

# Molformer: Motif-Based Transformer on 3D Heterogeneous Molecular Graphs

Fang Wu<sup>1,3</sup>, Dragomir Radev<sup>2</sup>, Stan Z. Li<sup>1\*</sup>

<sup>1</sup> School of Engineering, Westlake University

<sup>2</sup> Department of Computer Science, Yale University

<sup>3</sup> Institute of AI Industry Research, Tsinghua University

fw2359@columbia.edu, dragomir.radev@yale.edu, stan.zq.li@westlake.edu.cn

## Abstract

Procuring expressive molecular representations underpins AI-driven molecule design and scientific discovery. The research mainly focuses on atom-level homogeneous molecular graphs, ignoring the rich information in subgraphs or motifs. However, it has been widely accepted that substructures play a dominant role in identifying and determining molecular properties. To address such issues, we formulate heterogeneous molecular graphs (HMGs), and introduce a novel architecture to exploit both molecular motifs and 3D geometry. Precisely, we extract functional groups as motifs for small molecules and employ reinforcement learning to adaptively select quaternary amino acids as motif candidates for proteins. Then HMGs are constructed with both atom-level and motif-level nodes. To better accommodate those HMGs, we introduce a variant of the Transformer named Molformer, which adopts a heterogeneous self-attention layer to distinguish the interactions between multi-level nodes. Besides, it is also coupled with a multi-scale mechanism to capture fine-grained local patterns with increasing contextual scales. An attentive farthest point sampling algorithm is also proposed to obtain the molecular representations. We validate Molformer across a broad range of domains, including quantum chemistry, physiology, and biophysics. Extensive experiments show that Molformer outperforms or achieves the comparable performance of several state-of-the-art baselines. Our work provides a promising way to utilize informative motifs from the perspective of multi-level graph construction. The code is available at <https://github.com/smiles724/Molformer>.

## Introduction

The past decade has witnessed the extraordinary success of deep learning (DL) in many scientific areas. Inspired by these achievements, researchers have shown increasing interest in exploiting DL for drug discovery and material design with the hope of rapidly identifying desirable molecules. A key aspect of fast screening is how to represent molecules effectively, where graphs are a natural choice to preserve their internal structures. As a consequence, a number of *graph neural networks* (GNNs) (Gilmer et al. 2017; Ishida et al. 2021) have been invented and applied to molecular representation learning with noticeable performance.

\*The corresponding author.

However, most existing GNNs only take atom-level information in homogeneous molecular graphs as input, which fails to adequately exploit rich semantic information in motifs. Remarkably, motifs are significant subgraph patterns that frequently occur, and can be leveraged to uncover molecular properties. For instance, a carboxyl group (*COOH*) acts as a hydrogen-bond acceptor, which contributes to better stability and higher boiling points. Besides, similar to the role of N-gram in natural language, molecular motifs can promote the segmentation of atomic semantic meanings. While some regard motifs as additional features of atoms (Maziarka et al. 2020, 2021), these methods increase the difficulty to separate the semantic meanings of motifs from atoms explicitly and hinder models from viewing motifs from an integral perspective. Others (Huang et al. 2020) take motifs as the only input, but ignore the influence of single atoms and infrequent substructures.

To overcome these problems, we formulate a novel heterogeneous molecular graph (HMG) that is comprised of both atom-level and motif-level nodes as the model input. It provides a clean interface to incorporate nodes of different levels and prevents the error propagation caused by incorrect semantic segmentation of atoms. As for the determination of motifs, we adopt different strategies for different types of molecules. On the one hand, for small molecules, the motif lexicon is defined by functional groups based on chemical domain knowledge. On the other hand, for proteins that are constituted of sequential amino acids, a reinforcement learning (RL) motif mining technique is introduced to discover the most meaningful amino acid subsequences for downstream tasks.

In order to better align with HMGs, we present Molformer, an equivariant geometric model based on Transformer (Vaswani et al. 2017). Molformer differs from preceding Transformer-based models in two major aspects. First, it utilizes heterogeneous self-attention (HSA) to distinguish the interactions between nodes of different levels and incorporates them into the self-attention computation. Second, an Attentive Farthest Point Sampling (AFPS) algorithm is introduced to aggregate node features and obtain a comprehensive representation of the entire molecule.

To summarize, our contributions are as follows:

- To the best of our knowledge, we are the foremost to incorporate motifs and construct 3D heterogeneous molec-

ular graphs for molecular representation learning.

- We propose a novel Transformer architecture to perform on these heterogeneous molecular graphs. It has a modified self-attention to take into account the interactions between multi-level nodes, and an AFPS algorithm to integrate molecular representations.
- We empirically outperform or achieve competitive results compared to state-of-the-art baselines on several benchmarks of small molecules and proteins.

## Preliminaries

**Problem Definition.** Suppose a molecule  $\mathcal{S} = (\mathbf{P}, \mathbf{H})$  has  $N$  atoms, where  $\mathbf{P} = \{\mathbf{p}_i\}_{i=1}^N \in \mathbb{R}^{N \times 3}$  describes 3D coordinates associated to each atom and  $\mathbf{H} = \{\mathbf{h}_i\}_{i=1}^N \in \mathbb{R}^{N \times h}$  contains a set of  $h$ -dimension roto-translationally invariant features (e.g. atom types and weights).  $\mathbf{h}_i$  can be converted to a dense vector  $\mathbf{x}_i \in \mathbb{R}^{\psi_{\text{embed}}}$  via a multi-layer perceptron (MLP). A representation learning model  $f$  acts on  $\mathcal{S}$ , and obtains its representation  $\mathbf{r} = f(\mathcal{S})$ . Then  $\mathbf{r}$  is forwarded to a predictor  $g$  and attains the prediction of a biochemical property  $\hat{y} = g(\mathbf{r})$ .

**Self-Attention Mechanism.** Given input  $\{\mathbf{x}_i\}_{i=1}^N$ , the standard dot-product self-attention layer is computed as follows:

$$\mathbf{q}_i = f_Q(\mathbf{x}_i), \mathbf{k}_i = f_K(\mathbf{x}_i), \mathbf{v}_i = f_V(\mathbf{x}_i) \quad (1)$$

$$a_{ij} = \mathbf{q}_i \mathbf{k}_j^T / \sqrt{\psi_{\text{model}}}, \mathbf{z}_i = \sum_{j=1}^N \sigma(a_{ij}) \mathbf{v}_j \quad (2)$$

where  $\{f_Q, f_K, f_V\}$  are embedding transformations, and  $\{\mathbf{q}_i, \mathbf{k}_i, \mathbf{v}_i\}$  are respectively the query, key, and value vectors with the same dimension  $\psi_{\text{model}}$ .  $a_{ij}$  is the attention that the token  $i$  pays to the token  $j$ .  $\sigma$  denotes the *Softmax* function and  $\mathbf{z}_i$  is the output embedding of the token  $i$ . This formula conforms to a non-local network (Wang et al. 2018), which indicates its poor ability to capture fine-grained patterns in a local context.

