

MiniMol: A Parameter-Efficient Foundation Model for Molecular Learning

Kerstin Kläser^{*1} Błażej Banaszewski^{*1} Samuel Maddrell-Mander² Callum McLean¹
Luis Müller³ Ali Parviz^{4,5} Shenyang Huang⁵ Andrew Fitzgibbon¹

Abstract

In biological tasks, data is rarely plentiful as it is generated from hard-to-gather measurements. Therefore, pre-training foundation models on large quantities of available data and then transfer to low-data downstream tasks is a promising direction. However, how to design effective foundation models for molecular learning remains an open question, with existing approaches typically focusing on models with large parameter capacities. In this work, we propose MiniMol, a foundational model for molecular learning with 10 million parameters. MiniMol is pre-trained on a mix of roughly 3300 sparsely defined graph- and node-level tasks of both quantum and biological nature. The pre-training dataset includes approximately 6 million molecules and 500 million labels. To demonstrate the generalizability of MiniMol across tasks, we evaluate it on downstream tasks from the Therapeutic Data Commons (TDC) AD-MET group showing significant improvements over the prior state-of-the-art foundation model across 17 tasks. MiniMol will be a public and open-sourced model for future research.

1. Introduction

Accurate prediction of molecular properties plays an essential role in many applications, including novel drug discovery (Stokes et al., 2020; Jin et al., 2021; Wallach et al., 2015), efficient catalyst development (Zitnick et al., 2020), and materials design (Reiser et al., 2022). Traditionally, Density Functional Theory (DFT) methods (Nakata & Shimazaki, 2017) accurately compute molecular properties by physics simulation, but are computationally demanding even for small molecules, and becomes intractable in large scale of biological systems (Stevens et al., 2023). Consequently, deep learning methods such as Graph Neural

^{*}Equal contribution ¹Graphcore ²Work carried out while at Graphcore ³RWTH Aachen University ⁴New Jersey Institute of Technology ⁵Mila - Quebec AI Institute. Correspondence to: Kerstin Kläser <kerstink@graphcore.ai>, Błażej Banaszewski <blazejb@graphcore.ai>.

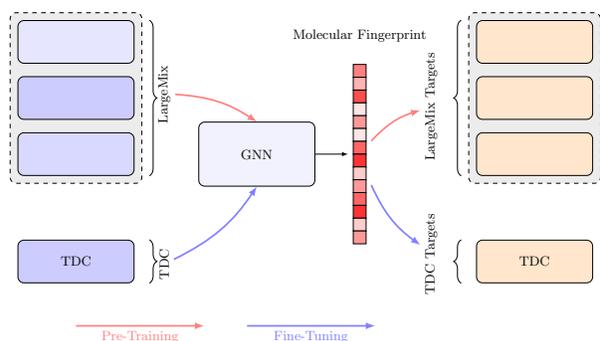


Figure 1. Workflow overview of the MiniMol pre-training and downstream task evaluation. MiniMol is pre-trained on the LargeMix datasets. Then MiniMol embeddings are used as molecular fingerprints for downstream tasks such as those in TDC.

Networks (GNNs) (Masters et al., 2023a; Gasteiger et al., 2019) and graph transformers (Rampásek et al., 2022) have achieved significant success in molecular representation learning. This is demonstrated in the recent Open Graph Benchmark (OGB) Large Scale Challenge where on the PCQM4Mv2 dataset, deep learning models produce accurate approximations of DFT while being significantly faster (Lu et al., 2023; Masters et al., 2023a). In addition, by training on DFT calculations and biological tasks, ML models can predict complex biochemical properties not possible with DFT alone. Therefore, building foundation models that capture chemical and biological knowledge from large amounts of pre-training data and adapt it for a wide range of downstream low-data tasks is a promising next step.

Prior work on foundational models for molecular learning has typically adopted the common practices used in computer vision and natural language processing, aiming for large model capacity and pre-training dataset size. Molecular foundation models such as MoLE (Méndez-Lucio et al., 2022), ChemBERTa-2 (Ahmad et al., 2022) and Galactica (Taylor et al., 2022) directly receive SMILES strings (Weininger et al., 1989) of molecules as input. However, a number of equally valid SMILES strings can denote the same molecule thus SMILES string based models are

unable to represent the symmetries underlying the molecular graphs. Therefore, large amounts of pre-training data and a large model capacity is required to properly learn these symmetries. Recently, ULTRA (Galkin et al., 2023), a foundational model for knowledge graphs, demonstrated that properly respecting the symmetries of the underlying data alleviates the need for large models in order to beat task-specific baselines. Specifically, ULTRA improves the state-of-the-art on a wide variety of knowledge graph reasoning tasks with only 177K parameters. Therefore, we leverage the permutation invariant property of GNNs to build a generally-capable and parameter-efficient foundation model for molecular fingerprinting.

In this work, we propose MiniMol, a parameter-efficient foundation model for molecular learning based on GNN backbone. MiniMol is pre-trained on the Graphium LargeMix dataset (Beaini et al., 2024) with around 6 million molecules and 526 million data labels. The pre-training strategy of Graphium is multi-level and multi-task meaning that over 3300 sparsely defined tasks on both graph and node level are trained jointly. MiniMol demonstrates strong downstream task performance on the Therapeutic Data Commons (TDC) ADMET group (Huang et al., 2021). Figure 1 shows a simplified diagram of this idea. Our main contributions are as follows:

- In this work, we propose MiniMol, a parameter-efficient foundation model with a GINE backbone of 10 million parameters, pre-trained on around 3300 biological and quantum tasks on the graph and node level of molecules.
- We demonstrate that the molecular fingerprint from MiniMol is highly transferable to downstream tasks. On the TDC benchmark, the current state of the art for a single model applied to all tasks (a foundation model) is MolE, which achieves a mean rank of 5.4, when compared against the specialized per-task models on the leaderboard. MiniMol achieves a mean rank of 3.6 and outperforms MolE on 17 tasks.
- We conduct a thorough performance correlation analysis between the pre-training datasets and downstream tasks. We found that PCQM4M_G25 dataset often has a negative correlation with downstream tasks thus highlighting the importance of understanding the correlation between pre-training tasks and downstream tasks.
- MiniMol will be a public and open-sourced model for future research. With only 10% of the parameters of prior state-of-the-art, MiniMol offers strong downstream performance and lower compute requirement to adapt.

Reproducibility: We include the code and weight checkpoint for MiniMol needed to reproduce the experiments as supplementary materials.

