

PCEvo: Path-Consistent Molecular Representation via Virtual Evolutionary

Kun Li^{1,2}, Longtao Hu¹, Jiajun Yu³, Yida Xiong¹, Hongzhi Zhang¹,
Jiameng Chen¹, Xiantao Cai¹, Jia Wu⁴ and Wenbin Hu^{1,2,5,*}

¹School of Computer Science, Wuhan University

²Hubei Jiangxia Laboratory

³College of Computer Science and Technology, Zhejiang University

⁴Department of Computing, Macquarie University

⁵Shenzhen Research Institute, Wuhan University

{likun98, hlt_2003, yidaxiong, zhanghongzhi, jiameng.chen, caixiantao, hwb}@whu.edu.cn,
jiajunyu1999@gmail.com, jia.wu@mq.edu.au

Abstract

Molecular representation learning aims to learn vector embeddings that capture molecular structure and geometry, thereby enabling property prediction and downstream scientific applications. In many AI for science tasks, labeled data are expensive to obtain and therefore limited in availability. Under the few-shot setting, models trained with scarce supervision often learn brittle structure-property relationships, resulting in substantially higher prediction errors and reduced generalization to unseen molecules. To address this limitation, we propose **PCEvo**, a path-consistent representation method that learns from virtual paths through dynamic structural evolution. PCEvo enumerates multiple chemically feasible edit paths between retrieved similar molecular pairs under topological dependency constraints. It transforms the labels of the two molecules into stepwise supervision along each virtual evolutionary path. It introduces a path-consistency objective that enforces prediction invariance across alternative paths connecting the same two molecules. Comprehensive experiments on the QM9 and MoleculeNet datasets demonstrate that PCEvo substantially improves the few-shot generalization performance of baseline methods. The code is available at <https://github.com/DrugD/PCEvo>.

1 Introduction

Recent years have witnessed rapid progress in molecular representation learning driven by LLMs-based [Zheng *et al.*, 2025; Yu *et al.*, 2025b; Li *et al.*, 2025a], graph-based [Yu *et al.*, 2025a; Yu *et al.*, 2024], and geometry-aware neural architectures [Rao *et al.*, 2026; Lin *et al.*, 2026; Xiong *et al.*, 2025]. Message passing networks and 3D geometric models have become mainstream for learning structure-dependent signals [Qin *et al.*, 2026; Li *et al.*, 2025b], with representative backbones spanning MPNNs [Gilmer *et al.*, 2017], SchNet [Schütt *et al.*, 2018], DimeNet [Gasteiger *et*




Figure 1: Comparison of molecular representation paradigms: static end-to-end modeling, multi-stage learning, and virtual path modeling.

al., 2020], and more recent equivariant Transformer families such as Equiformer [Liao and Smidt, 2023] and ViSNet [Wang *et al.*, 2024b]. In parallel, self-supervised and contrastive pretraining has been increasingly adopted to leverage large unlabeled corpora and improve data efficiency [Li *et al.*, 2025c], including general graph contrastive objectives [Wang *et al.*, 2024a] and molecule-specific pretraining paradigms such as GROVER [Rong *et al.*, 2020], MolCLR [Wang *et al.*, 2022], and UniMol [Zhou *et al.*, 2023]. These developments are routinely evaluated on established benchmarks such as QM9 [Ramakrishnan *et al.*, 2014] and MoleculeNet [Wu *et al.*, 2018], and are further expanding toward cross-modal alignment between molecular structures and natural language descriptions [Wang *et al.*, 2025].

Despite this progress, the dominant learning paradigm remains largely static: each molecule is treated as a single end-point sample (i.e., only its final observed structure is used), and the model is trained to regress properties via direct supervision on sparse labels [An *et al.*, 2025; Zhang *et al.*, 2025] (as shown in Fig. 1). In few-shot settings, this formulation is prone to overfitting because chemical space is sparsely cov-




Figure 2: Virtual evolutionary editing uses the HOMO energy from the QM9 dataset to illustrate two evolution paths between a pair of molecules. Despite exhibiting different intermediate fluctuations, both paths inevitably converge to the same final HOMO value, highlighting the path-independent nature of the property.

ered; as a result, models can rely on dataset-specific correlations and frequent motifs, leading to unstable generalization. More fundamentally, most static molecular representations do not explicitly model how properties change under chemically feasible local edits to molecular structure [Li *et al.*, 2026]. In practical drug discovery, where labeled assays are costly and task distributions are sparse, capturing such regularities can improve the robustness of learned representations under the few-shot setting.

From a chemical perspective, molecules are discrete combinatorial objects whose structures can be naturally described as graphs and encoded by rule-based line notations [Gilmer *et al.*, 2017; Weininger, 1988]. More importantly, property changes in medicinal chemistry are often analyzed through small, chemically valid transformations, where a minimal structural modification induces a measurable shift in activity or physicochemical profiles [Ishikura *et al.*, 2025; Hussain and Rea, 2010a]. These observations motivate a virtual evolutionary path view of the structure–property relationship (Fig. 2) [Maggiora and Bajorath, 2014; Stumpfe and Bajorath, 2012]. By treating minimal edit operations, such as atom substitution, local connectivity adjustment, and functional group modification, as fundamental units [Walz *et al.*, 2025; Hussain and Rea, 2010b], we can construct virtual evolutionary paths connecting different molecular states [Jensen, 2019]. Modeling property variation along such paths enables the models to leverage incremental structural signals that are often underutilized when training solely on static endpoint samples.

Based on this perspective, we propose **PCEvo**, a path-consistent representation method via virtual evolutionary. PCEvo does not treat molecules as isolated, static samples; instead, it represents the structural differences between two similar molecules as an evolutionary sequence of chemically

feasible minimal edit operations. PCEvo further converts the labels of the two endpoint molecules into stepwise supervision distributed along the path, enabling the model to learn how incremental structural changes drive property variations during training, rather than performing the regression only at the endpoints. Importantly, there are typically multiple feasible editing paths connecting the same molecular pair; yet, the cumulative property change should remain consistent regardless of which path is taken. Leveraging this principle, PCEvo introduces a path-consistency objective that aligns the accumulated predictions across alternative feasible paths. This objective integrates seamlessly with standard static supervised training, thereby improving generalization in few-shot settings. We evaluate PCEvo on QM9 and MoleculeNet by integrating it into a range of representative backbone methods, thereby systematically demonstrating its generality and effectiveness; results show that PCEvo consistently reduces prediction error under few-shot settings on QM9 and achieves state-of-the-art (SOTA) performance on the three MoleculeNet regression tasks under the standard split, while maintaining reliable gains in both accuracy and stability when labeled data are limited. Our contributions are summarized as follows:

- We propose PCEvo, a path-consistent molecular representation method that learns from virtual evolutionary edit paths between similar molecules, transforming sparse end-point supervision into stepwise training signals along chemically feasible paths.
- Extensive experiments on QM9 and MoleculeNet under few-shot settings show that PCEvo consistently reduces prediction error and improves predictive stability, achieving SOTA performance compared with strong baselines.