**Position Encoding.** Self-attention is invariant to the permutation of the input, and position encoding (PE) is the only technique to reveal position information. PE can be based on absolute positions or relative distances. The former uses raw positions and is not robust to spatial transformations. The latter manipulates the attention score by incorporating relative distances:  $a_{ij} = \mathbf{q}_i \mathbf{k}_j^T / \sqrt{\psi_{\text{model}}} + f_{\text{PE}}(\mathbf{p}_i - \mathbf{p}_j)$ , where  $f_{\text{PE}}$  is a translation-invariant PE function. The rotation invariance can be further accomplished by taking a L2-norm  $d_{ij} = \|\mathbf{p}_i - \mathbf{p}_j\|_2$  between the  $i$ -th and  $j$ -th atom.

## Heterogeneous Molecular Graphs

Motifs are frequently-occurring substructure patterns and serve as the building blocks of complex molecular structures. They have great expressiveness of the biochemical characteristics of the whole molecules. In this section, we first describe how to extract motifs from small molecules and proteins respectively, and then present how to formulate HMGs.

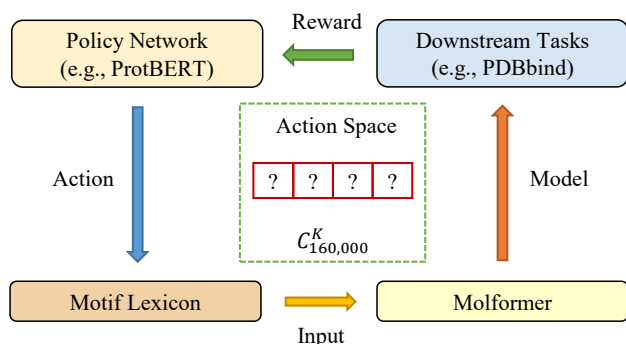


Figure 1: The workflow of RL motif mining method in proteins. In each iteration, the policy network is responsible for producing a motif lexicon. Molformer’s performance on downstream tasks and the diversity of the lexicon are considered as the reward.

## Motifs in Small Molecules

In the chemical community, a set of standard criteria have been developed to recognize motifs with essential functionalities in small molecules. There, we build motif templates of four categories of functional groups, containing groups that contain only carbon and hydrogen (Hydrocarbons), groups that contain halogen (Haloalkanes), groups that contain oxygen, and groups that contain nitrogen (see Appendix for more discussion). Practically, we rely on RD-Kit to draw them from SMILES of small molecules.

## Motifs in Proteins

In large protein molecules, motifs are local regions of 3D structures or amino acid sequences shared among proteins that influence their functions (Somnath, Bunne, and Krause 2021). Each motif usually consists of only a few elements, such as the ‘helix-turn-helix’ motif, and can describe the connectivity between secondary structural elements. The detection of protein motifs has been long studied. Nevertheless, existing tools are either from the context of a protein surface (Somnath, Bunne, and Krause 2021) or are task-independent and computationally expensive (Mackenzie, Zhou, and Grigoryan 2016). On the basis of this peculiarity, we design an RL mining method to discover task-specific protein motifs heuristically.

Therefore, we consider motifs with four amino acids because they make up the smallest polypeptide and have special functionalities in proteins. For instance,  $\beta$ -turns are composed of four amino acids and are a non-regular secondary structure that causes a change in the direction of the polypeptide chain. Each amino acid can be of 20 different possibilities, such as Alanine, Isoleucine, and Methionine, so there are  $1.6 \times 10^5 (= 20^4)$  potential quaternary motifs.

Our goal is to find the most effective lexicon  $\mathcal{V}^* \in \mathcal{V}$  composed of  $K$  quaternary amino acid templates, where  $\mathcal{V}$  denotes the space of all  $C_{1.6 \times 10^5}^K$  potential lexicons. Since we aim to mine the optimal task-specific lexicon, it is practically feasible to only consider the existing quaternions in

the downstream datasets instead of all  $1.6 \times 10^5$  possible quaternions.

In each iteration of the parameter update, we use a pre-trained ProtBERT (Elnaggar et al. 2020) with an MLP as the policy network  $\pi_\theta$ . Specifically, all possible quaternions are fed into ProtBERT to obtain their corresponding representations  $\{e_i\}_{i=1}^{1.6 \times 10^5}$ , which are subsequently sent to the MLP to acquire their scores  $\{s_i\}_{i=1}^{1.6 \times 10^5}$ . These scores illustrate each quaternion’s significance to benefit the downstream tasks if they are chosen as a part of the vocabulary. Then top- $K$  motifs with the highest scores are selected to comprise  $\mathcal{V} \in \mathbb{V}$  in accordance with  $\{s_i\}_{i=1}^{1.6 \times 10^5}$ , and  $\mathcal{V}$  is used as templates to extract motifs and construct HMGs in downstream tasks. After that, a Molformer is trained based on these HMGs. Its validation performance is regarded as the reward  $r$  to update parameters  $\theta$  by means of policy gradients. With adequate iterations, the agent can select the optimal task-specific quaternary motif lexicon  $\mathcal{V}^*$ .

Remarkably, our motif mining process is a one-step game, since the policy network  $\pi_\theta$  only generates the vocabulary  $\mathcal{V}$  once in each iteration. Thus, the trajectory consists of only one action, and the performance of Molformer based on the chosen lexicon  $\mathcal{V}$  composes a part of the total reward. Moreover, we also consider the diversity of motif templates within the lexicon and calculate it as:

$$d_{\text{div}}(\mathcal{V}) = \frac{1}{|\mathcal{V}|} \sum_{\mathbf{m}_i \in \mathcal{V}} \sum_{\mathbf{m}_j \in \mathcal{V}} d_{\text{lev}}(\mathbf{m}_i, \mathbf{m}_j) \quad (3)$$

where  $d_{\text{lev}}$  is the Levenshtein distance of two quaternary sequences  $\mathbf{m}_i$  and  $\mathbf{m}_j$ . The final reward, therefore, becomes  $R(\mathcal{V}) = r + \gamma d_{\text{div}}(\mathcal{V})$ , where  $\gamma$  is a weight to balance two reward terms, and the policy gradient is computed as the following objective:

$$\nabla_\theta J(\theta) = \mathbb{E}_{\mathcal{V} \in \mathbb{V}} [\nabla_\theta \log \pi_\theta(\mathcal{V}) R(\mathcal{V})] \quad (4)$$

## Formulation of Heterogeneous Molecular Graphs

Most prior studies (Maziarka et al. 2020; Rong et al. 2020; Maziarka et al. 2021) simply incorporate motifs into atom features. For instance, they differentiate carbons into aromatic or non-aromatic and deem them as extra features. We argue its ineffectiveness for two reasons. First, the fusion of multi-level features increases the difficulty to summarize the functionality of motifs. Second, it hinders models to see motifs from a unitary perspective. To fill these gaps, we separate apart motifs and atoms, regarding motifs as new nodes to formulate an HMG. This way disentangles motif-level and atom-level representations, thus alleviating the difficulty for models to properly mine the motif-level semantic meanings.

