Overfitting. *J. Mach. Learn. Res.*, 15(1):1929–1958, 2014. 25
- [112] Mario Krenn, Florian Häse, AkshatKumar Nigam, Pascal Friederich, and Alan Aspuru-Guzik. Self-referencing embedded strings (selfies): A 100% robust molecular string representation. *Machine Learning: Science and Technology*, 1(4):045024, 2020. 27
- [113] Aditya R Thawani, Ryan-Rhys Griffiths, Arian Jamasb, Anthony Bourached, Penelope Jones, William McCorkindale, Alexander A Aldrick, and Alpha A Lee. The photoswitch dataset: A molecular machine learning benchmark for the advancement of synthetic chemistry. *arXiv preprint arXiv:2008.03226*, 2020. 27
- [114] Ruoxi Sun. Does gnn pretraining help molecular representation? *arXiv preprint arXiv:2207.06010*, 2022. 27
- [115] Deepak Pathak, Philipp Krähenbühl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context Encoders: Feature Learning by Inpainting. In *CVPR*, pages 2536–2544, 2016. 28
- [116] Mehdi Noroozi and Paolo Favaro. Unsupervised Learning of Visual Representations by Solving Jigsaw Puzzles. In *ECCV*, pages 69–84, 2016. 28

- [117] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful Image Colorization. In *ECCV*, pages 649–666, 2016. 28
- [118] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Colorization as a Proxy Task for Visual Understanding. In *CVPR*, pages 840–849, 2017. 28
- [119] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised Representation Learning by Predicting Image Rotations. In *ICLR*, 2018. 28
- [120] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep Clustering for Unsupervised Learning of Visual Features. In *ECCV*, pages 139–156, 2018. 28
- [121] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeffrey Dean. Distributed Representations of Words and Phrases and their Compositionality. In *NIPS*, pages 3111–3119, 2013. 28
- [122] Ryan Kiros, Yukun Zhu, Ruslan R. Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Skip-Thought Vectors. In *NIPS*, pages 3294–3302, 2015. 28
- [123] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *ACL*, pages 7871–7880, 2020. 28
- [124] Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation Learning with Contrastive Predictive Coding. *arXiv.org*, 2018. 28
- [125] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised Learning of Visual Features by Contrasting Cluster Assignments. In *NeurIPS*, pages 9912–9924, 2020. 28
- [126] Xinlei Chen and Kaiming He. Exploring Simple Siamese Representation Learning. In *CVPR*, pages 15745–15753, 2021. 28
- [127] Tianyu Gao, Xingcheng Yao, and Danqi Chen. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In *EMNLP*, pages 6894–6910, 2021. 28
- [128] Tete Xiao, Xiaolong Wang, Alexei A. Efros, and Trevor Darrell. What Should Not Be Contrastive in Contrastive Learning. In *ICLR*, 2021. 28
- [129] Yonglong Tian, Chen Sun, Ben Poole, Dilip Krishnan, Cordelia Schmid, and Phillip Isola. What Makes for Good Views for Contrastive Learning? In *NeurIPS*, pages 6827–6839, 2020. 28
- [130] Senthil Purushwalkam and Abhinav Gupta. Demystifying Contrastive Self-Supervised Learning: Invariances, Augmentations and Dataset Biases. In *NeurIPS*, pages 3407–3418, 2020. 28
- [131] Julius von Kügelgen, Yash Sharma, Luigi Gresele, Wieland Brendel, Bernhard Schölkopf, Michel Besserve, and Francesco Locatello. Self-Supervised Learning with Data Augmentations Provably Isolates Content from Style. In *NeurIPS*, pages 16451–16467, 2021. 28
- [132] Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep Contextualized Word Representations. In *NAACL-HLT*, pages 2227–2237, 2018. 28
- [133] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. In *ICLR*, 2020. 28
- [134] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving Language Understanding by Generative Pre-Training. Technical report, OpenAI Blog, 2018. 28