2. Related Work

2.1. Molecular Fingerprints

Traditional molecular fingerprints such as Extended Connectivity Fingerprint (ECFP) (Rogers & Hahn, 2010), RDkit fingerprints (Landrum et al., 2013) and MAP4 (Capecchi et al., 2020) are designed for molecular characterization, similarity searching, and structure-activity modelling, with wide applications in drug discovery. However, they encode the presence of particular substructures within the molecule and should be manually customized for specific applications. In addition, it is shown that different types of fingerprints perform better for specific categories of molecules. For example, substructure fingerprint (Kim, 2021) has the best performance on small molecules such as drugs while atom-pair fingerprints (Awale & Reymond, 2014) are best suited for large molecules such as peptides. Even SOTA fingerprints can suffer from embedding collisions due to how sub-structures are resolved (Probst et al., 2022). From large molecular datasets, foundation models aim to learn a universal and descriptive molecular representation as practical molecular fingerprints for downstream tasks. Our MiniMol model generates strong fingerprints for various tasks on the TDC benchmark as indicated by downstream task performance.

2.2. Foundational Models in ML.

In Natural Language Processing (NLP) and Computer Vision (CV), foundation models have achieved significant progress, especially for Large Language Models (LLMs) (Achiam et al., 2023). Foundation models are often pre-trained on a mixture of data across a wide range of tasks. Downstream tasks with low amounts of data can then be solved using low-resource, inexpensive fine-tuning (Tian et al., 2023; Borzunov et al., 2023). Multi-modal datasets have been used with LLMs to align a single model representation across domains (Team et al., 2023; Betker et al., 2023). In most areas, these general properties have typically been emergent from extremely large models.

2.3. Foundational Models in Molecular Learning.

Initial unsupervised transformer-based designs most closely replicate work in NLP (Honda et al., 2019; Wang et al., 2019; Honda et al., 2019), representing molecules as SMILES strings (Weininger, 1988). These models leverage extremely large but low-fidelity unlabelled molecule datasets. Early FMs can achieve strong results on some small, out-of-domain tasks. However, generalizability to a wide variety of tasks remains limited (Zhu et al., 2021; Liu et al., 2023; Méndez-Lucio et al., 2022; Luo et al., 2023). Recent state-of-the-art in molecular property prediction employs geometric deep learning, often with an MPNN or Graph Transformer, trained with supervised labels of molec-

ular properties (Ying et al., 2021; Veličković et al., 2017; Dwivedi & Bresson, 2020). Recent work by (Masters et al., 2023b) has shown promising potential for MPNN scaling in-depth and total parameters for predicting the quantum properties of molecules. Recent work COATI (Kaufman et al., 2023) is based on SMILES and point clouds, utilizing a multi-modal encoder-decoder scheme designed for molecule regression tasks. Significantly, Shoghi et al. (2023) demonstrates the importance of multi-task pre-training on a range of length scales for low-resource fine-tuning and extracting learned representations for molecular property prediction (MPP). In this work, we use graph-based representations of molecules, which provide rich chemical and structural information. MiniMol is pre-trained on a large number of molecular properties of both quantum and biological tasks while demonstrating strong performance across many downstream tasks.

3. Method

Here, we present our architecture for pre-training on the LargeMix datasets (Beaini et al., 2023), extracting fingerprints and subsequently fine-tuning to downstream tasks (see Fig. 2).

3.1. Molecular Representation

Each molecule is modelled as a graph \mathcal{G} with N nodes representing the atoms and M edges representing the bonds. We denote the set of edges with \mathcal{E} . The atom and bond features are generated using RDKit, providing a set of categorical and floating values, and the atomic features are concatenated with positional and structural embeddings. From (Masters et al., 2023b; Rampásek et al., 2022) the Laplacian eigenvectors and eigenvalues, and the random walk probabilities, were found to be most beneficial. The input node feature vectors are the concatenation of these features

$$X^0 = [X^{\text{atom}} | X^{\text{LapVec}} | X^{\text{LapVal}} | X^{\text{RW}}],$$

and edge features are the bond features $E^0 = [E^{\text{bond}}]$. A global node is added to each graph, providing an additional connection to every node. It was shown in (Li et al., 2017) that the global node dramatically improves graph-level representation. This acts both as routing between otherwise distant portions of the graph and as a readout node for the graph property. It is initialized with a random vector. Each of the nodes, edges, and global features are initially embedded into the model dimensions using a two-layer MLP each (eq 1 - 3).

$$x^0 = \text{MLP}_x(X^0) \in \mathbb{R}^{N \times d_{\text{node}}} \quad (1)$$

$$e^0 = \text{MLP}_e(E^0) \in \mathbb{R}^{M \times d_{\text{edge}}} \quad (2)$$

$$g^0 = \text{MLP}_g(\text{rand}_g(0)) \in \mathbb{R}^{d_{\text{global}}} \quad (3)$$

3.2. Model

Given the initial node, edge and graph embeddings we update them through multiple layers of message-passing to obtain final node embeddings

$$x^{\text{final}} = \text{GNN}(x^0, e^0, g^0),$$

where GNN is a chosen GNN backbone. As described later in Section 4, we try three different backbone GNNs, namely GCN (Kipf & Welling, 2017), GINE (Hu et al., 2020b; Xu et al., 2019) and MPNN++ (Masters et al., 2023a). We briefly describe the different architectures here.

GCN The GCN (Kipf & Welling, 2016) incorporates only the node embeddings. Concretely, the ℓ -th layer of a GCN is defined as

$$x_i^{\ell+1} = \sum_{j \in \mathcal{N}(i) \cup \{i\}} \frac{1}{\sqrt{d_i d_j}} x_j^\ell,$$

where $\mathcal{N}(i)$ denotes the set of neighbors of i and d_i and d_j denote the degree of nodes i and j , respectively.

GINE GINE is an extension of GIN (Xu et al., 2019) to additionally incorporate edge embeddings (Hu et al., 2020b). Concretely, the ℓ -th layer of GINE is defined as

$$x_i^{\ell+1} = \text{MLP} \left((1 - \epsilon) \cdot x_i^\ell \sum_{j \in \mathcal{N}(i)} \text{ReLU}(x_j^\ell + e_{ij}^\ell) \right),$$

where ϵ is a learnable scalar. Compared to GCN, GINE is more expressive as it has the same expressive power for graph isomorphism test as the 1-Weisfeiler-Leman algorithm (Xu et al., 2019). Further, GINE has been used previously for pre-training graph neural networks (Hu et al., 2020b).

MPNN++ The Message Passing Neural Network (MPNN) module in (Masters et al., 2023a) called MPNN++, is designed to incorporate and update node-, edge- and graph-level embeddings. The MPNN++ is a complex architecture as it aggregates messages from adjacent nodes, incident edges and global graph-level embeddings. In addition, the MPNN++ uses layer-wise skip connections for node-, edge- and graph-level embeddings; see Appendix A.1 for a detailed description of the MPNN++ layer. We include MPNN++ as a backbone in our experiments.