Similar to the relation between phrases and single words in natural language, motifs in molecules carry higher-level semantic meanings than atoms. Therefore, they play an essential part in identifying the functionalities of their atomic constituents. Inspired by the employment of dynamic lattice structures in named entity recognition, we treat each category of motifs as a new type of node and build HMGs as the input of our Molformer. To begin with, motifs are extracted according to a motif vocabulary  $\mathcal{V}$ . We assume  $M$

motifs  $\{\mathbf{m}_i\}_{i=1}^M$  are detected in the molecule  $\mathcal{S}$ . Consequently, an HMG includes both the motif-level and atom-level nodes as  $\{\mathbf{x}_1, \dots, \mathbf{x}_N, \mathbf{x}_{m_1}, \dots, \mathbf{x}_{m_M}\}$ , where  $\mathbf{x}_{m_i} \in \mathbb{R}^{\psi_{\text{embed}}}$  is obtained through a learnable embedding matrix  $\mathbf{W}^M \in \mathbb{R}^{C' \times \psi_{\text{embed}}}$  and  $C'$  denotes the number of motif categories. As for positions of each motif, we adopt a weighted sum of 3D coordinates of its components as  $\mathbf{p}_{m_i} = \sum_{\mathbf{x}_i \in \mathbf{m}_i} \left( \frac{w_i}{\sum_{\mathbf{x}_i \in \mathbf{m}_i} w_i} \right) \cdot \mathbf{p}_i$ , where  $w_i$  are the atomic weights. Analogous to word segmentation, our HMGs composed of multi-level nodes avoid error propagation due to inappropriate semantic segmentation while leveraging the atom information for molecular representation learning.

## Molformer

In this section, we propose Molformer, which modifies Transformer with several novel components specifically designed for 3D HMGs. First, we build a motif vocabulary and match each molecule with this lexicon to obtain all its contained motifs. Then both atoms and motifs acquire their corresponding embeddings and are forwarded into  $L$  feature learning blocks. Each block consists of an HSA, a feed-forward network (FFN), and two-layer normalizations. After that, an AFPS is followed to adaptively produce the molecular representation, which is later fed into a dense predictor to forecast properties in a broad range of downstream tasks.

## Heterogeneous Self-attention

After formulating an HMG with  $N$  atom-level nodes and  $M$  motif-level nodes, it is important to endow the model with the capacity of separating the interactions between multi-order nodes. To this end, we utilize a function  $\phi(i, j) : \mathbb{R}^{(N+M) \times (N+M)} \rightarrow \mathbb{Z}$ , which identifies the relationship between any two nodes into three sorts: atom-atom, atom-motif, and motif-motif. Then a learnable scalar  $b_{\phi(i, j)} : \mathbb{Z} \rightarrow \mathbb{R}$  indexed by  $\phi(i, j)$  is introduced so that each node can adaptively attend to all other nodes according to their hierarchical relationship inside our HMGs.

In addition, we consider exploiting 3D molecular geometry (see Figure 2). Since robustness to global changes such as 3D translations and rotations is an underlying principle for molecular representation learning, we seek to satisfy roto-translation invariance. There, we borrow ideas from SE(3)-Transformer (Schütt, Unke, and Gastegger 2021) and AlphaFold2 (Jumper et al. 2021), and apply a convolutional operation to the pairwise distance matrix  $\mathbf{D} = [d_{i, j}]_{i, j \in [N+M]} \in \mathbb{R}^{(N+M) \times (N+M)}$  as  $\hat{\mathbf{D}} = \text{Conv}_{2d}(\mathbf{D})$ , where  $\text{Conv}_{2d}$  denotes a 2D shallow convolutional network with a kernel size of  $1 \times 1$ . Consequently, the attention score is computed as follows:

$$\hat{a}_{ij} = \left( \mathbf{q}_i \mathbf{k}_j^T / \sqrt{\psi_{\text{model}}} \right) \cdot \hat{d}_{ij} + b_{\phi(i, j)}, \quad (5)$$

where  $\hat{d}_{i, j} \in \hat{\mathbf{D}}$  controls the impact of interatomic distance over the attention score, and  $b_{\phi(i, j)}$  is shared across all layers.

Moreover, exploiting local context has proven to be important in sparse 3D space. However, it has been pointed out

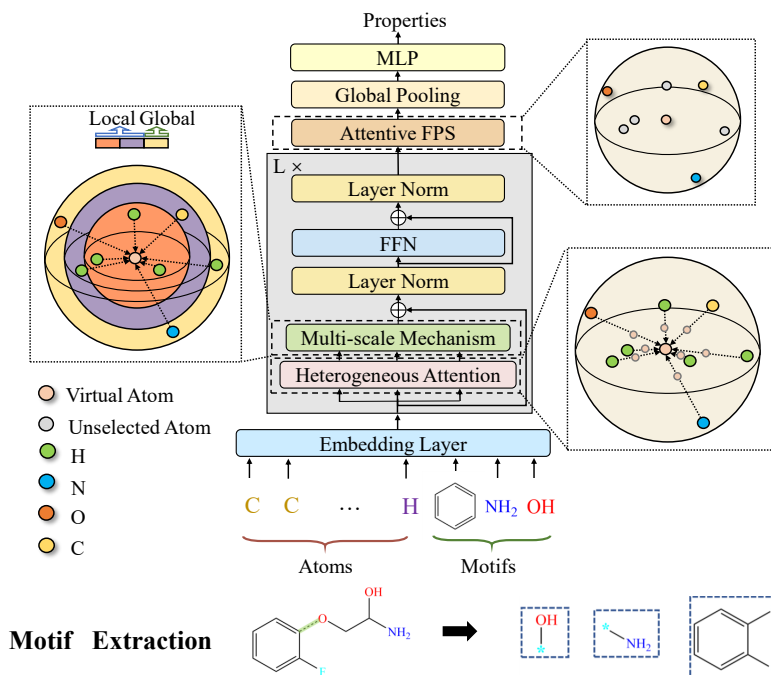


Figure 2: The architecture of Molformer. Given a heterogeneous molecular graph with both atom-level and motif-level nodes, stacked feature learning blocks composed of a heterogeneous self-attention module and an FFN compute their updated features. Afterward, an attentive subsampling module integrates the molecular representation for downstream predictions. Local features with different scales are purple and orange; yellow corresponds to global features.