2 Methods


2.1 Overview

To address the generalization limitations of static end-to-end learning in few-shot scenarios, as illustrated in Fig. 3, we propose a path-consistent molecular representation method, PCEvo. Deviating from conventional static structural mapping paradigms, PCEvo decomposes complex structure–property correlations into a sequence of learnable, cumulative edit steps. By imposing these path constraints, PCEvo could capture the intrinsic logic underlying molecular evolution.


2.2 Virtual Evolutionary Path Generation

To transform static molecular graphs into dynamic evolutionary paths, we propose a virtual evolutionary path generation pipeline. This pipeline converts a standard dataset $\mathcal{D} = \{(G_i, y_i)\}_{i=1}^N$ into an evolutionary dataset $\mathcal{D}_{\text{evo}} = \{(G_s, G_t, \tau) \mid \tau \in \mathcal{T}(G_s, G_t)\}$. For each molecular pair (G_s, G_t) , $\mathcal{T}(G_s, G_t)$ denotes the set of all valid evolutionary paths connecting the source molecule G_s to the target molecule G_t . Each path $\tau \in \mathcal{T}(G_s, G_t)$ is a sequence of edit operations. Explicitly modeling multiple valid paths for each molecular pair allows the PCEvo to account for the intrinsic diversity of structural transformations, facilitating the

1 - Virtual Evolutionary Path Generation



2 - Differential Encoding Unit



3 - Learning Objective




Figure 3: Overview of PCEvo method. Our method constructs virtual evolutionary paths between molecular states by identifying a sequence of valid graph edit operations. These discrete operations are then mapped into a continuous representation space to guide the learning of chemically interpretable molecular representations.

learning of property changes that are invariant to the ordering of operations.

Structural Neighbor Retrieval

Constructing evolutionary paths between arbitrary pairs of molecules is chemically implausible and computationally inefficient. To model realistic structural optimizations, we restrict the source molecules to the chemical neighborhood of the target.

For each target molecule G_t in the training or validation set, we query a candidate pool $\mathcal{G}_s$ to identify the top-K nearest neighbors. We utilize the Tanimoto similarity $f_{\text{Tanimoto}}(G_s, G_t)$ coefficient [Bajusz *et al.*, 2015] based on extended-connectivity fingerprints as the metric:

$$\mathcal{N}(G_t) = \text{Top}_K(\{G_s \in \mathcal{G}_{src} \mid G_s \neq G_t\}, f_{\text{Tanimoto}}). \quad (1)$$

This retrieval-based pairing strategy ensures that G_s and G_t share a significant structural scaffold, thereby defining a chemically meaningful optimization landscape.

Fundamental Path Operation Edit Units

Identifying the minimal set of edit operations ensures that the molecular transformation includes only essential structural changes, improving path generation efficiency and facilitating the learning of structure-property relationships. Given a pair (G_s, G_t) , we seek the minimal set of edit operations $\mathcal{S}$ required to transform G_s into G_t . We formulate this as a maximum common subgraph (MCS) problem [Fuchs and Riesen, 2025; Raymond and Willett, 2002; Bunke, 1997].

We establish an atom-level alignment $\pi : V_s \rightarrow V_t$ by minimizing the matching cost matrix defined by atomic num-

bers and valence constraints. Based on the atom-level alignment π , the symmetric difference between the two graphs is decomposed into a set of atomic operations. The definitions of these operations are detailed in Table 1. Specifically, unmapped atoms in the source molecule G_s correspond to REMOVE_ATOM operations, while unmapped atoms in the target molecule G_t correspond to ADD_ATOM operations. At the bond level, discrepancies between mapped atom pairs (u, v) in G_s and $(\pi(u), \pi(v))$ in G_t give rise to ADD_BOND, REMOVE_BOND, or CHANGE_BOND operations.

Applying this procedure yields an unordered set of operations $\mathcal{S} = \{o_1, \dots, o_M\}$, which constitutes the minimal set of edits required to transform G_s into G_t .

Topologically Constrained Path Sampling

While the set $\mathcal{S}$ specifies the structural changes, it does not determine the order in which they should be applied. Performing the operations in an arbitrary sequence may violate topological constraints, such as adding a bond to an atom that does not yet exist. To resolve this, we construct a directed acyclic graph, termed the dependency graph: $\mathcal{G}_{dep} = (\mathcal{S}, \mathcal{E}_{dep})$. An edge $o_i \rightarrow o_j$ exists if operation o_i is a topological prerequisite for o_j . The core constraints include:

- Existence precondition: An ADD_BOND(u, v) operation depends on the ADD_ATOM operations that create nodes u and v .
- Deletion latency: A REMOVE_ATOM(u) operation must succeed all bond removals or modifications involving node u .

Any linear extension (topological sort) of $\mathcal{G}_{dep}$ constitutes a valid path τ . Importantly, a single edit set $\mathcal{S}$ often allows

Algorithm 1 Virtual Evolutionary Path Generation

Require: Standard dataset $\mathcal{D}$, Number of neighbors K , Max paths P_{max}
Ensure: Evolutionary dataset $\mathcal{D}_{evo}$