- [135] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language Models are Unsupervised Multitask Learners. Technical report, OpenAI Blog, 2019. [28](#)
- [136] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language Models are Few-Shot Learners. In *NeurIPS*, pages 1877–1901, 2020. [28](#)
- [137] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*, 2021. [28](#)
- [138] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked Autoencoders Are Scalable Vision Learners. In *CVPR*, 2022. [28](#)
- [139] Chen Wei, Haoqi Fan, Saining Xie, Chao-Yuan Wu, Alan Yuille, and Christoph Feichtenhofer. Masked Feature Prediction for Self-Supervised Visual Pre-Training. *arXiv.org*, 2021. [28](#)
- [140] Lirong Wu, Haitao Lin, Zhangyang Gao, Cheng Tan, and Stan Z. Li. Self-supervised on Graphs: Contrastive, Generative, or Predictive. *IEEE Trans. Knowl. Data Eng.*, 2022. [28](#)
- [141] Yaochen Xie, Zhao Xu, Zhengyang Wang, and Shuiwang Ji. Self-Supervised Learning of Graph Neural Networks: A Unified Review. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2022. [28](#)
- [142] Yixin Liu, Shirui Pan, Ming Jin, Chuan Zhou, Yu Zheng, Feng Xia, and Philip S. Yu. Graph Self-Supervised Learning: A Survey. *IEEE Trans. Knowl. Data Eng.*, 2022. [28](#)
- [143] Yuning You, Tianlong Chen, Zhangyang Wang, and Yang Shen. When Does Self-Supervision Help Graph Convolutional Networks? In *ICML*, pages 10871–10880, 2020. [28](#)
- [144] Ke Sun, Zhouchen Lin, and Zhanxing Zhu. Multi-Stage Self-Supervised Learning for Graph Convolutional Networks on Graphs with Few Labeled Nodes. In *AAAI*, pages 5892–5899, 2020. [28](#)
- [145] Thomas N. Kipf and Max Welling. Variational Graph Auto-Encoders. In *BDL@NIPS*, 2016. [28](#)
- [146] Petar Veličković, William Fedus, William L. Hamilton, Pietro Liò, Yoshua Bengio, and R. Devon Hjelm. Deep Graph Infomax. In *ICLR*, 2019. [28](#)
- [147] Kaveh Hassani and Amir Hosein Khasahmadi. Contrastive Multi-View Representation Learning on Graphs. In *ICML*, pages 4116–4126, 2020. [28](#)
- [148] Shuai Lin, Pan Zhou, Zi-Yuan Hu, Shuojia Wang, Ruihui Zhao, Yefeng Zheng, Liang Lin, Eric P. Xing, and Xiaodan Liang. Prototypical Graph Contrastive Learning. *arXiv.org*, 2021. [28](#)
- [149] Yanqiao Zhu, Yichen Xu, Feng Yu, Qiang Liu, Shu Wu, and Liang Wang. Graph Contrastive Learning with Adaptive Augmentation. In *WWW*, pages 2069–2080, 2021. [28](#)
- [150] Jun Xia, Lirong Wu, Jintao Chen, Bozhen Hu, and Stan Z. Li. SimGRACE: A Simple Framework for Graph Contrastive Learning without Data Augmentation. In *WWW*, pages 1070–1079, 2022. [28](#)
- [151] Puja Trivedi, Ekdeep Singh Lubana, Yujun Yan, Yaoqing Yang, and Danai Koutra. Augmentations in Graph Contrastive Learning: Current Methodological Flaws & Towards Better Practices. In *WWW*, pages 1538–1549, 2022. [28](#)

The supporting information is composed of (a) detailed implementation of four view encoders, (b) descriptions of pre-training and fine-tuning datasets, along with data preprocessing methodologies, (3) specifications of experimental settings, including computing infrastructures, hyperparameters specifications, and computational complexity analysis, (4) additional experimental results, incorporating more experiments on molecular property regression tasks, ablation studies using alternative featurizations, experiments on variant fingerprint encoders, and explanation for the impact of pre-training data scale, and (5) a more extensive review of related work.

A Implementation of View Encoders

In this section, we introduce the detailed implementation of the four view encoders. We denote the representation for node (atom) v_i as \mathbf{h}_i and the representation at the graph (molecule) level as \mathbf{z} . As each encoder is independent to each other, we omit the superscript representing the specific view $m \in \{2D, 3D, FP, SM\}$ is clear for notation simplicity. Also, for clarity, when the context is clear, we omit the subscript j that indexes the molecule.