3.3. Pre-training

MiniMol is jointly pre-trained with many supervised tasks on both the graph and node levels. The total loss minimized during training is a summation of each of the pre-training tasks, accounting for label sparsity per molecule. There is a

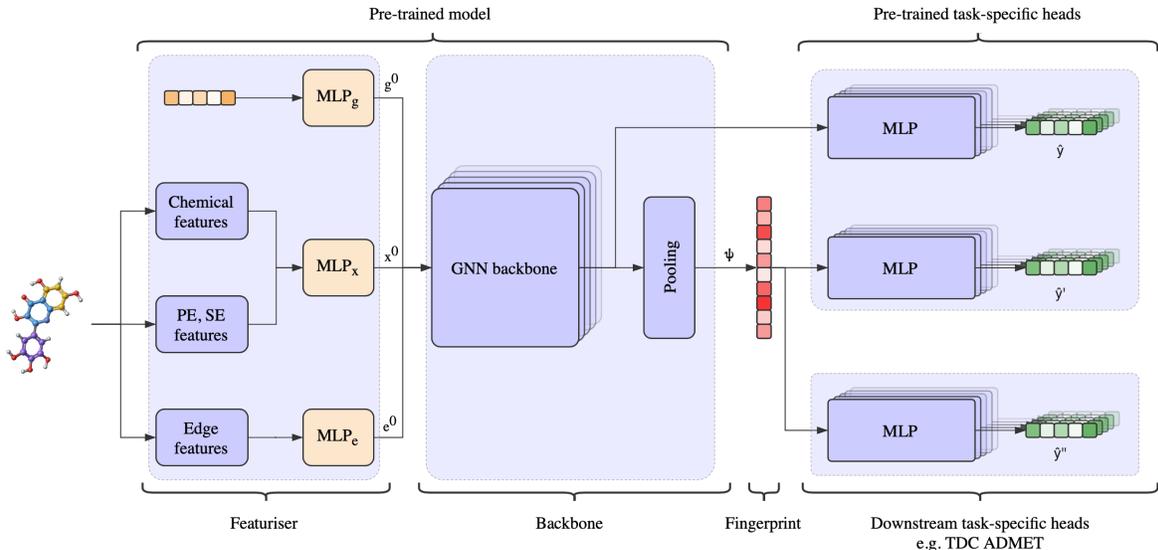


Figure 2. Schematic of the architecture of MiniMol. An example molecule is featurized in the first block. Node feature vectors are created by combining chemical features with positional and structural encodings, edge features are generated using RDKit and a random initial global vector is generated. Each initial vector is processed with a separate embedding MLP. The backbone of the model is a stack of MPNN layers, which output the molecular fingerprint ψ after pooling. The pre-pooling output is used for pre-training on node-level tasks, in our case, PCQM4M_N4. The fingerprint ψ is used either for pre-training multiple graph-level task heads or as an input to the downstream tasks, including the full set of ADMET tasks from the TDC benchmarks.

rich literature on how to combine losses for multi-task learning (Crawshaw, 2020), although in this case, the tasks are a combination of regression and classification, meaning the losses are not naturally commensurate. The mean absolute error (MAE) loss is used for the PCQM dataset (N4 and G25 tasks), the binary cross-entropy (BCE) loss is used for the PCBA tasks, and the hybrid cross-entropy (HCE) loss from (Beaini et al., 2023) is used for the L1000 datasets. Concretely, we compute the final loss \mathcal{L} as

$$\mathcal{L} = \mathcal{L}_{L1000} + \mathcal{L}_{PCBA} + \mathcal{L}_{N4} + \frac{1}{k} \mathcal{L}_{G25}.$$

where k is a scaling constant and we set $k = 5$ to account for data imbalance on the G25 dataset (see more discussion in Section 4).

3.4. Downstream Fingerprinting

For downstream tasks, we generate the global embeddings of the final layer of MiniMol from a given molecule which is also referred to as molecular fingerprints. This is more compute efficient and easy to use when compared to fine-tuning the entire model from end-to-end. In addition, generating meaningful molecular fingerprints for downstream tasks is an important aspect of the foundation model of molecular learning.

More specifically, we first extract fingerprints for all unique molecules in a given downstream task and subsequently use

the fingerprints as molecular representation to train a small Multilayer Perceptron (MLP) for making task-specific predictions. To generate the fingerprints from MiniMol, we compute the final node-, edge- and graph-level embeddings as described in Section 3.2 and subsequently obtain fingerprints by pooling the final node embeddings, e.g., via max pooling, obtaining

$$\psi = \sum_{i \in \mathcal{G}} x_i^{\text{final}}$$

for the fingerprint vector ψ .

There are several advantages of using molecular fingerprints for downstream tasks when compared to end-to-end fine-tuning. First, the above procedure allows a significantly more efficient way to train a model on low-data downstream tasks. Molecular fingerprints are pre-computed for a given set of molecules with a single forward pass from the model and then used for many downstream tasks or hyperparameter sweeps. Second, the downstream model is less likely to overfit when only the downstream MLP is trainable as there are fewer trainable parameters. Next, the use of molecular fingerprints for downstream tasks in this way matches existing workflows in the bio-chemistry domain increasing the practical utility of the method. Finally, as fingerprints are simply embedding vectors, practitioners are not required to access the architecture or weight of the pre-trained model thus only the expertise to train a task head (or MLP) is required.

4. Experimental Details

In our experiments, we pre-train MiniMol on LargeMix (Beaini et al., 2023) for various GNN backbones and subsequently fine-tune to all 22 tasks in the ADMET Group of the TDC benchmark.

4.1. Pre-training

MiniMol uses the LargeMix datasets from (Beaini et al., 2023) for pre-training, consisting of approximately 6M molecules and a total of 526M targets, which have been summarized in Table 1. The datasets are described below.

PCQM4M_G25_N4. This dataset contains 3.8M molecules from the PCQM dataset (Hu et al., 2021), from the OGB-LSC challenge. The dataset consists of quantum chemistry calculations for 25 molecular graph-level properties, and 4 node-level properties per atom, resulting in about 400M labelled data points.

PCBA. This dataset contains 1.5M molecules from the OGBG-PCBA dataset (Hu et al., 2020a). This bioassay dataset, derived experimentally from high-throughput screening methods, details the impact of the molecules on living cells across 1328 sparse labels. This results in about 100M labelled data points.

L1000_VCAP and L1000_MCF7. These datasets contain 26k molecules from the L1000 dataset (Subramanian et al., 2017) which details the change to gene expression profiles and cellular processes when exposed to the molecules in the dataset across about 1000 labels and 26M data points.

Table 1. Overview of the datasets in LargeMix.

Dataset	# Molecules	# Labels	# Data Points	% of All Data Points
PCQM4M_G25	3.81M	25 (G)	93M	17%
PCQM4M_N4	3.81M	4 (N)	197.7M	37%
PCBA_132B	1.56M	1328 (G)	224.4M	41%
L1000_VCAP	15K	978 (G)	15M	3%
L1000_MCF7	12K	978 (G)	11M	2%

These diverse labels from fundamental quantum chemistry properties to macro-scale cellular impact encourage a single general representation of the molecule suitable for downstream tasks. The combined LargeMix contains multiple task labels per molecule. The datasets only partially overlap thus requiring the model to generalize across domains from sparse labels on molecules. Following (Méndez-Lucio et al., 2022), we filter out molecules with more than 100 heavy atoms. In addition, we remove molecules in the ADMET group test sets from our pre-training data to avoid potential leakage of test labels (7% of MCF7, 4% of VCAP, 0.6% of PCBA, 0.07% of PCQM4M_G25/N4). During pre-training we split the dataset into 92% training, 4% validation, and 4% test data.