that self-attention is good at capturing global data patterns but ignores local context (Wu et al. 2020). Based on this fact, we impose a distance-based constraint on self-attention in order to extract multi-scaled patterns from both local and global contexts. Guo et al. (2020) propose to use integer-based distances to limit attention to local word neighbors, which cannot be used in molecules. This is because different types of molecules have different densities and molecules of the same type have different spatial regularity, which results in the non-uniformity of interatomic distances. Normally, small molecules have a mean interatomic distance of 1-2 Å (Angstrom,  $10^{-10}m$ ), which is denser than large molecules like proteins with approximately 5 Å on average. To tackle that, we design a multi-scale methodology to robustly capture details. Specifically, we mask nodes beyond a certain distance  $\tau_s$  at each scale  $s$ . The attention calculation is modified as:

$$a_{ij}^{\tau_s} = \hat{a}_{ij} \cdot \mathbf{1}_{\{d_{ij} < \tau_s\}}, \mathbf{z}_i^{\tau_s} = \sum_{j=1}^N \sigma(a_{ij}^{\tau_s}) \mathbf{v}_j, \quad (6)$$

where  $\mathbf{1}_{\{d_{ij} < \tau_s\}}$  is the indicator function. Notably, Equation 6 can be complementarily combined with Equation 5. Then, features extracted from  $S$  different scales  $\{\tau_s\}_{s=1}^S$ , as well as the informative global feature, are concatenated together to form a multi-scale representation, denoted by  $\mathbf{z}'_i = \mathbf{z}_i^{\tau_1} \oplus \dots \oplus \mathbf{z}_i^{\tau_S} \oplus \mathbf{z}_i^{global} \in \mathbb{R}^{(S+1)\psi_{model}}$ . After that,  $\mathbf{z}'_i$  is forwarded into a FFN to obtain  $\mathbf{z}''_i$  with the original dimension  $\psi_{model}$ .

## Attentive Farthest Point Sampling

After having the node embeddings  $\{\mathbf{z}''_i\}_{i=1}^{N+M}$ , we study how to obtain the molecular representation  $\mathbf{r}$ . For GNNs, several readout functions such as set2set (Vinyals, Bengio, and Kudlur 2015) and GG-NN (Gilmer et al. 2017) are invented. For Transformers, one way is via a virtual node. Though Ying et al. (2021) state that it significantly improves the performance of existing models in the leaderboard of Open Graph Benchmark (OGB), this way concentrates more on its close adjacent nodes and less on distant ones, and may lead to inadvertent over-smoothing of information propagation. Besides, it is difficult to locate a virtual node in 3D space and build connections to existing vertices. The other way selects a subset of nodes via a downsampling algorithm named Farthest Point Search (FPS), but it ignores nodes' differences and has the sensitivity to outlier points as well as uncontrollable randomness. To address these issues, we propose a new algorithm named AFPS. It aims to sample vertices by not merely spatial distances, but also their significance in terms of attention scores.

Specifically, we choose the virtual atom  $\mathbf{x}_\#$  as the starting point and initialize two lists  $\mathcal{P} = \{\mathbf{x}_\#\}$  and  $\mathcal{M} = \{\mathbf{x}_i\}_{i=1}^{N+M}$  to store remaining candidate points. Then the process begins with the attention score matrix  $\hat{\mathbf{A}} = [\hat{a}_{i,j}]_{i,j \in [N+M]} \in \mathbb{R}^{(N+M) \times (N+M)}$  and the interatomic distance matrix  $\mathbf{D} \in \mathbb{R}^{(N+M) \times (N+M)}$ . It can be easily proved that each row of  $\hat{\mathbf{A}}$  sums up to 1 after the *Softmax* operation along columns, i.e.  $\sum_j \hat{a}_{ij} = 1, \forall i \in [N+M]$ .

---

**Algorithm 1: Attentive Farthest Point Sampling**

---

**Input:** The attention score matrix  $\mathbf{A} \in \mathbb{R}^{(N+M) \times (N+M)}$ , a Euclidean distance matrix  $\mathbf{D} \in \mathbb{R}^{(N+M) \times (N+M)}$ .  
**Output:**  $K$  sampled points.  
 $\tilde{\mathbf{A}} \leftarrow \sum_i \hat{a}_{ij} \in \mathbb{R}^{N+M}$   $\triangleright$  sum up along rows  
 $\tilde{\mathbf{D}} \leftarrow \frac{\mathbf{D} - \min \mathbf{D}}{\max \mathbf{D} - \min \mathbf{D}} \in \mathbb{R}^{(N+M) \times (N+M)}$   $\triangleright$  normalize the distance matrix  
 $\mathcal{P} = \{\mathbf{x}_\#\}$   
 $\mathcal{M} = \{\mathbf{x}_i\}_{i=1}^{N+M}$   
**while** length( $\mathcal{P}$ )  $< k$  **do**  
     $\mathbf{x}_{\text{new}} \leftarrow \operatorname{argmax}_{i \in \mathcal{M}} (\min_{j \in \mathcal{P}} \tilde{D}_{ij} + \lambda \tilde{A}_i)$   $\triangleright$  pick up the node that maximize the objective  
     $\mathcal{P}.\text{append}(\mathbf{x}_{\text{new}})$   
     $\mathcal{M}.\text{remove}(\mathbf{x}_{\text{new}})$   
**end while return**  $\mathcal{P}$

---

In order to obtain the importance of each atom in the self-attention computation, we accumulate  $\tilde{\mathbf{A}}$  along rows and get  $\tilde{\mathbf{A}} = \sum_i \hat{a}_{ij} \in \mathbb{R}^{N+M}$ . Besides, we adopt the min-max normalization to rescale the distance matrix  $\mathbf{D}$  into values between 0 and 1, and obtain  $\tilde{\mathbf{D}} = \frac{\mathbf{D} - \min \mathbf{D}}{\max \mathbf{D} - \min \mathbf{D}}$ .

After the above preprocess, we repeatedly move a point  $\mathbf{x}_{\text{new}}$  from  $\mathcal{M}$  to  $\mathcal{P}$ , which ensures that  $\mathbf{x}_{\text{new}}$  is as far from  $\mathcal{P}$  as possible by maximizing  $\tilde{D}_{ij}$  and also plays a crucial role in attention computation by maximizing  $\tilde{A}_i$ . Mathematically, the AFPS aims to achieve the objective:

$$\max_{i \in \mathcal{M}} \sum_{j \in \mathcal{P} \setminus \{i\}} (\min_{j \in \mathcal{P} \setminus \{i\}} \tilde{D}_{ij} + \lambda \tilde{A}_i), \quad (7)$$

where  $\lambda$  is a hyperparameter to balance those two different goals. This process is repeated until  $\mathcal{P}$  has reached  $K$  points. Algorithm 1 provides a greedy approximation solution to this AFPS optimization objective for sake of computational efficiency.

After that, sampled features  $\{\mathbf{z}_i''\}_{i \in \mathcal{P}}$  are gathered by a Global Average Pooling layer to attain the molecular representation  $\mathbf{r} \in \mathbb{R}^{\psi_{\text{model}}}$ .

Remarkably, our proposed AFPS has considerable differences and superiority over a body of previous hierarchical approaches (Eismann et al. 2021). Their subsampling operations are mainly designed for protein complexities, which often have uniform structures. To be specific, they hierarchically use alpha carbons as the intermediate set of points and aggregate information at the level of those carbons for the entire complex. However, the structures of small molecules have no such stable paradigm, and we provide a universal method to adaptively subsample atoms without any prior assumptions on the atom arrangement.