Embedding 2D graphs. Graph Isomorphism Network (GIN) [25] is a simple and effective model to learn discriminative graph representations, which is proved to have the same representational power as the Weisfeiler-Lehman test [94]. Since GIN has been widely adopted for 2D graph representation learning [7, 8, 45], we leverage a GIN model to obtain the representations for the 2D molecular graphs. Recall that each molecule is represented as $\mathcal{G} = (\mathbf{A}, \mathbf{X}, \mathbf{E})$, where \mathbf{A} is the adjacency matrix, \mathbf{X} and \mathbf{E} are features for atoms and bonds respectively. The layer-wise propagation rule of GIN can be written as:

$$\mathbf{h}_i^{(k+1)} = f_{\text{atom}}^{(k+1)} \left(\mathbf{h}_i^{(k)} + \sum_{j \in \mathcal{N}(i)} \left(\mathbf{h}_j^{(k)} + f_{\text{bond}}^{(k+1)}(\mathbf{E}_{ij}) \right) \right), \quad (9)$$

where the input features $\mathbf{h}_i^{(0)} = \mathbf{x}_i$, $\mathcal{N}(i)$ is the neighborhood set of atom v_i , and f_{atom} , f_{bond} are two MultiLayer Perceptron (MLP) layers for transforming atoms and bonds features, respectively. By stacking K layers, we can incorporate K -hop neighborhood information into each center atom in the molecular graph. Then, we take the output of the last layer as the atom representations and further use the mean pooling to get the graph-level molecular representation:

$$\mathbf{z}^{2D} = \frac{1}{N} \sum_{i \in \mathcal{V}} \mathbf{h}_i^{(K)}. \quad (10)$$

Embedding 3D graphs. Following GraphMVP [12], we use the SchNet [26] as the encoder for the 3D geometry graphs. SchNet models message passing in the 3D space as continuous-filter convolutions, which is composed of a series of hidden layers, given as follows:

$$\mathbf{h}_i^{(k+1)} = f_{\text{MLP}} \left(\sum_{j=1}^N f_{\text{FG}}(\mathbf{h}_j^{(t)}, \mathbf{r}_i, \mathbf{r}_j) \right) + \mathbf{h}_i^{(t)}, \quad (11)$$

where the input $\mathbf{h}_i^{(0)} = \mathbf{a}_i$ is an embedding dependent on the type of atom v_i , $f_{\text{FG}}(\cdot)$ denotes the filter-generating network. To ensure rotational invariance of a predicted property, the message passing function is restricted to depend only on rotationally invariant inputs such as distances, which satisfying the energy properties of rotational equivariance by construction. Moreover, SchNet adopts radial basis functions to avoid highly correlated filters. The filter-generating network is defined as follow:

$$f_{\text{FG}}(\mathbf{x}_j, \mathbf{r}_i, \mathbf{r}_j) = \mathbf{x}_j \cdot e_k(\mathbf{r}_i - \mathbf{r}_j) = \mathbf{x}_j \cdot \exp(-\gamma \|\|\mathbf{r}_i - \mathbf{r}_j\|_2 - \mu\|_2^2). \quad (12)$$

Similarly, for non-quantum properties prediction concerned in this work, we take the average of the node representations as the 3D molecular embedding:

$$\mathbf{z}^{3D} = \frac{1}{N} \sum_{i \in \mathcal{V}} \mathbf{h}_i^{(K)}, \quad (13)$$

where K is the number of hidden layers.

Embedding fingerprints. Due to the discrete and extremely sparse nature of fingerprint vectors, we first transform all F binary feature fields into a dense embedding matrix $\mathbf{F}' \in \mathbb{R}^{F \times D_F}$ via embedding lookup. Then, we introduce a positional embedding matrix $\mathbf{P} \in \mathbb{R}^{F \times D_F}$ to capture the positional relationship among bits in the fingerprint vector, which is defined as:

$$\mathbf{P}_{p,2i} = \sin(p/10000^{2i/D_F}), \quad (14)$$

$$\mathbf{P}_{p,2i+1} = \cos(p/10000^{2i/D_F}), \quad (15)$$

where p denotes the corresponding bit position and i corresponds to the i -th embedding dimension. The positional embedding matrix will be added to the transformed embedding matrix:

$$\mathbf{F} = \mathbf{F}' + \mathbf{P}. \quad (16)$$

Thereafter, we use a multihead Transformer [27] to model the interaction among those feature fields. Specifically, we first transform each feature into a new embedding space as:

$$\mathbf{Q}^{(h)} = \mathbf{F}\mathbf{W}_Q^{(h)}, \quad (17)$$

$$\mathbf{K}^{(h)} = \mathbf{F}\mathbf{W}_K^{(h)}, \quad (18)$$

$$\mathbf{V}^{(h)} = \mathbf{F}\mathbf{W}_V^{(h)}, \quad (19)$$

where the three linear transformation matrices $\mathbf{W}_Q^{(h)}, \mathbf{W}_K^{(h)}, \mathbf{W}_V^{(h)} \in \mathbb{R}^{D_F \times D/H}$ parameterize the query, key, and value transformations for the h -th attention head, respectively. Following that, we compute the attention scores among all feature pairs and then linearly combine the value matrix from all H attention heads:

$$\mathbf{W}_A^{(h)} = \text{softmax} \left(\frac{\mathbf{Q}^{(h)}(\mathbf{K}^{(h)})^\top}{\sqrt{D_H}} \right), \quad (20)$$

$$\hat{\mathbf{Z}} = \left[\mathbf{W}_A^{(1)}\mathbf{V}^{(1)}; \mathbf{W}_A^{(2)}\mathbf{V}^{(2)}; \dots; \mathbf{W}_A^{(H)}\mathbf{V}^{(H)} \right], \quad (21)$$

Finally, we perform sum pooling on the resulting embedding matrix $\hat{\mathbf{Z}} \in \mathbb{R}^{F \times D_F}$ and use a linear model f_{LIN} to obtain the final fingerprint embedding $\mathbf{z}^{\text{FP}} \in \mathbb{R}^D$:

$$\mathbf{z}^{\text{FP}} = f_{\text{LIN}} \left(\sum_{d=1}^{D_F} \hat{\mathbf{Z}}_d \right). \quad (22)$$

Embedding SMILES strings. Given ASCII-encoded SMILES strings, we first tokenize them with the Byte-Pair Encoder (BPE) tokenizer [95], which strikes a balance among character- and word-level representations and allows to handle large vocabularies in molecular corpora. Specifically, BPE finds the best word segmentation by iteratively and greedily merging frequent pairs of characters. In our implementation, we use a max vocabulary size of 52K tokens for both pretraining and downstream datasets.

After tokenization, we first pretrain a RoBERTa [28] model on the pretraining dataset with the masking language model as the sole training objective, as SMILES strings do not possess sequential relationships. To be specific, 15% tokens in a SMILES string are randomly selected and replaced with a special token [MASK]. We also insert a special token [CLS] to each string to represent the whole molecule. The training objective function is to independently predict the original tokens given the output on masked tokens. Finally, the representation of the [CLS] token is regarded as the molecular embedding.

After pretraining the RoBERTa backbone, we freeze its parameters and leverage an additional MLP layer on top of each molecular embedding to obtain the final representation for each SMILES string. This strategy improves memory efficiency and thus enables larger batch sizes for our contrastive pretraining framework.

B Dataset Description

In this section, we briefly introduce the datasets used for pretraining and fine-tuning, as well as details of dataset preprocessing. Basic dataset statistics is summarized in Table 6.

Table 6: Statistics of datasets used in experiments. The first section describes the datasets with 3D information which is used for pre-training; the later two sections describe datasets for fine-tuning.

	Dataset	#Molecules	Avg. #atoms	Avg. #bonds	#Tasks	Avg. degree
	GEOM-Drug	304,466	44.40	46.40	—	2.09
Classification	BBBP	2,039	24.06	25.95	1	2.16
	Tox21	7,831	18.57	19.29	12	2.08
	ToxCast	8,576	18.78	19.26	617	2.05
	SIDER	1,427	33.64	35.36	27	2.10
	ClinTox	1,477	26.16	27.88	2	2.13
	MUV	93,087	24.23	26.28	17	2.17
	HIV	41,127	25.51	27.47	1	2.15
	BACE	1,513	34.09	36.86	1	2.16
Regression	ESOL	1,128	13.30	13.69	1	2.06
	Lipophilicity	4,200	27.04	29.50	1	2.18
	Malaria	9,999	30.36	33.20	1	2.19
	CEP	29,978	27.66	33.39	1	2.41
	QM9	130,831	18.03	18.65	8	2.07

B.1 Pretraining Datasets

We choose GEOM-Drugs⁴ [30] as the pre-training dataset, which contains high-quality conformers for 304,466 mid-sized organic molecules with experimental data. The conformer information in GEOM-Drugs is generated using the CREST [96] program, which provides reliable and accurate structure generation. Note that atoms usually have multiple conformations resulting in potentially different chemical properties. In our work, we focus on the conformations of the lowest energy, as they are more likely to occur naturally [11, 12]. Moreover, since the original full set is large (317K molecules with over 9M conformations), we follow GraphMVP [12] to sample a subset of 50K molecules, each with its top-5 conformations, for pretraining. We use the same random seeds to ensure data as with GraphMVP.