To cover a range of GNN backbones with increasing com-

plexity, we pre-train GCN, GINE, and finally MPNN++ models and subsequently evaluate their downstream performance on the TDC ADMET group datasets. Each model consists of 16 GNN layers with hidden dimensions adjusted such that all models have around 10M parameters. We train each model for 100 epochs using the Adam optimizer, with a maximum learning rate of $3e^{-4}$, 5 warm-up epochs and linear learning rate decay.

4.2. Benchmarking on TDC ADMET Group

Here, we describe the ADMET group of the TDC benchmark which we use to evaluate the downstream performance of MiniMol. The Therapeutics Data Commons (TDC) (Huang et al., 2021) is a platform designed to facilitate the assessment and development of AI methods in drug discovery. It particularly emphasizes identifying the most effective AI techniques for this purpose. Within TDC, the ADMET Benchmark Group specializes in single-instance prediction, offering a standardized suite of 22 datasets for molecular property prediction. These datasets vary in size, ranging from 475 to 13,130 molecules, and encompass tasks in both regression and classification. The datasets span a breadth of molecular properties, categorized into Absorption, Distribution, Metabolism, Excretion, and Toxicity. To ensure fair comparability, scaffold splits are employed, with 80% of data for training and 20% for testing.

A diverse array of models, including random forests, GNNs, and CNNs, are evaluated on these tasks. Their performance is showcased on the TDC leaderboard which includes various models trained with SMILES or other encoding strategies. To benchmark MiniMol, we first build an ensemble by training a distinct model on each fold in 5-fold cross-validation. While building the ensemble, the best epoch is selected based on validation loss, and to distinguish which ensemble to select for testing (e.g. while choosing one out of the sweep), the ensemble’s mean validation metric is used. Final test scores are derived from the top ensemble, with error bars reported from five trials (see Appendix A.4 for pseudo code). Table 3 presents the performance of three GNN architectures (GCN, GINE, MPNN++) across various datasets and two computational budgets for sweeping hyperparameters of the dataset-specific models. Since our fingerprinting approach permits fast evaluation of downstream predictors, we conduct extensive hyperparameter sweeping across all tasks (see Appendix A.2 for more details about the hyperparameter selection). For CPU-only runs, training a downstream model only takes 1 to 10 minutes per model per dataset.

5. Empirical Results

Here, we present our experimental results for both pre-training and fine-tuning on fingerprints.

Table 2. Results on downstream evaluation of MiniMol (GINE) with max pooling on TDC ADMET benchmarks, and comparison to the TDC leaderboard and MoIE. The rank is determined for each dataset individually, on a set of 7 scores, which include the test results from the TOP5 leaderboard, MoIE and MiniMol. The best result is shown in green and the top three results are highlighted in purple.

		TDC Dataset		Leaderboard Jan. 2024	MoIE		MiniMol (GINE)		
		Name	Size	Metric	SOTA Result	Result	Rank	Result	Rank
ABSORPTION	Caco2 Wang	906	MAE (↓)	0.276 ± .005	0.310 ± .010	6	0.324 ± .012	7	
	Bioavailability Ma	640	AUROC (↑)	0.748 ± .033	0.654 ± .028	7	0.699 ± .008	6	
	Lipophilicity AZ	4,200	MAE (↓)	0.467 ± .006	0.469 ± .009	3	0.455 ± .001	1	
	Solubility AqSolDB	9,982	MAE (↓)	0.761 ± .025	0.792 ± .005	5	0.750 ± .012	1	
	HIA Hou	578	AUROC (↑)	0.989 ± .001	0.963 ± .019	7	0.994 ± .003	1	
	Pgp Broccatelli	1,212	AUROC (↑)	0.938 ± .002	0.915 ± .005	7	0.994 ± .002	1	
DISTRIB.	BBB Martins	1,975	AUROC (↑)	0.920 ± .006	0.903 ± .005	7	0.923 ± .002	1	
	PPBR AZ	1,797	MAE (↓)	7.526 ± .106	8.073 ± .335	6	7.807 ± .188	4	
	VDss Lombardo	1,130	Spearman (↑)	0.713 ± .007	0.654 ± .031	3	0.570 ± .015	7	
METABOLISM	CYP2C9 Veith	12,092	AUPRC (↑)	0.859 ± .001	0.801 ± .003	5	0.819 ± .001	4	
	CYP2D6 Veith	13,130	AUPRC (↑)	0.790 ± .001	0.682 ± .008	6	0.718 ± .003	5	
	CYP3A4 Veith	12,328	AUPRC (↑)	0.916 ± .000	0.867 ± .003	7	0.878 ± .001	5	
	CYP2C9 Substrate	666	AUPRC (↑)	0.441 ± .033	0.446 ± .062	2	0.481 ± .013	1	
	CYP2D6 Substrate	664	AUPRC (↑)	0.736 ± .024	0.699 ± .018	7	0.726 ± .006	2	
	CYP3A4 Substrate	667	AUROC (↑)	0.667 ± .019	0.670 ± .018	1	0.644 ± .006	6	
EXCRET.	Half Life Obach	667	Spearman (↑)	0.576 ± .025	0.549 ± .024	4	0.493 ± .002	7	
	Clearance Hepatocyte	1,102	Spearman (↑)	0.536 ± .020	0.381 ± .038	7	0.448 ± .006	4	
	Clearance Microsome	1,020	Spearman (↑)	0.630 ± .010	0.607 ± .027	6	0.652 ± .007	1	
TOXICITY	LD50 Zhu	7,385	MAE (↓)	0.552 ± .009	0.823 ± .019	7	0.588 ± .010	3	
	hERG	648	AUROC (↑)	0.880 ± .002	0.813 ± .009	7	0.849 ± .007	6	
	Ames	7,255	AUROC (↑)	0.871 ± .002	0.883 ± .005	1	0.856 ± .001	5	
	DILI	475	AUROC (↑)	0.925 ± .005	0.577 ± .021	7	0.944 ± .007	1	
TDC Leaderboard Mean Rank:					5.4		3.6		

5.1. Pre-training on LargeMix

We present our pre-training results for MiniMol with GCN, GINE and MPNN++ as backbone GNNs on LargeMix in Table 3. Here, we observe that pre-training performance is only marginally affected by the choice of the backbone GNN. Moreover, the pre-training performance of a given GNN backbone also varies with tasks. For example, on the graph- and node-level regression tasks of PCQM4M_G25 and PCQM4M_N4 the MPNN++ backbone performs consistently better while GINE performs best on the classification tasks of PCBA_1328. On L1000_VCAP we observe roughly comparable results across all metrics with a slight advantage of GINE in terms of AUROC. Finally, on L1000_MCF7, GCN and GINE are largely on par with the MPNN++ showing worse performance in terms of AUROC. Most importantly, no single backbone is superior to the other two. Next, we present our results of fine-tuning these models to the

ADMET group benchmark of TDC using our fingerprinting approach.