## Experiments

We conduct extensive experiments on 7 datasets of both small and large protein molecules from three different domains, including quantum chemistry, physiology, and biophysics. The appendix summarises statistical information of these 7 benchmark datasets, such as the number of tasks and

task types, the number of molecules and atom classes, the minimum and the maximum number of atoms, and the density (mean interatomic distances) of all molecules.

**Datasets.** We test Molformer on QM7 (Blum and Raymond 2009), QM8 (Ramakrishnan et al. 2015), QM9 (Ramakrishnan et al. 2014), BBBP (Martins et al. 2012), ClinTox (Gayvert, Madhukar, and Elemento 2016), and BACE (Subramanian et al. 2016)<sup>1</sup>. QM7 is a subset of GDB-13 with 7K molecules. QM8 and QM9 are subsets of GDB-17 with 22K and 133K molecules respectively. BBBP involves records of whether a compound carries the permeability property of penetrating the blood-brain barrier. ClinTox compares drugs approved through FDA and drugs eliminated due to toxicity during clinical trials. BACE is collected for recording compounds that could act as the inhibitors of human  $\beta$ -secretase 1 (BACE-1).

We also inspect Molformer’s ability to learn mutual relations between proteins and molecules on the PDBbind dataset (Wang et al. 2005). Following Townshend et al. (2020), we split protein-ligand complexes by protein sequence identity at 30%. As for the target, we predict  $pS = -\log(S)$ , where  $S$  is the binding affinity in Molar unit. In addition, we only use the pocket of each protein and put pocket-ligand pairs together as the input.

For QM9, we use the exact train/validation/test split as Townshend et al. (2020). For PDBbind, 90% of the data is used for training and the rest is divided equally between validation and test like Chen et al. (2019). For others, we adopt the scaffold splitting method with a ratio of 8:1:1 for train/validation/test as Rong et al. (2020). More implementing details can be found in Appendix.

**Baselines** For small molecules, baselines are as follows. TF\_Robust (Ramsundar et al. 2015) takes molecular fingerprints as the input. Weave (Kearnes et al. 2016), MPNN (Gilmer et al. 2017), SchNet (Schütt et al. 2018), MEGNet (Chen et al. 2019), GROVER (Rong et al. 2020), DMPNN (Yang et al. 2019), MGCN (Lu et al. 2019), AttentiveFP (Xiong et al. 2019), DimeNet (Klicpera, Groß, and Günnemann 2020), DimeNet++ (Klicpera et al. 2020), PaiNN (Schütt, Unke, and Gastegger 2021), and SphereNet (Liu et al. 2021) are graph convolutional models. Graph Transformer (Chen, Barzilay, and Jaakkola 2019), MAT (Maziarka et al. 2020), R-MAT (Maziarka et al. 2021), SE(3)-Transformer (Fuchs et al. 2020), and LieTransformer (Hutchinson et al. 2021) are Transformer-based Equivariant Neural Networks (ENNs) (Thomas et al. 2018).

For PDBbind, we choose seven baselines. DeepDTA (Öztürk, Özgür, and Ozkirimli 2018) and DeepAffinity (Karimi et al. 2019) take in pairs of ligand and protein SMILES as input. Cormorant (Anderson, Hy, and Kondor 2019) is an ENN that represents each atom by its absolute 3D coordinates. HoloProt (Somnath, Bunne, and Krause 2021) captures higher-level fingerprint motifs on the protein surface. SchNet, 3DCNN and 3DGCN (Townshend et al. 2020) are 3D methods.

---

<sup>1</sup>For BBBP, ClinTox, and BACE, we use RDKit to procure 3D coordinates from SMILES.

Method	QM7	QM8	BBBP	ClinTox	BACE
TF-Robust	120.6	0.024	0.860	0.765	0.824
Weave	94.7	0.022	0.837	0.823	0.791
MPNN	113.0	0.015	0.913	0.879	0.815
Schnet	74.2	0.020	0.847	0.717	0.750
DMPNN	105.8	0.014	0.919	0.897	0.852
MGCN	77.6	0.022	0.850	0.634	0.734
Attentive FP	126.7	0.028	0.908	0.933	0.863
Graph Transformer	47.8	0.010	0.913	-	0.880
MAT	102.8	-	0.728	-	0.846
R-MAT	68.6	-	0.746	-	0.871
GROVER <sub>large</sub>	89.4	0.017	0.911	0.884	0.858
Molformer	<b>43.2</b>	<b>0.009</b>	<b>0.926</b>	<b>0.937</b>	<b>0.884</b>

Table 1: For regression tasks in QM7 and QM8, lower MAE is better. For classification tasks in BBBP, ClinTox, and Bace, higher values are better.

## Overall Results on Benchmark Tasks

**Molecules.** Table 1 and Table 2 document the overall results on small molecules datasets, where the best performance is marked bold and the second best is underlined for clear comparison. It can be discovered that Molformer achieves the lowest MAE of 43.2 on QM7 and 0.009 on QM8, beating several strong baselines including DMPNN and Graph Transformer. Besides, Molformer offers competitive performance in all property regression tasks on QM9. Particularly, we outperform all Transformer-based ENNs, including SE(3)-Transformer and LieTransformer. As for classification problems, we surpass all baselines mostly by a fairly large margin.

**Proteins.** Table 3 reports the Root-Mean-Squared Deviation (RMSD), the Pearson correlation ( $R_p$ ), and the Spearman correlation ( $R_s$ ) on PDBbind. Molformer achieves the lowest RMSD and the best Pearson and Spearman correlations. As Wu et al. (2018) claim, appropriate featurization which holds pertinent information is significant for PDBbind. However, an important observation in our work is that deep learning approaches with the exploitation of 3D geometric information can perform better than conventional methods like DeepDTA and DeepAffinity which use a set of physicochemical descriptors but ignore 3D structures.

## Ablation Study and Discussion

**What Is the Effect of Each Component?** We investigate the effectiveness of each component of our Molformer in Table 4. It can be observed that HSA along with HMGs substantially boosts the model’s performance compared with the naive method that immediately adds 3D coordinates as the atom input feature. MAE declines from 132.2 to 46.5 in QM7 while decreasing from 0.0205 to 0.0097 in QM8. In addition, AFPS produces better predictions than the counterpart that utilizes the virtual node as the molecular representation (a case study of AFPS is in the Appendix). We also discover that the multi-scale mechanism significantly reduces RMSD from 50.1 to 46.5 on QM7, but its improvements in QM8 are much smaller. This phenomenon indicates that the

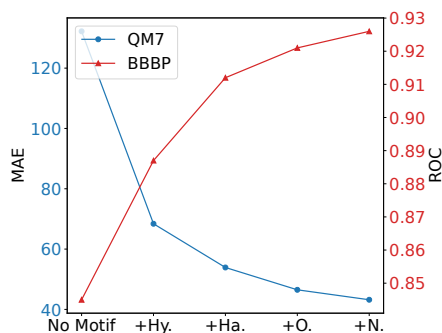


Figure 3: The ablation study of consecutively adding motif groups (from left to right) in the QM7 and BBBP datasets.

multi-scale mechanism is an appropriate way to alleviate the problem of inadequate training in small datasets. It endows Molformer with the capability to extract local features by regulating the scope of self-attention. However, as the data size gets larger, Molformer does not require the assistance of the multi-scale mechanism to abstract local patterns, since the parameters of convolution operators are properly trained.