- 1: Initialize $\mathcal{D}_{evo} \leftarrow \emptyset$
- 2: **for all** $G_t \in \mathcal{D}$ **do**
- 3: Candidate pool $\mathcal{G}_s \leftarrow (G_t, \mathcal{G}_s, K)$
- 4: **for all** $G_s \in \mathcal{G}_s$ **do**
- 5: $\pi \leftarrow \text{ALIGNMENT}(G_s, G_t)$
- 6: $\mathcal{S} \leftarrow \text{EXTRACTMINIMALEDITS}(G_s, G_t, \pi)$
- 7: $\mathcal{G}_{dep} \leftarrow \text{BUILDDependencyGRAPH}(\mathcal{S})$
- 8: $\mathcal{T}_{pair} \leftarrow \emptyset$
- 9: **while** $|\mathcal{T}_{pair}| < P_{max}$ **do**
- 10: $\tau \leftarrow \text{SAMPLING}(\mathcal{G}_{dep})$
- 11: $\mathcal{T}_{pair} \leftarrow \mathcal{T}_{pair} \cup \{\tau\}$
- 12: **end while**
- 13: **for all** $\tau \in \mathcal{T}_{pair}$ **do**
- 14: $\mathcal{D}_{evo} \leftarrow \mathcal{D}_{evo} \cup \{(G_s, G_t, \tau)\}$
- 15: **end for**
- 16: **end for**
- 17: **end for**
- 18: **return** $\mathcal{D}_{evo}$

exponentially many valid permutations. We leverage this by randomly sampling up to P_{max} distinct topological sorts for each molecular pair. This serves as a powerful combinatorial data augmentation strategy: the model learns to associate the same net property change Δy with multiple distinct paths $\{\tau_1, \dots, \tau_{P_{max}}\}$, thereby enforcing path-independence of the physical property and improving generalization. Each path $\tau \in \mathcal{T}(G_s, G_t)$ is a sequence of edit operations, and its length may vary across different paths.

2.3 Differential Evolutionary Path Encoder

To represent the feature changes induced by each virtual edit operation along a molecular structural evolution path, a differential evolutionary path encoder is proposed. Unlike the standard paradigm that encodes molecules as isolated static data points, this module explicitly models the incremental changes between consecutive intermediate states along the virtual path τ .

Let $G_s^{(t)}$ denote the intermediate molecular state after the t -th edit operation in the path, with the boundary conditions $G_s^{(0)} = G_s$ and $G_s^{(T)} = G_t$, where T is the total path length.

To bridge the discrete symbolic space of chemical edits and the continuous latent space, we first define a bijective mapping $\psi: \mathcal{O} \rightarrow \mathbb{R}^D$. The t -th edit operation o_t , which triggers the transition $G_s^{(t)} \rightarrow G_s^{(t+1)}$, is vectorized via concatenated encodings:

$$\mathbf{x}_t = \psi(o_t) = [\mathbf{e}_{type} \parallel \mathbf{e}_{atom} \parallel \mathbf{e}_{bond} \parallel \mathbf{e}_{pos}], \quad (2)$$

where components encode the operation type, atomic semantics, bond attributes, and positional indices, respectively. This vectorization provides the initial condition for the specific structural transformation at step t .

Then, we employ a shared-weight molecular encoder $f_\theta(\cdot)$ to project the adjacent states $G_s^{(t)}$ and $G_s^{(t+1)}$ into a continu-

Operation	Description
ADD/REM/REP_ATOM	Adds, removes, or replaces an atom node.
ADD/REM/CHG_BOND	Adds, removes, or modifies a bond edge.

Table 1: Elementary Graph Edit Operations.

ous latent manifold:

$$\mathbf{h}_t = f_\theta(G_s^{(t)}), \quad \mathbf{h}_{t+1} = f_\theta(G_s^{(t+1)}). \quad (3)$$

Our method PCEvo is architecture-agnostic and shared across all steps to ensure consistent feature extraction along the path. For each virtual edit, we model its incremental effect by differencing adjacent representations and predicting the corresponding property change:

$$\mathbf{d}_t = \mathbf{h}_{t+1} - \mathbf{h}_t, \quad \Delta \hat{P}_t = \phi_\psi(\mathbf{d}_t), \quad (4)$$

which learns edit-wise contributions in an additive manner.

2.4 Learning Objective

Our optimization jointly anchors absolute property prediction in the global chemical space and enforces summation consistency of property variations along virtual evolutionary paths. The model consists of a shared-weight backbone encoder $f_\theta(\cdot)$ and a non-linear delta predictor $\phi_\psi(\cdot)$. Given a labeled molecular pair (G_s, G_t) with properties (y_s, y_t) and a feasible edit path $\tau = (G_s^{(0)}, \dots, G_s^{(T)})$ satisfying $G_s^{(0)} = G_s$ and $G_s^{(T)} = G_t$, we denote the latent representation at step t as

$$\mathbf{h}_t = f_\theta(G_s^{(t)}), \quad t = 0, 1, \dots, T. \quad (5)$$

Static Property Loss. We supervise absolute predictions at the observed endpoints:

$$\mathcal{L}_{\text{static}} = \|f_\theta(G) - y\|_1 \quad (6)$$

Path Summation Consistency Loss. We require that the cumulative sum of predicted stepwise property increments along the path matches the ground-truth endpoint difference:

$$\mathcal{L}_{\text{cons}} = \left\| \widehat{\Delta P}(\tau) - (y_t - y_s) \right\|_1, \quad (7)$$

$$\widehat{\Delta P}(\tau) = \sum_{t=0}^{T-1} \Delta \hat{P}_t = \sum_{t=0}^{T-1} \phi_\psi(\mathbf{h}_{t+1} - \mathbf{h}_t). \quad (8)$$

where the non-linear predictor $\phi_\psi(\cdot)$ ensures the summation does not mathematically collapse into a simple difference of endpoints. This forces the gradient signal to propagate through every intermediate state, compelling the model to learn the specific physical contribution of each elementary edit operation. The final objective is $\mathcal{L} = \mathcal{L}_{\text{static}} + \mathcal{L}_{\text{cons}}$.

3 Theoretical Analysis

This section explains why PCEvo generalizes better by decomposing sparse static endpoint regression supervision into edit-level signals along evolutionary paths. Two mechanisms are emphasized: (i) edit decomposition increases the effective number of supervised signals; (ii) learning on edit-local representation increments reduces hypothesis complexity.