5.2. Downstream performance on TDC

Here, we present our fine-tuning results on the ADMET group datasets of TDC; see Table 4 for a comparison of different backbone GNNs in terms of downstream performance. GINE demonstrates a significant empirical advantage as the GNN backbone for downstream tasks. Thus, in Table 2, we compare MiniMol (GINE), to the TDC leaderboard and the current state-of-the-art, MoIE (Méndez-Lucio et al., 2022). MiniMol with GINE backbone, achieves top 1 performance on 8 tasks, setting a new state-of-the-art on these datasets. Moreover, MiniMol (GINE) achieves top-3 performance on 10 tasks. Therefore, MiniMol (GINE) is shown to be a versatile model across a wide range of tasks, competing with or exceeding the performance of the

Table 3. Results for GNN 10M baselines on LARGEMIX dataset. We report performance metrics on the test set for each dataset in LARGEMIX separately. The best scores per metric per dataset are marked in bold.

Dataset	Metric	Model		
		GCN	GINE	MPNN
PCQM4M_G25	MAE ↓	0.218	0.208	0.200
	Pearson ↑	0.884	0.889	0.892
	R^2 ↑	0.790	0.799	0.803
PCQM4M_N4	MAE ↓	0.025	0.022	0.021
	Pearson ↑	0.975	0.979	0.980
	R^2 ↑	0.952	0.959	0.961
PCBA_1328	CE ↓	0.033	0.033	0.033
	AUROC ↑	0.777	0.784	0.782
	AP ↑	0.286	0.302	0.287
L1000_VCAP	CE ↓	0.061	0.061	0.061
	AUROC ↑	0.500	0.514	0.500
	AP ↑	0.504	0.504	0.506
L1000_MCF7	CE ↓	0.059	0.058	0.059
	AUROC ↑	0.533	0.531	0.519
	AP ↑	0.513	0.516	0.514

Table 4. The effect of specific GNN architectures in the backbone of the fingerprinting model on the downstream performance. The rank is determined for each dataset individually, on a set of 7 scores, which include the test results from the TOP5 TDC leaderboard, MoIE and MiniMol. Here, all models used sum pooling, whereas our best model uses max pooling.

MiniMol backbone	Mean Rank	# Top1 Results	# Top3 Results
MPNN++	4.8	3	6
GCN	4.4	4	8
GINE	3.9	5	10

best task-specialized architectures. We report TDC leaderboard results up until January 2024. In addition, MiniMol (GINE) outperforms MoIE on 17 datasets, indicating that with only 10% of the parameters, our MiniMol approach is favourable to MoIE in downstream performance across many molecular tasks. For our best model, we explored different pooling methods, see Appendix A.3.

In Table 5, we compare MiniMol to other fingerprinting methods using the same evaluation downstream adaptation method as ours, so the only difference is the quality of generated fingerprints.

5.3. Correlation analysis

We also conduct a comprehensive correlation analysis to determine the impact of various pre-training datasets on the performance across 22 ADMET-group benchmarks from the

Table 5. Comparison of MiniMol to other molecular fingerprinting models using the same evaluation method as ours, including ensembles.

Model	Mean rank	>MoIE	TOP1	TOP3
MiniMol	3.6	17	8	10
AGBT (Chen et al., 2021a)	5.7	10	1	4
MolFormer (Ross et al., 2022)	5.7	7	0	4
BET (Chen et al., 2021b)	6.1	7	1	2

TDC benchmark. We checkpoint on two critical points in the training process: the epoch with the lowest total validation loss and at the end of training (of which these two epochs can differ). This approach yielded 32 unique combinations of pre-training validation and downstream test metrics.

Spearman’s rho correlation coefficients (Sedgwick, 2014) were calculated for each pair of pre-training and downstream metrics (see Tab. 6). The aim was to systematically identify which pre-training datasets have the most impact on downstream task performance, either positively or negatively. The p-value threshold was set at 0.1. Considering that some metrics indicate improvement when either increased or decreased, the correlation values were multiplied by the sign indicative of the direction in which improvement is registered for each metric. This adjustment provided a more nuanced understanding of whether enhancements in a pre-training metric are correlated with improvements in downstream performance.

We find an overall positive impact of our pre-training metrics on downstream performance thus demonstrating the effectiveness of MiniMol pre-training step. Interestingly for graph-level tasks on PCQM4M_G25, we find high pre-training results to be inversely correlated with strong performance on the ADMET group datasets. Specifically, a low MAE on PCQM4M_G25 results in negative effects across many downstream tasks. At the same time, we also find that the majority of pre-training metrics of LargeMix are highly informative of downstream performance, including the node-level tasks on PCQM4M. Further discussion is provided in Section 6.2 on possible explanations for these findings.

6. Discussion

Here, we discuss the differences in pre-training and fine-tuning performance for the different backbone GNNs of MiniMol as well as a discussion on the results of our correlation analysis.

6.1. Effect of model complexity

For the fine-tuning results in Section 4.2, our analysis reveals that while the three GNN backbone, namely, GCN,

Table 6. Correlation analysis (Spearman’s rho) between pre-training validation and downstream performance. The green color indicates positive correlation, and red a negative correlation. Results with a p-value over 10% are blank. See §5.3 for discussion.

Dataset	Metric	Overall Loss	L1000_MCF		L1000_VCAP		PCBA		PCQM4M_G25		PCQM4M_N4
			Loss	AUROC	Loss	AUROC	Loss	AUROC	Loss	MAE	Loss (MAE)
Caco2 Wang	MAE			0.590		0.651	0.762	0.718			
Bioavailability Ma	AUROC				0.397						
Lipophilicity AZ	MAE	0.459	0.561	0.568		0.539	0.683	0.627		-0.389	
Solubility AqSolDB	MAE	0.434	0.472	0.588		0.7	0.739	0.704			
HIA Hou	AUROC	0.427	0.489	0.603	0.382	0.548	0.768	0.645		-0.337	
Pgp Broccatelli	AUROC	0.444	0.577			0.361	0.497			-0.387	
BBB Martins	AUROC	0.634	0.436	0.583		0.378	0.481	0.483	0.364	-0.492	
PPBR AZ	MAE		0.476		0.436						
VDss Lombardo	Spearman	0.408				0.343	0.351				
CYP2C9 Veith	AUPRC	0.432		0.649		0.711	0.747	0.829	0.557		0.551
CYP2D6 Veith	AUPRC	0.566		0.641		0.487	0.624	0.704	0.616		0.585
CYP3A4 Veith	AUPRC	0.523		0.649		0.713	0.72	0.818	0.584		0.608
CYP2C9 Substrate	AUPRC		0.523		0.558	-0.377		-0.445	-0.566		-0.586
CYP2D6 Substrate	AUPRC										
CYP3A4 Substrate	AUROC					0.409	0.468				
Half Life Obach	Spearman					0.503	0.671	0.498			
Clearance Hepatocyte	Spearman		0.374								
Clearance Microsome	Spearman		0.599		0.496						
LD50 Zhu	MAE	0.501		0.543		0.522	0.586	0.617	0.339		0.342
hERG	AUROC					0.57	0.42	0.453			
AMES	AUROC	0.791		0.591		0.486	0.629	0.643	0.604	-0.628	0.528
DILI	AUROC		0.376	0.49		0.416	0.567	0.454			
	Sum	5.622	4.883	6.496	2.269	7.959	9.712	7.749	2.499	-2.232	2.028