**What Is the Contribution of HMGs?** The ideology of constructing heterogeneous graphs has achieved success in not only chemical knowledge graphs but named entity recognition (Li et al. 2020). The former views the chemical characteristics obtained from domain knowledge of elements as shared nodes, while the latter converts the lattice structure into a flat structure consisting of spans. To verify its efficacy, we compare our motif-based HMGs with the naive fusion of multi-level features. Table 5 shows a noticeable improvement in our HMGs over the other two variants.

**Have We Found Good Candidates of Motifs?** How to determine motifs is critical to HMGs. Concerning small molecules, we define motifs on the basis of functional groups, which refers to a substituent or moiety that causes molecules’ characteristic chemical reactions. To further explore their contributions, we divide functional groups into four categories: Hydrocarbons (Hy.), Haloalkanes (Ha.), groups that contain oxygen (O.), and groups that contain nitrogen (N.) and please refer to Appendix for more details. The ablation studies (see Figure 3) demonstrate that Molformer can gain improvements from all four groups of motifs, where Hydrocarbons and Haloalkanes are the most and the least effective types respectively. This is in line with the fact that Hydrocarbons occur most frequently in organic molecules. Moreover, the best performance is achieved when all categories are considered, implying a promising direction to discover more effective motifs. As for proteins, motifs discovered by our RL mining method share the same backbone as CC (C (NC (C) C (O) =O) =O) NC (CNC (CN) =O) =O (see Appendix), which is a hydrogen bond donor and implies a mark to distinguish potential binding site. Moreover, the portion of those motifs in the pocket (1.38%) is nearly twice that in other locations (0.73%), conforming to the fact that pockets are the most preferable part for ligands to bind with.

Target Unit	$\epsilon^{\text{HOMO}}$ eV	$\epsilon^{\text{LUMO}}$ eV	$\Delta\epsilon$ eV	$\mu$ D	$\alpha$ bohr <sup>3</sup>	$R^2$ $a_0^2$	ZPVE meV	$U_0$ meV	$U$ meV	$H$ meV	$G$ meV	$c_v$ cal/mol K
MPNN	.043	.037	.069	.030	.092	.150	1.27	45	45	39	44	.800
Schnet	.041	.034	.063	.033	.235	<u>.073</u>	1.7	14	19	14	14	.033
MEGNet <sub>full</sub>	.038	.031	.061	.040	.083	.265	1.4	9	10	10	10	.030
MGCN	.042	.057	.064	.056	<b>.030</b>	.110	<b>1.12</b>	12.9	14.4	14.6	16.2	.038
DimeNet	.027	.019	.034	.028	.046	.331	1.29	8.02	7.89	8.11	8.98	.024
DimeNet++	<u>.024</u>	<u>.019</u>	<u>.032</u>	<u>.029</u>	<u>.043</u>	.331	1.21	6.32	<u>6.28</u>	6.53	<u>7.56</u>	<u>.023</u>
SphereNet	<b>.024</b>	<b>.019</b>	<b>.032</b>	<u>.026</u>	<u>.047</u>	.292	<u>1.12</u>	<u>6.26</u>	<u>7.33</u>	<u>6.40</u>	<u>8.0</u>	<b>.021</b>
PaiNN	.028	.020	.046	<b>.012</b>	.045	<b>.066</b>	1.28	<b>5.85</b>	<b>5.53</b>	<b>5.98</b>	<b>7.35</b>	.024
SE(3)-Transformer	.035	.033	.053	.051	.142	-	-	-	-	-	-	-
LieTransformer	.033	.029	.052	.061	.104	2.29	3.55	17	16	27	23	.041
Molformer	.025	.026	.039	.028	.041	.350	2.05	7.52	7.46	7.38	8.11	.025

Table 2: Comparison of MAE on QM9. The last three methods are Transformer-based methods.

Method	Geometry	RMSD	$R_p$	$R_s$
DeepDTA	Non-3D	1.565	0.573	0.574
DeepAffinity	Non-3D	1.893	0.415	0.426
Schnet	3D	1.892	<u>0.601</u>	-
Cormorant	3D	<u>1.429</u>	0.541	0.532
3DCNN	3D	1.520	0.558	0.556
3DGCN	3D	1.963	0.581	<u>0.647</u>
HoloProt	3D	1.464	0.509	0.500
Molformer	3D	<b>1.386</b>	<b>0.623</b>	<b>0.651</b>

Table 3: Comparison of RMSD,  $R_p$ , and  $R_s$  on PDBbind.

	HSA	AFPS	HMG	QM7	QM8	PDBbind
1	-	-	-	132.2	0.0205	1.925
2	✓	-	✓	46.5	<u>0.0097</u>	1.441
3	✓	✓	✓	<b>43.2</b>	<b>0.0095</b>	<b>1.386</b>

Table 4: Effects of each module on QM7, QM8 and PDBbind (RMSD).

## Related Works

**Motifs in Molecular Graphs.** Motifs have been proven to benefit many tasks from exploratory analysis to transfer learning. Various algorithms have been proposed to exploit them for contrastive learning (Zhang et al. 2020), self-supervised pretraining (Zhang et al. 2021), generation (Jin, Barzilay, and Jaakkola 2020), protein design (Li et al. 2022) and drug-drug interaction prediction (Huang et al. 2020). But none of them take advantage of motifs to build a heterogeneous graph for molecular property prediction. Moreover, previous motif mining methods either depend on exact counting (Cantoni, Gatti, and Lombardi 2011) or sampling and statistical estimation (Wernicke 2006) to extract motifs. No prior studies extract task-specific motifs to enhance model performance.

**Molecular Representation Learning.** Molecules are usually represented as 1D sequences, including amino acid se-

	QM7	QM8	PDBbind
No Motifs	132.2	0.0205	1.925
Multi-Level Fusion	89.7	<u>0.0154</u>	<u>1.427</u>
Heterogeneous Graphs	<b>43.2</b>	<b>0.0095</b>	<b>1.386</b>

Table 5: Comparison of HMGs with simple feature fusion on QM7, QM8 and PDBbind (RMSD).

quences and SMILES (Xu et al. 2017), and 2D graphs (Duvinaud et al. 2015). Despite that, more evidence indicates that 3D structures lead to better modeling and superior performance. 3D CNNs (Anand-Achim et al. 2021) and GNNs (Cho and Choi 2018) become popular to capture these complex geometries. Nonetheless, the aforementioned methods are inefficient at grabbing local contextual feature and long-range dependencies.