B.2 Fine-Tuning Datasets

For fine-tuning, we use 12 datasets collected from MoleculeNet⁵ [19], ChEMBL [39], and CEP [40], which target on different properties and distinct tasks. These properties can be divided into four main categories: physical chemistry, biophysics, physiology, and quantum properties.

Physical chemistry. ESOL [97] consists of water solubility data recording whether molecules are water-soluble. The Lipophilicity dataset is a subset of ChEMBL [39] measuring the molecule octanol/water distribution coefficient. The CEP dataset is a subset of the Harvard Clean Energy Project (CEP) [40], which estimates the organic photovoltaic efficiency.

Biophysics. The HIV dataset [98] is introduced by Drug Therapeutics Program (DTP) AIDS Antiviral Screen, which tests the molecular ability to inhibit HIV replication. The Maximum Unbiased Validation (MUV) group [99] is another benchmark dataset selected from PubChem BioAssay by applying a refined nearest neighbor analysis. The BACE dataset provides qualitative binding results for a set of inhibitors of human β -secretase 1 (BACE-1). The Malaria dataset [100] assesses the drug efficacy in inhibiting parasites that cause malaria.

Physiology. The Blood-brain barrier penetration (BBBP) dataset [101] models the barrier permeability of molecules targeting central nervous system. Tox21 [102], ToxCast [103], and ClinTox [104] are all related to the toxicity of molecular compounds. The Side Effect Resource (SIDER) [105] is a dataset measuring the adverse drug reactions of 27 system organ classes of marketed drugs.

⁴<https://github.com/learningmatter-mit/geom>

⁵<https://github.com/deepchem/deepchem>

Quantum properties. The QM9 dataset [20] consists of 134K stable small organic molecules made up of four kinds of heavy atoms (C, O, F, N) and hydrogen (H), and each atom is associated with 12 quantum mechanical properties as regression targets, which are closely related to geometries minimal in energy. In our experiments, we delete 3,054 uncharacterized molecules which failed the geometry consistency check [20] and adopt the remaining 130,831 molecules for the downstream task with 8 quantum-related property regression targets. Specifically, the property μ represents the dipole moment of the molecule, the property α represents the corresponding isotropic polarizability, and the $\langle R^2 \rangle$ denotes the electronic spatial extent. The property ϵ_{HOMO} , ϵ_{LUMO} , and ϵ_{gap} represent the energy of HOMO, the energy of LUMO, and the energy gap of the former two energies respectively. The property U_0 and U denote the internal energy of the molecule at 0K and 298.15K, respectively.

A note on overlapping molecules. We find that only 14,087 ($\approx 28.2\%$) molecules in the MUV dataset overlap with those in GEOM-Drugs and there is no overlapping molecules in the other fine-tuning datasets. Considering that *no* labels is used during pretraining, for those overlapped molecules, we are evaluating the in-domain performance (unsupervised learning performance), while for those non-overlapped ones, we are evaluating the transfer learning performance.

B.3 Dataset Preprocessing

For classification tasks, we leverage atom types and chirality tags as atom attributes, while the type and direction of the bond are corresponding bond attributes. Both the atom and bond attributes are expressed in the form of discrete indices without further embedding. For regression tasks, we first transform discrete atom and bond attributes through learnable embedding lookup layers following OGB [41]. Since molecules in the fine-tuning datasets do not have 3D information available, we use ETKDG [50] in RDKit [32] to generate molecular conformations.

Constructing fingerprints. For both pretraining and fine-tuning datasets, we use RDKit to generate molecular fingerprints. Morgan fingerprints [23, 24] encode molecules in fixed-length binary strings, with bits indicating presence or absence of specific substructures. The algorithm assigns an initial identifier to each non-hydrogen atom according to a set of atomic invariants, iteratively updates the identifiers among neighborhood atoms within certain hops, and encodes the identifiers using a hash function. After hashing all of these identifiers into a fixed-length binary string, the representation provides information on topological characteristics of the molecule.

In our implementation, we set the diameter of neighborhood to 2, the length of fingerprints to 1024, and follow the default configuration of ECFP4 [51], which uses the following connectivity invariants to construct the initial node features:

- The atomic number
- The number of heavy (non-hydrogen) neighbor atoms
- The number of attached hydrogens
- The formal charge
- Atom isotopes
- Whether the atom is part of at least one ring

C Implementation Details

C.1 Computing Infrastructures

Software infrastructures. The proposed MOCO is implemented in Python 3.7, with the following supporting libraries: PyTorch 1.10.2 [106], PyG 2.0.3 [107], RDKit 2022.03.1 [32], OGB 1.3.3 [41] and HuggingFace’s Transformers 4.17.0 [108].