Methods		HOMO (100-shot)			HOMO (1000-shot)			LUMO (100-shot)			LUMO (1000-shot)		
		MAE ↓	MSE ↓	PCC ↑	MAE ↓	MSE ↓	PCC ↑	MAE ↓	MSE ↓	PCC ↑	MAE ↓	MSE ↓	PCC ↑
GCN [Kipf and Welling, 2017]		0.5031	0.4298	-0.2760	0.4318	0.3393	0.3345	0.7651	0.8399	0.6577	0.6330	0.6617	0.7859
GraphSAGE [Hamilton <i>et al.</i> , 2017]		0.5121	0.4523	-0.1318	0.4759	0.4077	-0.0035	0.7618	0.8946	0.6397	0.7134	0.7961	0.7322
GAT [Veličković <i>et al.</i> , 2018]		0.5016	0.4290	-0.2948	0.7613	0.8418	0.1641	0.7682	0.8763	0.6408	0.7234	0.8132	0.7238
GIN [Xu <i>et al.</i> , 2019]		0.4928	0.4190	-0.1601	0.3529	0.2176	0.6347	0.6768	0.7474	0.7225	0.6001	0.6568	0.8257
EMPP [An <i>et al.</i> , 2025]		0.4058	0.2872	0.5820	0.2733	0.1244	0.8115	0.7328	0.7804	0.6983	0.2713	0.1223	0.8150
EMPP* [An <i>et al.</i> , 2025]		0.7520	0.8391	0.0587	2.0454	4.6518	-0.2126	0.9696	1.4393	0.6016	1.0356	1.5991	0.6211
Moleculeformer [Qin <i>et al.</i> , 2025]		0.3353	0.2165	0.6966	0.2710	0.1274	0.8068	0.8985	1.2785	0.4202	0.5642	0.5420	0.8192
GotenNet [Aykent and Xia, 2025]		0.5565	0.6819	-0.0591	0.3743	0.2860	0.5950	0.8360	1.1393	0.5379	0.4194	0.2938	0.9067
SchNet [Schütt <i>et al.</i> , 2018]	+PCEvo	0.3330	0.1853	0.7694	0.2057	0.0758	0.8954	0.5275	0.5253	0.8494	0.2269	0.0924	0.9712
	Origin.	0.4688	0.4917	0.3413	0.3617	0.2580	0.6336	0.6159	0.6006	0.8050	0.2944	0.1551	0.9514
	Improve.	28.97%	62.30%	125.41%	43.13%	70.61%	41.31%	14.35%	12.54%	5.52%	22.94%	40.45%	2.08%
DimeNet [Gasteiger <i>et al.</i> , 2020]	+PCEvo	0.5502	0.4888	0.3362	0.3711	0.2373	0.6824	0.6955	0.8305	0.7117	0.5553	0.5210	0.8242
	Origin.	0.6162	0.6871	0.1406	0.4392	0.3322	0.4835	0.7173	0.7877	0.6814	0.5642	0.5420	0.8192
	Improve.	10.71%	28.86%	139.04%	15.50%	28.55%	41.13%	3.04%	-5.43%	4.45%	1.58%	3.88%	0.61%
Equiformer [Liao and Smidt, 2023]	+PCEvo	0.3971	0.3152	0.5319	0.2292	0.0963	0.8637	0.5463	0.5210	0.8055	0.2592	0.1292	0.9593
	Origin.	0.5189	0.5251	0.0726	0.3098	0.1698	0.7366	0.8261	1.0350	0.5609	0.4264	0.2920	0.9057
	Improve.	23.46%	39.98%	631.98%	26.02%	43.29%	17.25%	33.87%	49.66%	43.60%	39.19%	55.72%	5.91%
ViSNet [Wang <i>et al.</i> , 2024b]	+PCEvo	0.4570	0.4725	0.4609	0.2346	0.1001	0.8696	0.4073	0.2905	0.9035	0.2490	0.1178	0.9635
	Origin.	0.5889	0.7685	-0.1158	0.3181	0.2159	0.7107	0.4852	0.3693	0.8791	0.2585	0.1297	0.9603
	Improve.	22.39%	38.52%	497.95%	26.25%	53.62%	22.36%	16.04%	21.33%	2.78%	3.70%	9.19%	0.32%

* denotes EMPP trained with an auxiliary objective that additionally predicts masked atomic coordinates.

Table 2: Overall experiments on QM9 dataset for HOMO and LUMO (units: eV) properties under few-shot settings. **Origin** denotes the baseline backbone, and **Improve.(%)** indicates the performance gain.

Methods	ESOL	FreeSolv	Lipophilicity	Avg.
GCN [Kipf and Welling, 2017]	1.431	2.870	0.712	1.671
GIN [Xu et al., 2019]	1.452	2.765	0.850	1.689
N-GRAM [Liu et al., 2019]	1.100	2.510	0.880	1.497
Hu et al. [Hu et al., 2020]	1.100	2.764	0.739	1.534
GROVER [Rong et al., 2020]	1.423	2.947	0.823	1.731
MolCLR [Wang et al., 2022]	1.113	2.301	0.789	1.401
Equiformer [Liao and Smidt, 2023]	1.037	2.672	0.989	1.566
SchNet [Schütt et al., 2018]	0.912	2.255	0.923	1.363
DimeNet [Gasteiger et al., 2020]	1.135	2.589	1.043	1.589
ViSNet [Wang et al., 2024b]	1.049	3.158	0.944	1.717
PCEvo SchNet (Ours)	1.060	1.958	0.762	1.260
PCEvo DimeNet (Ours)	1.117	1.843	0.925	1.295
PCEvo ViSNet (Ours)	0.969	2.539	0.798	1.436

Table 3: Overall experiments on MoleculeNet under standard data splits. RMSE is reported for each task, and **Avg.** denotes the mean RMSE over all three tasks.

Methods	ESOL (1128)			FreeSolv (642)			Lipophilicity (4200)		
	R=0.1	R=0.2	R=0.8	R=0.1	R=0.2	R=0.8	R=0.1	R=0.2	R=0.8
SchNet									
+PCEvo	1.484	1.568	1.060	3.490	2.383	1.958	1.112	0.854	0.762
Origin.	1.098	1.046	0.912	3.972	4.042	2.255	1.112	1.067	0.923
Improve.	-35.15	-49.94	-16.21	12.15	41.06	13.18	-0.02	19.96	17.45
DimeNet									
+PCEvo	1.489	1.663	1.117	3.623	2.555	1.843	1.111	0.980	0.925
Origin.	2.156	2.249	1.135	5.287	4.910	2.589	1.083	1.065	1.043
Improve.	30.92	26.06	1.54	31.47	47.96	28.82	-2.54	7.99	11.29
ViSNet									
+PCEvo	1.235	1.165	0.969	2.857	2.684	2.539	1.080	0.837	0.798
Origin.	1.341	1.333	1.049	4.255	3.801	3.158	1.104	1.052	0.944
Improve.	7.93	12.56	7.70	32.87	29.40	19.59	2.18	20.43	15.51

Table 4: Overall experiments on MoleculeNet under different train/test splits. RMSE is reported for each task, where R denotes the training sample ratio.

3.1 Learning Formulation

Given a labeled dataset $\mathcal{D} = \{(G_i, y_i)\}_{i=1}^N$ and the induced evolutionary dataset $\mathcal{D}_{\text{evo}}$, we compare static endpoint learn-

ing and path-consistent learning. Here $\hat{\mathcal{R}}(\cdot)$ denotes the empirical risk (sample average of the loss) minimized during training.