Transformer attempts to address that issue, assuming full connection and using self-attention to capture long-term dependencies. Some researchers feed SMILES in Transformer to obtain representations (Honda, Shi, and Ueda 2019) and conduct pretraining (Chithrananda, Grand, and Ramsundar 2020). Others employ Transformer to solve generative tasks (Ingraham et al. 2019) or fulfill equivariance (Fuchs et al. 2020) via spherical harmonics. However, they are either incapable to encode 3D geometry, non-sensitive to local contextual patterns, or inefficient to aggregate atom features. More essentially, they are not specially designed to operate on heterogeneous graphs of molecules.

## Conclusion

This paper presents Molformer for 3D molecular representation learning on heterogeneous molecular graphs. First, we extract motifs to formulate heterogeneous molecular graphs. After that, Molformer adopts a heterogeneous self-attention to distinguish the interactions between multi-level nodes. Then a simple but efficient downsampling algorithm is introduced to better accumulate molecular representations. Experiments demonstrate the superiority of Molformer in various domains, which beats a group of baseline algorithms.

## References

- Anand-Achim, N.; Eguchi, R. R.; Mathews, I. I.; Perez, C. P.; Derry, A.; Altman, R. B.; and Huang, P. 2021. Protein sequence design with a learned potential. *bioRxiv*, 2020–01.
- Anderson, B.; Hy, T.-S.; and Kondor, R. 2019. Cormorant: Covariant molecular neural networks. *arXiv preprint arXiv:1906.04015*.
- Blum, L. C.; and Raymond, J.-L. 2009. 970 million drug-like small molecules for virtual screening in the chemical universe database GDB-13. *Journal of the American Chemical Society*, 131(25): 8732–8733.
- Cantoni, V.; Gatti, R.; and Lombardi, L. 2011. 3D Protein Surface Segmentation through Mathematical Morphology. In *International Joint Conference on Biomedical Engineering Systems and Technologies*, 97–109. Springer.
- Chen, B.; Barzilay, R.; and Jaakkola, T. 2019. Path-augmented graph transformer network. *arXiv preprint arXiv:1905.12712*.
- Chen, C.; Ye, W.; Zuo, Y.; Zheng, C.; and Ong, S. P. 2019. Graph networks as a universal machine learning framework for molecules and crystals. *Chemistry of Materials*, 31(9): 3564–3572.
- Chithrananda, S.; Grand, G.; and Ramsundar, B. 2020. Chemberta: Large-scale self-supervised pretraining for molecular property prediction. *arXiv preprint arXiv:2010.09885*.
- Cho, H.; and Choi, I. S. 2018. Three-dimensionally embedded graph convolutional network (3dgcn) for molecule interpretation. *arXiv preprint arXiv:1811.09794*.
- Duvenaud, D.; Maclaurin, D.; Aguilera-Iparraguirre, J.; Gómez-Bombarelli, R.; Hirzel, T.; Aspuru-Guzik, A.; and Adams, R. P. 2015. Convolutional networks on graphs for learning molecular fingerprints. *arXiv preprint arXiv:1509.09292*.
- Eismann, S.; Townshend, R. J.; Thomas, N.; Jagota, M.; Jing, B.; and Dror, R. O. 2021. Hierarchical, rotation-equivariant neural networks to select structural models of protein complexes. *Proteins: Structure, Function, and Bioinformatics*, 89(5): 493–501.
- Elnaggar, A.; Heinzinger, M.; Dallago, C.; Rihawi, G.; Wang, Y.; Jones, L.; Gibbs, T.; Feher, T.; Angerer, C.; Steinegger, M.; et al. 2020. ProtTrans: towards cracking the language of Life’s code through self-supervised deep learning and high performance computing. *arXiv preprint arXiv:2007.06225*.
- Fuchs, F. B.; Worrall, D. E.; Fischer, V.; and Welling, M. 2020. Se (3)-transformers: 3d roto-translation equivariant attention networks. *arXiv preprint arXiv:2006.10503*.
- Gayvert, K. M.; Madhukar, N. S.; and Elemento, O. 2016. A data-driven approach to predicting successes and failures of clinical trials. *Cell chemical biology*, 23(10): 1294–1301.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural message passing for quantum chemistry. In *International conference on machine learning*, 1263–1272. PMLR.
- Guo, Q.; Qiu, X.; Liu, P.; Xue, X.; and Zhang, Z. 2020. Multi-scale self-attention for text classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 7847–7854.
- Honda, S.; Shi, S.; and Ueda, H. R. 2019. Smiles transformer: Pre-trained molecular fingerprint for low data drug discovery. *arXiv preprint arXiv:1911.04738*.
- Huang, K.; Xiao, C.; Hoang, T.; Glass, L.; and Sun, J. 2020. Caster: Predicting drug interactions with chemical substructure representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 702–709.
- Hutchinson, M. J.; Le Lan, C.; Zaidi, S.; Dupont, E.; Teh, Y. W.; and Kim, H. 2021. Lietransformer: Equivariant self-attention for lie groups. In *International Conference on Machine Learning*, 4533–4543. PMLR.
- Ingraham, J.; Garg, V. K.; Barzilay, R.; and Jaakkola, T. 2019. Generative models for graph-based protein design. *Advances in neural information processing systems*.
- Ishida, S.; Miyazaki, T.; Sugaya, Y.; and Omachi, S. 2021. Graph Neural Networks with Multiple Feature Extraction Paths for Chemical Property Estimation. *Molecules*, 26(11): 3125.
- Jin, W.; Barzilay, R.; and Jaakkola, T. 2020. Hierarchical generation of molecular graphs using structural motifs. In *International Conference on Machine Learning*, 4839–4848. PMLR.
- Jumper, J.; Evans, R.; Pritzel, A.; Green, T.; Figurnov, M.; Ronneberger, O.; Tunyasuvunakool, K.; Bates, R.; Žídek, A.; Potapenko, A.; et al. 2021. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873): 583–589.
- Karimi, M.; Wu, D.; Wang, Z.; and Shen, Y. 2019. Deep-Affinity: interpretable deep learning of compound–protein affinity through unified recurrent and convolutional neural networks. *Bioinformatics*, 35(18): 3329–3338.
- Kearnes, S.; McCloskey, K.; Berndl, M.; Pande, V.; and Riley, P. 2016. Molecular graph convolutions: moving beyond fingerprints. *Journal of computer-aided molecular design*, 30(8): 595–608.
- Klicpera, J.; Giri, S.; Margraf, J. T.; and Günnemann, S. 2020. Fast and uncertainty-aware directional message passing for non-equilibrium molecules. *arXiv preprint arXiv:2011.14115*.
- Klicpera, J.; Groß, J.; and Günnemann, S. 2020. Directional message passing for molecular graphs. *arXiv preprint arXiv:2003.03123*.
- Li, A. J.; Sundar, V.; Grigoryan, G.; and Keating, A. E. 2022. TERMinator: a neural framework for structure-based protein design using tertiary repeating motifs. *arXiv preprint arXiv:2204.13048*.
- Li, X.; Yan, H.; Qiu, X.; and Huang, X. 2020. FLAT: Chinese NER using flat-lattice transformer. *arXiv preprint arXiv:2004.11795*.
- Liu, Y.; Wang, L.; Liu, M.; Zhang, X.; Oztekin, B.; and Ji, S. 2021. Spherical message passing for 3d graph networks. *arXiv preprint arXiv:2102.05013*.