GINE and MPNN++, all achieves similar pre-training performance, the GINE backbone shows a significant advantage when fine-tuning to downstream tasks. To give a potential explanation for this finding, recall that we adjust hidden dimensions of the different backbone GNNs to roughly align to 10M parameters. Here, the higher model complexity of MPNN++ leads to substantially smaller hidden dimension sizes than GINE. Our results thus indicate that the architectural complexity of the MPNN++ is less effective in downstream performance than a simple increase in hidden dimensions. If we match the hidden dimension size of GINE in MPNN, the model would reach the size of roughly 50M parameters.

At the same time, while less complex and allowing for even larger hidden dimensions, GCN layers might be expressive enough for strong downstream performance. Specifically, the GCN omits the use of edge features and is shown to be less powerful than the 1-WL test (Xu et al., 2019) while GINE is as expressive as the 1-WL test. As such, we hypothesize that the GINE allows for a trade-off between a sufficient level of architectural complexity and more effective use of parameter budget in terms of larger hidden dimensions that translate into stronger downstream performance.

Finally, we want to highlight that our results also reveal the general robustness of MiniMol pipeline to the choice of backbone GNNs, where all three variants were better

than the current state-of-the-art foundation model MolE on the ADMET group in many tasks (as seen in Table 4 and Table 2) while having significantly fewer parameters and being employed in an efficient fingerprinting pipeline, as opposed to fine-tuning all weights.

6.2. Detrimental effect of training on PCQM4M_g25

From the correlation analysis, results indicate that the PCQM4M_G25 dataset negatively impacts downstream performance in the ADMET group benchmark. This dataset is unique in that it is the only graph-level quantum task. As a result, we might conclude that atomic and sub-atomic properties of the molecules in PCQM4M_G25 are only marginally relevant in the context of biological tasks, resulting in a detriment to performance on downstream tasks.

However, this explanation is insufficient, so far that the PCQM4M_N4 dataset, which is a node-level quantum task that covers the same molecules as PCQM4M_G25, demonstrates positive downstream benefits. As a result, we hypothesize that the issue rather lies in the training dynamics. Concretely, models that perform well on PCQM4M_G25 might do so at the expense of overfitting, which is expected for multi-task learning on an unbalanced number of targets (Beaini et al., 2023). Indeed, as already described in Section 4.1, reducing the influence of PCQM4M_G25 on the overall loss led to improved overall performance, supporting our hypothesis.

Our findings highlight the importance of carefully designing the overall loss when pre-training on multi-task datasets such as LargeMix. Such a design should consider data imbalance to ensure positive impacts on downstream performance. This is particularly important for foundation models whose downstream applications are mostly unknown at the time of pre-training.

7. Conclusion

In this work, we propose a novel parameter-efficient foundation model for molecular learning called MiniMol. MiniMol is pre-trained on over 3300 biological and quantum tasks on graph- and node-level molecules and subsequently evaluated on the ADMET group of the TDC benchmark. MiniMol outperforms the previous state-of-the-art foundation model on the ADMET group, MoLE, with only 10M parameters, constituting only 10% of MoLE’s 100M parameters. In addition, fine-tuning with MLPs on the fingerprints of pre-trained MiniMol, allows for efficient fine-tuning.

Our empirical results showed that the molecular fingerprints extracted from MiniMol are highly transferable to downstream tasks. MiniMol achieves top 1 performance on 8 tasks and top 3 performance on 10 tasks on the ADMET group. In addition, our task correlation analysis highlighted the importance of carefully designing the overall loss for multi-task pre-training. With MiniMol, we showcased the potential for parameter-efficient multi-task multi-level pre-training and fine-tuning on fingerprints. One future direction is to design pre-training datasets that align with a wider variety of downstream tasks. We will open source MiniMol and believe it will be an important tool for future research on molecular foundation models.

8. Broader Impact

Releasing our Model may not have immediate direct societal impacts. However, it is crucial to acknowledge the potential implications that arise when providing access to a foundation model for molecular graphs. One concerning possibility is the misuse of this technology for the development of chemical weapons, toxins, or unregulated drugs. To address these potential risks, we are committed to implementing robust mitigation strategies. Central to our approach is the active promotion of beneficial applications, particularly in the fields of material and drug discovery. By highlighting the positive utilization of this technology, we aim to channel its potential toward scientific advancements that contribute to societal well-being.