Static model. The static objective $\mathcal{L}_{\text{static}}$ supervises absolute property prediction on $\mathcal{D}$. We denote by $\mathcal{H}_{\text{static}}$ the hypothesis class of all static predictors realizable by the backbone when parameters vary:

$$\hat{\mathcal{R}}_{\text{static}}(f_\theta) = \frac{1}{N} \sum_{i=1}^N \ell(f_\theta(G_i), y_i), \quad (9)$$

$$\mathcal{H}_{\text{static}} = \{f_\theta(\cdot) \mid \theta \in \Theta\}. \quad (10)$$

Path-consistent learning. The consistency objective $\mathcal{L}_{\text{cons}}$ enforces that the summed stepwise deltas along $\tau = (G_s^{(0)}, \dots, G_s^{(T)})$ match the endpoint difference $(y_t - y_s)$. With $h_t = f_\theta(G_s^{(t)})$, the predicted total change is

$$\widehat{\Delta P}(\tau) = \sum_{t=0}^{T-1} \phi_\psi(h_{t+1} - h_t). \quad (11)$$

We denote by $\mathcal{H}_{\text{cons}}$ the hypothesis class of all delta predictors realizable by the MLP head when its parameters vary:

$$\hat{\mathcal{R}}_{\text{cons}}(\phi_\psi) = \frac{1}{|\mathcal{D}_{\text{evo}}|} \sum_{\tau \in \mathcal{D}_{\text{evo}}} \ell(\widehat{\Delta P}(\tau), y_t - y_s), \quad (12)$$

$$\mathcal{H}_{\text{cons}} = \{\phi_\psi(\cdot) \mid \psi \in \Psi\}. \quad (13)$$

3.2 Generalization Bound

We adopt Rademacher complexity to compare the generalization of static endpoint learning and edit-path learning.

Assumption 1 (Effective sample size under mixing). *Edit-level samples induced by virtual paths form a weakly dependent sequence. For instance, they satisfy β -mixing with*

$\sum_{k \geq 1} \beta(k) < \infty$. Then the edit-level learning admits an effective sample size

$$n_{\text{eff}} \geq \frac{N\bar{L}}{c_{\text{mix}}}, \quad c_{\text{mix}} = 1 + 2 \sum_{k \geq 1} \beta(k), \quad (14)$$

where $\bar{L}$ is the average path length.

Assumption 1 states that, although edit-level samples generated along paths are weakly dependent, their statistical effect can be captured by an effective sample size n_{eff} , so the generalization bound applies by replacing n with n_{eff} .

Let $\mathfrak{R}_n(\mathcal{H})$ denote the empirical Rademacher complexity. By the standard Rademacher generalization theory [Yin *et al.*, 2019; Bartlett *et al.*, 2005], the generalization gap scales as $O(\mathfrak{R}_n(\mathcal{H}) + \sqrt{\log(1/\delta)/n})$ uniformly over $f \in \mathcal{H}$. Under Assumption 1, edit-level learning obeys the same scaling with n replaced by n_{eff} (up to constants absorbed into c_{mix}).

Theorem 1 (Edit decomposition yields a tighter generalization bound). Assume $\ell(\cdot, \cdot)$ is bounded in $[0, 1]$ and is 1-Lipschitz in its first argument. Assume $\|h\|_2 \leq B_{\text{static}}$ and $\|\Delta h\|_2 \leq B_{\text{edit}}$. If hypotheses in $\mathcal{H}_{\text{static}}$ and $\mathcal{H}_{\text{cons}}$ are L -Lipschitz with respect to their inputs, then:

(i) Complexity scaling. The empirical Rademacher complexity scales with the input radius and the sample size. In particular,

$$\mathfrak{R}_N(\mathcal{H}_{\text{static}}) = O\left(\frac{LB_{\text{static}}}{\sqrt{N}}\right), \quad (15)$$

$$\mathfrak{R}_{n_{\text{eff}}}(\mathcal{H}_{\text{cons}}) = O\left(\frac{LB_{\text{edit}}}{\sqrt{n_{\text{eff}}}}\right). \quad (16)$$

(ii) Generalization bounds. With probability at least $1 - \delta$,

$$\mathcal{R}_{\text{static}} \leq \hat{\mathcal{R}}_{\text{static}} + 2\mathfrak{R}_N(\mathcal{H}_{\text{static}}) + 3\sqrt{\frac{\log(2/\delta)}{2N}}, \quad (17)$$

$$\mathcal{R}_{\text{cons}} \leq \hat{\mathcal{R}}_{\text{cons}} + 2\mathfrak{R}_{n_{\text{eff}}}(\mathcal{H}_{\text{cons}}) + 3\sqrt{\frac{\log(2/\delta)}{2n_{\text{eff}}}}. \quad (18)$$

The two bounds share the same form, but differ in the effective sample size and the input radius. Edit-path learning replaces (N, B_{static}) with $(n_{\text{eff}}, B_{\text{edit}})$, where n_{eff} is induced by path decomposition and Δh is edit-local. In typical settings, a single edit produces a local perturbation in representation space, so $B_{\text{edit}} \leq B_{\text{static}}$. Therefore, the edit-path bound becomes tighter whenever

$$\frac{B_{\text{edit}}}{\sqrt{n_{\text{eff}}}} < \frac{B_{\text{static}}}{\sqrt{N}}. \quad (19)$$

In particular, the static generalization gap scales as

$$\mathcal{R}_{\text{static}} - \hat{\mathcal{R}}_{\text{static}} = O\left(\frac{LB_{\text{static}}}{\sqrt{N}} + \sqrt{\frac{1}{N}}\right), \quad (20)$$

while edit-path learning follows the same scaling with (N, B_{static}) replaced by $(n_{\text{eff}}, B_{\text{edit}})$. Finally, enforcing path consistency further restricts the admissible solutions within $\mathcal{H}_{\text{cons}}$, which can only reduce the effective complexity and tighten the above bound.

4 Experiments

Datasets. We conduct our evaluations on the QM9 dataset [Ramakrishnan *et al.*, 2014], focusing on the prediction of the highest occupied molecular orbital (ϵ_{HOMO}) and lowest unoccupied molecular orbital (ϵ_{LUMO}). For physicochemical property prediction, we further employ the ESOL, FreeSolv, and Lipophilicity datasets from MoleculeNet [Wu *et al.*, 2018].