- Lu, C.; Liu, Q.; Wang, C.; Huang, Z.; Lin, P.; and He, L. 2019. Molecular property prediction: A multilevel quantum interactions modeling perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 1052–1060.
- Mackenzie, C. O.; Zhou, J.; and Grigoryan, G. 2016. Tertiary alphabet for the observable protein structural universe. *Proceedings of the National Academy of Sciences*, 113(47): E7438–E7447.
- Martins, I. F.; Teixeira, A. L.; Pinheiro, L.; and Falcao, A. O. 2012. A Bayesian approach to in silico blood-brain barrier penetration modeling. *Journal of chemical information and modeling*, 52(6): 1686–1697.
- Maziarka, Ł.; Danel, T.; Mucha, S.; Rataj, K.; Tabor, J.; and Jastrzebski, S. 2020. Molecule attention transformer. *arXiv preprint arXiv:2002.08264*.
- Maziarka, Ł.; Majchrowski, D.; Danel, T.; Gaiński, P.; Tabor, J.; Podolak, I.; Morkisz, P.; and Jastrzebski, S. 2021. Relative Molecule Self-Attention Transformer. *arXiv preprint arXiv:2110.05841*.
- Öztürk, H.; Özgür, A.; and Ozkirimli, E. 2018. DeepDTA: deep drug–target binding affinity prediction. *Bioinformatics*, 34(17): i821–i829.
- Ramakrishnan, R.; Dral, P. O.; Rupp, M.; and von Lilienfeld, O. A. 2014. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific Data*, 1.
- Ramakrishnan, R.; Hartmann, M.; Tapavicza, E.; and Von Lilienfeld, O. A. 2015. Electronic spectra from TDDFT and machine learning in chemical space. *The Journal of chemical physics*, 143(8): 084111.
- Ramsundar, B.; Kearnes, S.; Riley, P.; Webster, D.; Konerding, D.; and Pande, V. 2015. Massively multitask networks for drug discovery. *arXiv preprint arXiv:1502.02072*.
- Rong, Y.; Bian, Y.; Xu, T.; Xie, W.; Wei, Y.; Huang, W.; and Huang, J. 2020. Self-supervised graph transformer on large-scale molecular data. *arXiv preprint arXiv:2007.02835*.
- Schütt, K. T.; Sauceda, H. E.; Kindermans, P.-J.; Tkatchenko, A.; and Müller, K.-R. 2018. SchNet—a deep learning architecture for molecules and materials. *The Journal of Chemical Physics*, 148(24): 241722.
- Schütt, K. T.; Unke, O. T.; and Gastegger, M. 2021. Equivariant message passing for the prediction of tensorial properties and molecular spectra. *arXiv preprint arXiv:2102.03150*.
- Somnath, V. R.; Bunne, C.; and Krause, A. 2021. Multi-Scale Representation Learning on Proteins. In *Thirty-Fifth Conference on Neural Information Processing Systems*.
- Subramanian, G.; Ramsundar, B.; Pande, V.; and Denny, R. A. 2016. Computational modeling of  $\beta$ -secretase 1 (BACE-1) inhibitors using ligand based approaches. *Journal of chemical information and modeling*, 56(10): 1936–1949.
- Thomas, N.; Smidt, T.; Kearnes, S.; Yang, L.; Li, L.; Kohlhoff, K.; and Riley, P. 2018. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219*.
- Townshend, R. J.; Vögele, M.; Suriana, P.; Derry, A.; Powers, A.; Laloudakis, Y.; Balachandar, S.; Anderson, B.; Eismann, S.; Kondor, R.; et al. 2020. ATOM3D: Tasks On Molecules in Three Dimensions. *arXiv preprint arXiv:2012.04035*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- Vinyals, O.; Bengio, S.; and Kudlur, M. 2015. Order matters: Sequence to sequence for sets. *arXiv preprint arXiv:1511.06391*.
- Wang, R.; Fang, X.; Lu, Y.; Yang, C.-Y.; and Wang, S. 2005. The PDBbind database: methodologies and updates. *Journal of medicinal chemistry*, 48(12): 4111–4119.
- Wang, X.; Girshick, R.; Gupta, A.; and He, K. 2018. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7794–7803.
- Wernicke, S. 2006. Efficient detection of network motifs. *IEEE/ACM transactions on computational biology and bioinformatics*, 3(4): 347–359.
- Wu, Z.; Liu, Z.; Lin, J.; Lin, Y.; and Han, S. 2020. Lite transformer with long-short range attention. *arXiv preprint arXiv:2004.11886*.
- Wu, Z.; Ramsundar, B.; Feinberg, E. N.; Gomes, J.; Geniesse, C.; Pappu, A. S.; Leswing, K.; and Pande, V. 2018. MoleculeNet: a benchmark for molecular machine learning. *Chemical science*, 9(2): 513–530.
- Xiong, Z.; Wang, D.; Liu, X.; Zhong, F.; Wan, X.; Li, X.; Li, Z.; Luo, X.; Chen, K.; Jiang, H.; et al. 2019. Pushing the boundaries of molecular representation for drug discovery with the graph attention mechanism. *Journal of medicinal chemistry*, 63(16): 8749–8760.
- Xu, Z.; Wang, S.; Zhu, F.; and Huang, J. 2017. Seq2seq fingerprint: An unsupervised deep molecular embedding for drug discovery. In *Proceedings of the 8th ACM international conference on bioinformatics, computational biology, and health informatics*, 285–294.
- Yang, K.; Swanson, K.; Jin, W.; Coley, C.; Eiden, P.; Gao, H.; Guzman-Perez, A.; Hopper, T.; Kelley, B.; Mathea, M.; et al. 2019. Analyzing learned molecular representations for property prediction. *Journal of chemical information and modeling*, 59(8): 3370–3388.
- Ying, C.; Cai, T.; Luo, S.; Zheng, S.; Ke, G.; He, D.; Shen, Y.; and Liu, T.-Y. 2021. Do Transformers Really Perform Bad for Graph Representation? *arXiv preprint arXiv:2106.05234*.
- Zhang, S.; Hu, Z.; Subramonian, A.; and Sun, Y. 2020. Motif-driven contrastive learning of graph representations. *arXiv preprint arXiv:2012.12533*.
- Zhang, Z.; Liu, Q.; Wang, H.; Lu, C.; and Lee, C.-K. 2021. Motif-based Graph Self-Supervised Learning for Molecular Property Prediction. *arXiv preprint arXiv:2110.00987*.