Hardware infrastructures. We conduct all experiments on a computer server with 8 NVIDIA GeForce RTX 3090 GPUs (with 24GB memory each) and 256 AMD EPYC 7742 CPUs.

C.2 Hyperparameter Specifications

All model parameters are initialized with the Glorot initialization [109] and trained using the Adam optimizer [110].

For hyperparameters of the GIN model and the SchNet model, we keep them in line with the baseline GraphMVP [12]. Specifically, for the GIN model, we set the number of convolutional layers to 5 and hidden dimension to 300; for the SchNet model, we set the hidden dimension and the number of filters in continuous-filter convolution to 128. The interatomic distances are measured with 50 radial basis functions. We stack 6 interaction layers in the SchNet architecture. For the RoBERTa model, we use the default configuration. Particularly, the hidden dimension is set to 768 and the number of attention heads is set to 12, and the dropout ratio between hidden layers is set to 0.1. In addition, we add a MLP layer with the hidden dimension of 300 as the prediction head.

For the fingerprint encoder, we set the bit embedding dimension and the hidden dimension to 64 and 300, respectively. Moreover, we adopt a one-layer Transformer model with 8 attention heads. For the representation aggregation module, similarly, we set the hidden dimension of all the trainable parameters to 300. For all activation functions, we stick to $\text{ReLU}(\cdot) = \max(0, \cdot)$. For the temperature parameter τ in the contrastive objective, we tune it from $\{0.1, 0.3, 0.5\}$. All these hyperparameters except for those in GIN and SchNet backbones are selected based on grid search on the validation set.

Moreover, we empirically find that the dropout ratio [111] and the learning rate are two important hyperparameters in our model. Due to different training dynamics of different view encoders, we do a hyperparameter search of the learning rates and dropout ratio for each encoder from $\{10^{-3}, 10^{-4}, \dots, 10^{-7}\}$ and $\{0, 0.3, 0.5\}$, respectively. We would like to emphasize that we do not tune hyperparameters for different downstream datasets.

C.3 Experimental Configurations

To obtain the SMILES embeddings, the RoBERTa backbone is firstly pretrained for 4 epochs with all of the available molecules in the GEOM-Drugs dataset. Then, during the pretraining stage, we set train the whole model for 100 epochs with a batch size of 256 and we ensure that the model converges.

In the fine-tuning stage, for fair comparison, we follow the scaffold split adopted by GraphMVP [12] and repeat the experiments for three times. Specifically, we fine-tune our model on downstream tasks with three different random splits and we report the averaged test performance with the best models obtained on the validation set for all the methods.

C.4 Computational Complexity Analysis

In MOCO, calculating featurizations (i.e. fingerprints, SMILES strings, and 2D/3D graphs) is very efficient: computing fingerprints based on SMILES strings can be done in an instant. For 3D structures, there are two types of conformations involved: for pretraining, we directly use the DFT-calculated conformations shipped with the dataset; for fine-tuning, as no such conformation is available, we sample conformation with RDKit which is efficient (≈ 0.01 sec/molecule). Note that it is costly to use DFT to optimize molecular conformation. An example is given by the OpenCatalyst project⁶, which reports that a standard relaxation using DFT takes 8 to 10 hours. In terms of pretraining time, our MOCO takes slightly longer than GraphMVP’s. To be specific, our pretraining on GEOM-Drugs takes $\approx 4\text{h}17\text{m}$ to finish whereas the baseline GraphMVP takes $\approx 4\text{h}03\text{m}$.

D Additional Experiments

In this section, we further perform experiments on non-quantum property regression tasks. We also include additional experiments with ablated models and model variants.

D.1 More Experiments on Molecular Property Regression Tasks

We further conduct experiments on four additional regression tasks for molecular property prediction, where the results are presented in Table 7. It can be clearly seen from the table that our MOCO considerably improves the performance of baselines on three datasets and achieves similar performance to the baseline approaches on the Lipophilicity dataset, which once again verifies the effectiveness of our framework and demonstrates the importance of integrating different molecular featurization techniques.

⁶<https://opencatalystproject.org/challenge.html>

Table 7: Additional results on four molecular property regression tasks in terms of Root-Mean-Square Error (RMSE). The lowest prediction error is highlighted in **boldface**.