Baselines. We evaluate PCEvo against a comprehensive set of backbone methods, grouped into three categories: **2D GNNs/Transformer:** GCN [Kipf and Welling, 2017], GraphSAGE [Hamilton *et al.*, 2017], GAT [Veličković *et al.*, 2018], GIN [Xu *et al.*, 2019], and Moleculeformer [Qin *et al.*, 2025]; **3D geometry-aware:** SchNet [Schütt *et al.*, 2018], DimeNet [Gasteiger *et al.*, 2020], Equiformer [Liao and Smidt, 2023], ViSNet [Wang *et al.*, 2024b], and GotenNet [Aykent and Xia, 2025]; **Pretraining:** N-GRAM [Liu *et al.*, 2019], GROVER [Rong *et al.*, 2020], MolCLR [Wang *et al.*, 2022], Hu *et al.* [Hu *et al.*, 2020], and EMPP [An *et al.*, 2025].

Implementation Details. We designed specific data splitting protocols to simulate varying degrees of data scarcity. For the QM9 dataset, we constructed extremely few-shot training/testing sets containing randomly sampled 100 and 1,000 molecules. For the smaller MoleculeNet datasets, we simulated few-shot learning by using 10%, 20%, and 80% of the total data for training; the 10% data were allocated for testing. Performance was assessed using three standard regression metrics: mean absolute error (MAE), mean squared error (MSE), and the Pearson correlation coefficient (PCC).

4.1 Overall Experiments

To evaluate the effectiveness and scalability of PCEvo for molecular representation, we conduct experiments on QM9 and MoleculeNet. Following our path construction procedure, for each target molecule, we retrieve the top- $N=5$ structural neighbors and enumerate up to $P_{\text{max}}=50$ valid edit paths for each molecular pair under topological dependency constraints.

On QM9, we adopt two low-resource settings with $N=100$ and $N=1000$ labeled training samples (results in Table 2). PCEvo yields consistent improvements across four 3D backbones (SchNet, DimeNet, Equiformer, and ViSNet), and the correlation recovery is particularly pronounced in few-shot settings ($N=100$): for ϵ_{HOMO} , Equiformer improves PCC from 0.07 to 0.53 (+0.46). Moreover, compared with recent strong baselines such as EMPP, Moleculeformer, and GotenNet, PCEvo achieves better performance under the same few-shot protocol. These results indicate that when the chemical space is sparsely covered, PCEvo can effectively enhance the generalization of backbone models.

On MoleculeNet, since the benchmarks are intrinsically small-scale, we evaluate PCEvo under both the standard split (80% training; Table 3) and explicit low-resource splits with training ratios $R \in 0.1, 0.2, 0.8$ (Table 4). The results show that integrating PCEvo into SchNet, DimeNet, and ViSNet leads to consistent gains across tasks and training ratios. The improvement on FreeSolv is especially evident; for example, with SchNet at $R=0.2$, RMSE decreases from 4.042 to

$P_{\max}$	N	$\mathcal{L}_{\text{cons}}$	MAE↓	MSE↓	PCC↑
5	1	✓	0.487	0.430	0.145
5	1	✗	0.672	0.881	-0.088
5	3	✓	0.385	0.265	0.659
5	3	✗	0.518	0.480	-0.108
5	5	✓	0.319	0.179	0.761
5	5	✗	0.390	0.273	0.639
1	1	✓	0.687	0.942	-0.085
*1	1	✗	0.469	0.492	0.341
1	3	✓	0.672	0.894	-0.088
1	3	✗	0.677	0.904	-0.086
1	5	✓	0.494	0.443	0.127
1	5	✗	0.486	0.428	0.153

Table 5: Ablation study. * denotes the vanilla SchNet results.

2.383. Meanwhile, under the standard split, PCEvo also outperforms several widely used pretraining baselines: for instance, PCEvo_{SchNet} achieves an average RMSE of 1.260, improving over MolCLR (1.401) by 0.141.


Overall, PCEvo exhibits stable and transferable improvements across backbones and few-shot settings, and compares favorably to a range of pretraining methods. These results suggest that moving beyond static endpoint regression by decomposing supervision into edit-wise steps and enforcing path-consistency can yield more robust molecular representations in few-shot settings.

4.2 Ablation Study

To isolate the effect of the path-consistency constraint $\mathcal{L}_{\text{cons}}$, we conduct a controlled ablation study. We use SchNet as the backbone, train under the 100-shot, and evaluate on a held-out validation set of 100 disjoint molecules, with ϵ_{HOMO} as the target property. Across all runs, we keep the static endpoint supervision $\mathcal{L}_{\text{static}}$ enabled. Therefore, within each $(N, P_{\max})$ configuration, the only difference is whether $\mathcal{L}_{\text{cons}}$ is enabled.

Main finding: $\mathcal{L}_{\text{cons}}$ is effective when multiple paths are available. Table 5 shows that relative to the single path setting ($P_{\max} = 1$), multi-path training ($P_{\max} = 5$) consistently improves performance when the path-consistency constraint $\mathcal{L}_{\text{cons}}$ is enabled. This trend holds across all neighbor budgets $N \in \{1, 3, 5\}$, indicating that the benefit of $\mathcal{L}_{\text{cons}}$ is robust to the choice of N . In contrast, simply increasing the number of paths without $\mathcal{L}_{\text{cons}}$ does not lead to reliable gains and can even degrade performance in few-shot settings.

On the role of naive duplication. In our training construction, endpoint molecules are reused across different molecular pairs and paths. Such reuse mainly increases the number of training instances, but does not necessarily introduce additional chemical diversity at the endpoints; therefore, the observed gains cannot be explained by sample count alone. Table 5 provides a controlled comparison supporting this point: once the path-consistency constraint is removed, increasing the path budget ($P_{\max} > 1$) or the neighbor budget ($N > 1$) no longer yields consistent improvements; instead, performance becomes unstable under some configurations.


 Figure 4: Hyperparameter sensitivity analysis on QM9 with SchNet under the 100-shot setting. (a) fixes $N=10$ and varies $P_{\max}$, (b) fixes $P_{\max}=50$ and varies N .