References

- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Ahmad, W., Simon, E., Chithrananda, S., Grand, G., and Ramsundar, B. ChemBERTa-2: Towards Chemical Foundation Models, September 2022. URL <http://arxiv.org/abs/2209.01712>. arXiv:2209.01712 [cs, q-bio].
- Awale, M. and Reymond, J.-L. Atom pair 2d-fingerprints perceive 3d-molecular shape and pharmacophores for very fast virtual screening of zinc and gdb-17. *Journal of chemical information and modeling*, 54(7):1892–1907, 2014.
- Beaini, D., Huang, S., Cunha, J. A., Li, Z., Moisescu-Pareja, G., Dymov, O., Maddrell-Mander, S., McLean, C., Wenkel, F., Müller, L., Mohamud, J. H., Parviz, A., Craig, M., Koziarski, M., Lu, J., Zhu, Z., Gabellini, C., Klaser, K., Dean, J., Wognum, C., Sypetkowski, M., Rabusseau, G., Rabbany, R., Tang, J., Morris, C., Koutis, I., Ravanelli, M., Wolf, G., Tossou, P., Mary, H., Bois, T., Fitzgibbon, A., Banaszewski, B., Martin, C., and Masters, D. Towards Foundational Models for Molecular Learning on Large-Scale Multi-Task Datasets, October 2023. URL <http://arxiv.org/abs/2310.04292>. arXiv:2310.04292 [cs].
- Beaini, D., Huang, S., Cunha, J. A., Moisescu-Pareja, G., Dymov, O., Maddrell-Mander, S., McLean, C., Wenkel, F., Müller, L., Mohamud, J. H., et al. Towards foundational models for molecular learning on large-scale multi-task datasets. 2024.
- Betker, J., Goh, G., Jing, L., Brooks, T., Wang, J., Li, L., Ouyang, L., Zhuang, J., Lee, J., Guo, Y., et al. Improving image generation with better captions. *Computer Science*. <https://cdn.openai.com/papers/dall-e-3.pdf>, 2:3, 2023.
- Borzunov, A., Ryabinin, M., Chumachenko, A., Baranchuk, D., Dettmers, T., Belkada, Y., Samygin, P., and Raffel, C. Distributed inference and fine-tuning of large language models over the internet. *arXiv preprint arXiv:2312.08361*, 2023.
- Capecchi, A., Probst, D., and Reymond, J.-L. One molecular fingerprint to rule them all: drugs, biomolecules, and the metabolome. *Journal of cheminformatics*, 12(1):1–15, 2020.
- Chen, D., Gao, K., Nguyen, D. D., et al. Algebraic graph-assisted bidirectional transformers for molecular property prediction. *Nat Commun*, 12:3521, 2021a. doi: 10.1038/s41467-021-23720-w.
- Chen, D., Zheng, J., Wei, G.-W., and Pan, F. Extracting predictive representations from hundreds of millions of molecules. *The Journal of Physical Chemistry Letters*, 12(44):10793–10801, 2021b. doi: 10.1021/acs.jpcclett.1c03058.
- Crawshaw, M. Multi-task learning with deep neural networks: A survey. *arXiv preprint arXiv:2009.09796*, 2020.
- Dwivedi, V. P. and Bresson, X. A Generalization of Transformer Networks to Graphs, December 2020. URL <https://arxiv.org/abs/2012.09699v2>.
- Galkin, M., Yuan, X., Mostafa, H., Tang, J., and Zhu, Z. Towards foundation models for knowledge graph reasoning. *CoRR*, abs/2310.04562, 2023. doi: 10.48550/ARXIV.2310.04562. URL <https://doi.org/10.48550/arXiv.2310.04562>.
- Gasteiger, J., Groß, J., and Günnemann, S. Directional message passing for molecular graphs. In *International Conference on Learning Representations*, 2019.
- Honda, S., Shi, S., and Ueda, H. R. SMILES Transformer: Pre-trained Molecular Fingerprint for Low Data Drug Discovery, November 2019. URL <http://arxiv.org/abs/1911.04738>. arXiv:1911.04738 [cs, stat].
- Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., and Leskovec, J. Open graph benchmark: Datasets for machine learning on graphs. *Advances in neural information processing systems*, 33:22118–22133, 2020a.
- Hu, W., Liu, B., Gomes, J., Zitnik, M., Liang, P., Pande, V., and Leskovec, J. Strategies for Pre-training Graph Neural Networks, February 2020b. URL <http://arxiv.org/abs/1905.12265>. arXiv:1905.12265 [cs, stat].
- Hu, W., Fey, M., Ren, H., Nakata, M., Dong, Y., and Leskovec, J. Ogb-lsc: A large-scale challenge for machine learning on graphs. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- Huang, K., Fu, T., Gao, W., Zhao, Y., Roohani, Y., Leskovec, J., Coley, C., Xiao, C., Sun, J., and Zitnik, M. Therapeutics data commons: Machine learning datasets and tasks for drug discovery and development. *Advances in neural information processing systems*, 2021.
- Jin, W., Stokes, J. M., Eastman, R. T., Itkin, Z., Zakharov, A. V., Collins, J. J., Jaakkola, T. S., and Barzilay, R. Deep learning identifies synergistic drug combinations for treating COVID-19. *Proceedings of the National Academy of Sciences of the United States of America*, 118(39):e2105070118, September 2021. ISSN 1091-6490. doi: 10.1073/pnas.2105070118.

- Kaufman, B., Williams, E., Underkoffler, C., Pederson, R., Mardirossian, N., Watson, I., and Parkhill, J. COATI: multi-modal contrastive pre-training for representing and traversing chemical space, August 2023. URL <https://chemrxiv.org/engage/chemrxiv/article-details/64e8137fdd1a738447f73f7aa>.
- Kim, S. Exploring chemical information in pubchem. *Current protocols*, 1(8):e217, 2021.
- Kipf, T. N. and Welling, M. Semi-Supervised Classification with Graph Convolutional Networks, September 2016. URL <https://arxiv.org/abs/1609.02907v4>.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. 2017.
- Landrum, G. et al. Rdkit: A software suite for cheminformatics, computational chemistry, and predictive modeling. *Greg Landrum*, 8:31, 2013.
- Li, J., Cai, D., and He, X. Learning Graph-Level Representation for Drug Discovery, September 2017. URL <http://arxiv.org/abs/1709.03741>. arXiv:1709.03741 [cs, stat].
- Liu, J., Yang, C., Lu, Z., Chen, J., Li, Y., Zhang, M., Bai, T., Fang, Y., Sun, L., Yu, P. S., and Shi, C. Towards Graph Foundation Models: A Survey and Beyond, December 2023. URL <http://arxiv.org/abs/2310.11829>. arXiv:2310.11829 [cs].
- Lu, S., Gao, Z., He, D., Zhang, L., and Ke, G. Highly accurate quantum chemical property prediction with unimol+. *arXiv preprint arXiv:2303.16982*, 2023.
- Luo, Y., Yang, K., Hong, M., Liu, X. Y., and Nie, Z. MolFM: A Multimodal Molecular Foundation Model, July 2023. URL <http://arxiv.org/abs/2307.09484>. arXiv:2307.09484 [physics, q-bio].
- Masters, D., Dean, J., Klaser, K., Li, Z., Maddrell-Mander, S., Sanders, A., Helal, H., Beker, D., Fitzgibbon, A., Huang, S., Rampášek, L., and Beaini, D. Gps++: Reviving the art of message passing for molecular property prediction. *Transactions on Machine Learning Research*, 2023a.
- Masters, D., Dean, J., Klaser, K., Li, Z., Maddrell-Mander, S., Sanders, A., Helal, H., Beker, D., Fitzgibbon, A., Huang, S., Rampášek, L., and Beaini, D. GPS++: Reviving the Art of Message Passing for Molecular Property Prediction, February 2023b. URL <http://arxiv.org/abs/2302.02947>. arXiv:2302.02947 [cs].
- Méndez-Lucio, O., Nicolaou, C., and Earnshaw, B. MolE: a molecular foundation model for drug discovery, November 2022. URL <http://arxiv.org/abs/2211.02657>. arXiv:2211.02657 [cs, q-bio].
- Nakata, M. and Shimazaki, T. Pubchemqc project: a large-scale first-principles electronic structure database for data-driven chemistry. *Journal of chemical information and modeling*, 57(6):1300–1308, 2017.
- Probst, D., Schwaller, P., and Reymond, J.-L. Reaction classification and yield prediction using the differential reaction fingerprint DRFP. *Digital Discovery*, 1(2):91–97, 2022. doi: 10.1039/D1DD00006C. URL <https://pubs.rsc.org/en/content/articlelanding/2022/dd/d1dd00006c>. Publisher: Royal Society of Chemistry.
- Rampášek, L., Galkin, M., Dwivedi, V. P., Luu, A. T., Wolf, G., and Beaini, D. Recipe for a general, powerful, scalable graph transformer. *Advances in Neural Information Processing Systems*, 35:14501–14515, 2022.
- Rampášek, L., Galkin, M., Dwivedi, V. P., Luu, A. T., Wolf, G., and Beaini, D. Recipe for a General, Powerful, Scalable Graph Transformer. Technical Report arXiv:2205.12454, arXiv, May 2022. URL <http://arxiv.org/abs/2205.12454>. arXiv:2205.12454 [cs] type: article.
- Reiser, P., Neubert, M., Eberhard, A., Torresi, L., Zhou, C., Shao, C., Metni, H., van Hoesel, C., Schopmans, H., Sommer, T., and Friederich, P. Graph neural networks for materials science and chemistry. *Communications Materials*, 3(1):1–18, November 2022. ISSN 2662-4443. doi: 10.1038/s43246-022-00315-6. URL <https://www.nature.com/articles/s43246-022-00315-6>. Number: 1 Publisher: Nature Publishing Group.
- Rogers, D. and Hahn, M. Extended-connectivity fingerprints. *Journal of chemical information and modeling*, 50(5):742–754, 2010.
- Ross, J., Belgodere, B., Chenthamarakshan, V., et al. Large-scale chemical language representations capture molecular structure and properties. *Nat Mach Intell*, 4:1256–1264, 2022. doi: 10.1038/s42256-022-00580-7.
- Sedgwick, P. Spearman’s rank correlation coefficient. *Bmj*, 349, 2014.
- Shoghi, N., Kolluru, A., Kitchin, J. R., Ulissi, Z. W., Zitnick, C. L., and Wood, B. M. From Molecules to Materials: Pre-training Large Generalizable Models for Atomic Property Prediction, October 2023. URL <http://arxiv.org/abs/2310.16802>. arXiv:2310.16802 [cs].