Pretraining	ESOL	Lipophilicity	Malaria	CEP	Avg.
—	1.364 \pm 0.016	0.736 \pm 0.006	1.122 \pm 0.011	1.380 \pm 0.033	1.15051
AttrMask	1.112 \pm 0.048	0.730 \pm 0.004	1.119 \pm 0.014	1.256 \pm 0.000	1.05419
ContextPred	1.196 \pm 0.037	0.702\pm0.020	1.101 \pm 0.015	1.243 \pm 0.025	1.06059
JOAO	1.120 \pm 0.019	0.708 \pm 0.007	1.145 \pm 0.010	1.293 \pm 0.003	1.06631
GraphMVP	1.091 \pm 0.021	0.718 \pm 0.016	1.114 \pm 0.013	1.236 \pm 0.023	1.03968
MOCO	0.984\pm0.034	0.707 \pm 0.001	1.093\pm0.009	1.101\pm0.007	0.97125

Table 8: Results for alternative featurizations in terms of ROC-AUC (% , \uparrow).


Featurization	BBBP	Tox21	SIDER	ClinTox	BACE	MUV	HIV	Toxcast	Avg.
SMILES	70.7	76.0	58.7	81.0	81.6	73.2	73.5	64.4	72.4
2D	65.4	74.9	58.0	58.8	72.6	71.0	75.3	61.6	67.2
3D	61.7	74.0	59.6	69.9	81.7	73.7	67.9	64.4	69.1
Fingerprint	66.3	73.4	51.0	64.4	76.3	76.9	70.5	61.1	67.5
MOCO	71.6	76.7	61.2	81.6	82.6	78.5	78.3	64.9	74.4

D.2 Ablation Studies on Multiple Views

To further verify the necessity of incorporating all four views (2D, 3D, FP and SM), we further compare our MOCO model with four featurization techniques with its ablated counterparts. Figure 5 summarizes the model performance obtained with only three views and Table 8 demonstrates the model performance obtained with single view. From the figure, it is observed that our MOCO that adaptively learns to optimize the combinations of four featurization techniques for different downstream tasks achieves the best performance for almost all datasets, which demonstrates the necessity of comprehensively considering four featurization. We kindly note that our framework is general and flexible; it is thus not limited to specifically incorporate these four views.

D.3 Experiments on Variant Fingerprint Encoders

Since there is a lack of proper neural encoders for fingerprints, we propose an attention-based network to model interactions of feature fields in fingerprint vectors, which considers the discrete and extremely sparse nature of fingerprints. To demonstrate its effectiveness, we leverage a simple MLP and a Long Short-Term Memory (LSTM) model as the encoder for the fingerprints respectively. The performance comparison is summarized in Figure 6. It is clear that our proposed attention encoder that models feature interactions for fingerprints achieves the best performance.

**Figure 5:** Ablation studies on four featurization techniques used in MOCO.




Figure 6: Results with two variant encoders for fingerprints: MLP and LSTM.

Table 9: Results for alternative featurizations in terms of ROC-AUC (% , \uparrow).

Featurization	BBBP	Tox21	SIDER	ClinTox	BACE	MUV	HIV	Toxcast	Avg.
SELFIES	70.6	76.5	62.0	79.8	82.1	77.5	78.4	64.1	73.9
Fragprints	68.1	75.0	58.3	78.6	71.1	77.0	77.1	63.4	71.1
MOCO	71.6	76.7	61.2	81.6	82.6	78.5	78.3	64.9	74.4

D.4 Experiments Using Alternative Featurizations

We further consider two more featurization techniques published in recent years: SELFIES [112], which is based on string representations and Fragprints [113], which is a path-based substructural fingerprints. We replace SMILES and fingerprints with SELFIES and Fragprints respectively and keep other experimental protocols the same. The results on molecular property prediction are shown in Table 9.

D.5 Explanation for Pre-training Data Scale

We chose 50K data for pretraining mainly for two reasons. Firstly, to ensure a fair comparison, we strictly followed the experimental settings of GraphMVP for pretraining. It is worth noting that our data selection method is random, rather than based on specific criteria, which would not lead to distribution shift. Moreover, all baseline methods were pre-trained using the same size and completely identical data and then fine-tuned, so we have maximally ensured the fairness of the comparison. Secondly, whether different pretraining data sizes are really effective in molecular pretraining research is still an issue that needs to be explored in depth. GraphMVP claim that “with 50K pretraining data scale, generally we can match with the original paper, even though most of them are using larger pre-training datasets, like ZINC-2m”. Moreover, previous study [114] has shown that increasing the data size does not bring significant gains, and they give several reasons of why the self-supervised pretraining may not be very effective. However, while ensuring fairness of comparison, we used a smaller pretraining data size to save computational costs and achieved promising results. We also conducted additional pre-training on a larger scale, and the results are shown in the Table 10.