4.3 Hyperparameter Sensitivity Analysis

To analyze the impact of the maximum number of virtual paths $P_{\max}$ and the number of source neighbors N on PCEvo, we conduct experiments on the QM9 dataset under the 100-shot setting using SchNet as the backbone, with results shown in Fig. 4. Regarding the number of virtual paths, we observe that a single path ($P_{\max} = 1$) yields negligible gains over the baseline, indicating that simple augmentation is insufficient. However, increasing $P_{\max}$ significantly boosts performance by enforcing consistency across diverse topological orderings. This confirms that path-consistent via virtual evolutionary is central to the PCEvo’s success. Regarding the context size N , a clear trade-off exists between diversity and relevance. Performance peaks at $N = 5$; utilizing larger neighborhoods ($N \geq 15$) degrades accuracy, as distant neighbors introduce chemically irrelevant scaffolds that generate noisy and overly complex edit paths.

5 Conclusion

In this work, we propose PCEvo, a path-consistent molecular representation method that augments conventional static regression supervision with a virtual evolutionary edit-path consistency constraint. This design enables the model to learn how local structural edits accumulate into global property variations, thereby yielding more robust structure–property mappings. Experiments on QM9 and MoleculeNet under few-shot settings show that PCEvo consistently reduces error, improves stability, and achieves SOTA performance against strong baselines.

Acknowledgments

This work was supported in part by the National Key Research and Development Program of China (2023YFC2705705), the Natural Science Foundation of China (No. 62476203), Key Project of Traditional Chinese Medicine Joint Fund of Hubei Provincial Natural Science Foundation (No.2025AFD47), Hubei Province Science and Technology Innovation Plan Project (No.2025BCB035), the Guangdong Provincial Natural Science Foundation General Project (No. 2025A1515012155), the Shenzhen Natural Science Foundation Project (No. JCYJ20250604122534006, JCYJ20230807090211021). The Hubei Jiangxia laboratory development project (No. 2025JXKF034), the New

Cornerstone Science Foundation through the XPLOER PRIZE, The Innovative Research Group Project of Hubei Province under Grants 2024AFA017.

References

- [An *et al.*, 2025] Junyi An, Chao Qu, Yun-Fei Shi, XinHao Liu, Qianwei Tang, Fenglei Cao, and Yuan Qi. Equivariant masked position prediction for efficient molecular representation. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [Aykent and Xia, 2025] Sarp Aykent and Tian Xia. GotenNet: Rethinking Efficient 3D Equivariant Graph Neural Networks. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [Bajusz *et al.*, 2015] Dávid Bajusz, Anita Rácz, and Károly Héberger. Why is tanimoto index an appropriate choice for fingerprint-based similarity calculations? *Journal of cheminformatics*, 7(1):20, 2015.
- [Bartlett *et al.*, 2005] Peter L Bartlett, Olivier Bousquet, and Shahar Mendelson. Local rademacher complexities. *The Annals of Statistics*, 33(4):1497–1537, 2005.
- [Bunke, 1997] Horst Bunke. On a relation between graph edit distance and maximum common subgraph. *Pattern recognition letters*, 18(8):689–694, 1997.
- [Fuchs and Riesen, 2025] Mathias Fuchs and Kaspar Riesen. Fast approximate maximum common subgraph computation. *Pattern Recognition Letters*, 190:66–72, 2025.
- [Gasteiger *et al.*, 2020] Johannes Gasteiger, Janek Groß, and Stephan Günnemann. Directional message passing for molecular graphs. In *International Conference on Learning Representations*, 2020.
- [Gilmer *et al.*, 2017] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International conference on machine learning*, pages 1263–1272. Pmlr, 2017.
- [Hamilton *et al.*, 2017] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.
- [Hu *et al.*, 2020] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. Strategies for pre-training graph neural networks. In *International Conference on Learning Representations*, 2020.
- [Hussain and Rea, 2010a] Jameed Hussain and Ceara Rea. Computationally efficient algorithm to identify matched molecular pairs (mmps) in large data sets. *Journal of Chemical Information and Modeling*, 50(3):339–348, 2010.
- [Hussain and Rea, 2010b] Jameed Hussain and Ceara Rea. Computationally efficient algorithm to identify matched molecular pairs (mmps) in large data sets. *Journal of chemical information and modeling*, 50(3):339–348, 2010.
- [Ishikura *et al.*, 2025] Hikaru Ishikura, Callum S Begg, Juan J Rojas, Luka Blagojevic, Gavin J Smith, Joyce Luk, Rosemary A Croft, Charles Romain, Chulho Choi, and James A Bull. Do amino-oxetanes resemble amides? a matched molecular pairs property and structural comparison. 2025.
- [Jensen, 2019] Jan H Jensen. A graph-based genetic algorithm and generative model/monte carlo tree search for the exploration of chemical space. *Chemical science*, 10(12):3567–3572, 2019.
- [Kipf and Welling, 2017] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017.
- [Li *et al.*, 2025a] Kun Li, Zhennan Wu, Shoupeng Wang, Jia Wu, Shirui Pan, and Wenbin Hu. Druggipilot: Llm-based parameterized reasoning agent for drug discovery. *arXiv preprint arXiv:2505.13940*, 2025.
- [Li *et al.*, 2025b] Kun Li, Yida Xiong, Hongzhi Zhang, Xiantao Cai, Jia Wu, Bo Du, and Wenbin Hu. Graph-structured small molecule drug discovery through deep learning: Progress, challenges, and opportunities. In *2025 IEEE International Conference on Web Services (ICWS)*, pages 1033–1042, 2025.
- [Li *et al.*, 2025c] Kun Li, Yue Zeng, Yi-da Xiong, Hao-chen Wu, Sui Fang, Zhi-yan Qu, Yan Zhu, Bo Du, Zhao-bing Gao, and Wen-bin Hu. Contrastive learning-based drug screening model for glun1/glun3a inhibitors. *Acta Pharmacologica Sinica*, pages 1–13, 2025.
- [Li *et al.*, 2026] Kun Li, Longtao Hu, Jiameng Chen, Hongzhi Zhang, Yida Xiong, Xiantao Cai, Wenbin Hu, and Jia Wu. Can molecular evolution mechanism enhance molecular representation? *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(18):15108–15116, Mar. 2026.
- [Liao and Smidt, 2023] Yi-Lun Liao and Tess Smidt. Equiformer: Equivariant graph attention transformer for 3d atomistic graphs. In *The Eleventh International Conference on Learning Representations*, 2023.
- [Lin *et al.*, 2026] Zhuohao Lin, Kun Li, Jiameng Chen, Jiajun Yu, Duanhua Cao, Yizhen Zheng, and Wenbin Hu. Rethinking molecular ood generalization via target-aware source selection. *arXiv preprint arXiv:2605.13932*, 2026.
- [Liu *et al.*, 2019] Shengchao Liu, Mehmet F Demirel, and Yingyu Liang. N-gram graph: Simple unsupervised representation for graphs, with applications to molecules. *Advances in neural information processing systems*, 32, 2019.
- [Maggiora and Bajorath, 2014] Gerald M Maggiora and Jürgen Bajorath. Chemical space networks: a powerful new paradigm for the description of chemical space. *Journal of computer-aided molecular design*, 28(8):795–802, 2014.
- [Qin *et al.*, 2025] Mingyuan Qin, Ziyang Sun, Lei Feng, Chongyin Han, Jingjing Xia, and Lianyi Han. Molecule-former is a gcn-transformer architecture for molecular