E More Related Work

The following section provides a more broad literature review across the spectrum of self-supervised representation learning.

Table 10: Results for alternative featurizations in terms of ROC-AUC (% , \uparrow).

Featurization	BBBP	Tox21	SIDER	ClinTox	BACE	MUV	HIV	Toxcast	Avg.
MOCO-100K	70.4	76.2	62.6	83.7	80.0	80.4	77.6	64.9	74.4
MOCO	71.6	76.7	61.2	81.6	82.6	78.5	78.3	64.9	74.4

E.1 Self-Supervised Representation Learning on Visual and Natural Language Data

A SSL model trains itself by learning a part of the input from another through pretext tasks. Depending on the pretext task, the existing SSL studies can be divided into three main categories.

Early SSL work studies *predictive training* on pseudo-labels directly computed from the raw data. In Computer Vision (CV) domains, typical pretext tasks include image inpainting [115], rearranging shuffled image patches [116], colorizing grayscale images [117, 118], recognizing geometric transformations [119], and predicting cluster assignments [120]. In Natural Language Processing (NLP), word2vec [121] popularizes this paradigm by proposing Continuous Bag-Of-Words (CBOW) and skip-gram models for predicting center and neighboring words, respectively. Other exemplary work includes Kiros et al. [122] that predicts neighborhood sentences and BART [123] that recovers sentence permutation.

The second group of SSL is *contrastive learning*, which seeks to maximize the agreement of embeddings in the latent space under stochastic data augmentations by contrasting positive and negative samples [4]. It has revolutionized unsupervised representation learning in recent years [34–36, 124–127] and has been witnessed to perform on par with its supervised counterparts [34, 35]. A key success to contrastive models is to leverage strong data augmentations that induce invariance irrelevant to properties of the end tasks [128–131]. The third line of development focuses on *generative modeling* of input data. Its core idea is to randomly remove a portion of data and train the model to recover the removed content. This so-called masked language modeling and its autoregressive counterparts are first pioneered in the NLP community [13, 28, 29, 132–136] and have since gained increasing popularity in the CV domain [137–139]. Unlike contrastive learning, generative approaches do not rely on curated data augmentations. It has been reported that they scale well and generalize to different downstream tasks [29, 136, 138].

E.2 Graph Self-Supervised Representation Learning

Analogous to the above studies on visual and natural language data, SSL approaches in the graph domain can also be organized into the same three categories. Due to the rapid development of graph SSL, we only review the most representative studies in each group. Readers may refer to recent surveys [140–142] for comprehensive reviews and Zhu et al. [38] for a benchmarking study.

Firstly, the pioneering *predictive* model Hu et al. [7] explores four strategies at both node and graph levels, including masked attribute prediction, context prediction, supervised attribute prediction, and structural similarity prediction. You et al. [143] study three SSL tasks through a multi-task framework to enable predictive training of graph-structured data. M3S [144] explores the use of cluster assignments [120] as pseudo-labels and proposes a self-training framework that incrementally adds high-confident nodes to the labeled dataset.

The second group of work studies *generative training*. GraphSAGE [42] performs the link prediction task to reconstruct the graph structure in a once-for-all manner, similar to graph autoencoders [145]. GPT-GNN [44] proposes to perform node and edge reconstruction iteratively.

Lastly, along the line of graph *contrastive learning*, some investigate contrasting modes for graph data, typical work of which includes cross-scale contrasting [146, 147], same-scale contrasting [8, 37, 43], and hierarchical contrasting [9, 148]. Another line of work investigates data augmentations. GraphCL [8] proposes four heuristic augmentation schemes including edge dropping, node dropping, attribute masking, and subgraph cropping; its follow-up JOYO [45] proposes to learn the augmentation priors via bi-level optimization. GCA [149] proposes adaptive augmentation that better preserves important semantics and structures of the underlying graph. SimGRACE [150] eschews the need of explicit augmentation; Trivedi et al. [151] propose content-aware augmentation to avoid corrupting task-relevant information.