- property prediction. *Communications Biology*, 8(1):1668, 2025.
- [Qin *et al.*, 2026] Xiaohan Qin, Chao Wang, Zhengyang Zhou, Linjiang Chen, Wenjie Du, and Yang Wang. Msanchor: De novo molecular generation from mass spectrometry data with anchor-extended molecular scaffolds. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 953–961, 2026.
- [Ramakrishnan *et al.*, 2014] Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole Von Lilienfeld. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1(1):1–7, 2014.
- [Rao *et al.*, 2026] Haojie Rao, Kun Li, Yida Xiong, Jiameng Chen, Wenbin Hu, Yizhen Zheng, Jiajun Yu, and Duanhua Cao. From single-step edit response to multi-step molecular optimization. *arXiv preprint arXiv:2605.10035*, 2026.
- [Raymond and Willett, 2002] John W Raymond and Peter Willett. Maximum common subgraph isomorphism algorithms for the matching of chemical structures. *Journal of computer-aided molecular design*, 16(7):521–533, 2002.
- [Rong *et al.*, 2020] Yu Rong, Yatao Bian, Tingyang Xu, Weiyang Xie, Ying Wei, Wenbing Huang, and Junzhou Huang. Self-supervised graph transformer on large-scale molecular data. *Advances in neural information processing systems*, 33:12559–12571, 2020.
- [Schütt *et al.*, 2018] Kristof T Schütt, Huziel E Sauceda, P-J Kindermans, Alexandre Tkatchenko, and K-R Müller. SchNet—a deep learning architecture for molecules and materials. *The Journal of Chemical Physics*, 148(24), 2018.
- [Stumpfe and Bajorath, 2012] Dagmar Stumpfe and Jurgen Bajorath. Exploring activity cliffs in medicinal chemistry: miniperspective. *Journal of medicinal chemistry*, 55(7):2932–2942, 2012.
- [Veličković *et al.*, 2018] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018.
- [Walz *et al.*, 2025] Carlo Walz, Moritz Spiske, Magnus Walter, Benjamin-Luca Keller, Mario Mezler, Carolin Hoft, Frauke Pohlki, Stella Vukelic, and Felix Hausch. Macrocyclization as a strategy for kinetic solubility improvement: A comparative analysis of matched molecular pairs. *Journal of Medicinal Chemistry*, 68(3):2639–2656, 2025.
- [Wang *et al.*, 2022] Yuyang Wang, Jianren Wang, Zhonglin Cao, and Amir Barati Farimani. Molecular contrastive learning of representations via graph neural networks. *Nature Machine Intelligence*, 4(3):279–287, 2022.
- [Wang *et al.*, 2024a] Mengzhu Wang, Jiao Li, Houcheng Su, Nan Yin, Liang Yang, and Shen Li. Graphcl: Graph-based clustering for semi-supervised medical image segmentation. *arXiv preprint arXiv:2411.13147*, 2024.
- [Wang *et al.*, 2024b] Yusong Wang, Tong Wang, Shaoning Li, Xinheng He, Mingyu Li, Zun Wang, Nanning Zheng, Bin Shao, and Tie-Yan Liu. Enhancing geometric representations for molecules with equivariant vector-scalar interactive message passing. *Nature Communications*, 15(1):313, Jan 2024.
- [Wang *et al.*, 2025] Runze Wang, Mingqi Yang, and Yanming Shen. Bridging molecular graphs and large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 21234–21242, 2025.
- [Weininger, 1988] David Weininger. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of Chemical Information and Computer Sciences*, 28(1):31–36, 1988.
- [Wu *et al.*, 2018] 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. *Chemical science*, 9(2):513–530, 2018.
- [Xiong *et al.*, 2025] Yida Xiong, Kun Li, Jiameng Chen, Hongzhi Zhang, Di Lin, Yan Che, and Wenbin Hu. Text-guided multi-property molecular optimization with a diffusion language model. *Information Fusion*, page 103907, 2025.
- [Xu *et al.*, 2019] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019.
- [Yin *et al.*, 2019] Dong Yin, Ramchandran Kannan, and Peter Bartlett. Rademacher complexity for adversarially robust generalization. In *International conference on machine learning*, pages 7085–7094. PMLR, 2019.
- [Yu *et al.*, 2024] Jiajun Yu, Zhihao Wu, Jinyu Cai, Adele Lu Jia, and Jicong Fan. Kernel readout for graph neural networks. In *IJCAI*, pages 2505–2514, 2024.
- [Yu *et al.*, 2025a] Jiajun Yu, Zhihao Wu, Jielong Lu, Tianyue Wang, and Haishuai Wang. A centrality-based graph learning framework. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, pages 3588–3596, 2025.
- [Yu *et al.*, 2025b] Jiajun Yu, Yizhen Zheng, Huan Yee Koh, Shirui Pan, Tianyue Wang, and Haishuai Wang. Collaborative expert llms guided multi-objective molecular optimization. *arXiv preprint arXiv:2503.03503*, 2025.
- [Zhang *et al.*, 2025] Hongzhi Zhang, Zhonglie Liu, Kun Meng, Jiameng Chen, Jia Wu, Bo Du, Di Lin, Yan Che, and Wenbin Hu. Zero-shot learning with subsequence reordering pretraining for compound-protein interaction, 2025.
- [Zheng *et al.*, 2025] Yizhen Zheng, Huan Yee Koh, Jiaxin Ju, Madeleine Yang, Lauren T May, Geoffrey I Webb, Li Li, Shirui Pan, and George Church. Large language models for drug discovery and development. *Patterns*, 6(10), 2025.
- [Zhou *et al.*, 2023] Gengmo Zhou, Zhifeng Gao, Qiankun Ding, Hang Zheng, Hongteng Xu, Zhewei Wei, Linfeng Zhang, and Guolin Ke. Uni-mol: A universal 3d molecular representation learning framework. In *The Eleventh International Conference on Learning Representations*, 2023.