- Stevens, J. A., Grünewald, F., van Tilburg, P. A. M., König, M., Gilbert, B. R., Brier, T. A., Thornburg, Z. R., Luthey-Schulten, Z., and Marrink, S. J. Molecular dynamics simulation of an entire cell. *Frontiers in Chemistry*, 11, 2023. ISSN 2296-2646. doi: 10.3389/fchem.2023.1106495. URL <https://www.frontiersin.org/articles/10.3389/fchem.2023.1106495>.
- Stokes, J. M., Yang, K., Swanson, K., Jin, W., Cubillos-Ruiz, A., Donghia, N. M., MacNair, C. R., French, S., Carfrae, L. A., Bloom-Ackermann, Z., Tran, V. M., Chiappino-Pepe, A., Badran, A. H., Andrews, I. W., Chory, E. J., Church, G. M., Brown, E. D., Jaakkola, T. S., Barzilay, R., and Collins, J. J. A Deep Learning Approach to Antibiotic Discovery. *Cell*, 180(4):688–702.e13, February 2020. ISSN 1097-4172. doi: 10.1016/j.cell.2020.01.021.
- Subramanian, A., Narayan, R., Corsello, S. M., Peck, D. D., Natoli, T. E., Lu, X., Gould, J., Davis, J. F., Tubelli, A. A., Asiedu, J. K., et al. A next generation connectivity map: L1000 platform and the first 1,000,000 profiles. *Cell*, 171(6):1437–1452, 2017.
- Taylor, R., Kardas, M., Cucurull, G., Scialom, T., Hartshorn, A., Saravia, E., Poulton, A., Kerkez, V., and Stojnic, R. Galactica: A large language model for science. *CoRR*, abs/2211.09085, 2022.
- Team, G., Anil, R., Borgeaud, S., Wu, Y., Alayrac, J.-B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A. M., Hauth, A., et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- Tian, K., Mitchell, E., Yao, H., Manning, C. D., and Finn, C. Fine-tuning language models for factuality. *arXiv preprint arXiv:2311.08401*, 2023.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. Graph Attention Networks, October 2017. URL <https://arxiv.org/abs/1710.10903v3>.
- Wallach, I., Dzamba, M., and Heifets, A. AtomNet: A Deep Convolutional Neural Network for Bioactivity Prediction in Structure-based Drug Discovery, October 2015. URL <https://arxiv.org/abs/1510.02855v1>.
- Wang, S., Guo, Y., Wang, Y., Sun, H., and Huang, J. SMILES-BERT: Large Scale Unsupervised Pre-Training for Molecular Property Prediction. *BCB '19*, pp. 429–436, New York, NY, USA, September 2019. Association for Computing Machinery. ISBN 978-1-4503-6666-3. doi: 10.1145/3307339.3342186. URL <https://doi.org/10.1145/3307339.3342186>.
- Weininger, D. 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, February 1988. ISSN 0095-2338. doi: 10.1021/ci00057a005. URL <https://doi.org/10.1021/ci00057a005>. Publisher: American Chemical Society.
- Weininger, D., Weininger, A., and Weininger, J. L. Smiles. 2. algorithm for generation of unique smiles notation. *Journal of chemical information and computer sciences*, 29(2):97–101, 1989.
- Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How Powerful are Graph Neural Networks? *arXiv:1810.00826 [cs, stat]*, February 2019. URL <http://arxiv.org/abs/1810.00826>. arXiv: 1810.00826.
- Ying, C., Cai, T., Luo, S., Zheng, S., Ke, G., He, D., Shen, Y., and Liu, T.-Y. Do Transformers Really Perform Bad for Graph Representation?, November 2021. URL <http://arxiv.org/abs/2106.05234>. arXiv:2106.05234 [cs].
- Zhu, J., Xia, Y., Qin, T., Zhou, W., Li, H., and Liu, T.-Y. Dual-view Molecule Pre-training, October 2021. URL <http://arxiv.org/abs/2106.10234>. arXiv:2106.10234 [cs, q-bio].
- Zitnick, C. L., Chanussot, L., Das, A., Goyal, S., Heras-Domingo, J., Ho, C., Hu, W., Lavril, T., Palizhati, A., Riviere, M., Shuaibi, M., Sriram, A., Tran, K., Wood, B., Yoon, J., Parikh, D., and Ulissi, Z. An Introduction to Electrocatalyst Design using Machine Learning for Renewable Energy Storage, October 2020. URL <http://arxiv.org/abs/2010.09435>. arXiv:2010.09435 [cond-mat].

A. Appendix

A.1. MPNN architecture

In what follows, we describe the MPNN architecture in (Masters et al., 2023b) in detail. Here, the embeddings are incrementally updated with each MPNN layer in the model as:

$$x^{\ell+1}, e^{\ell+1}, g^{\ell+1} = \text{MPNN}(x^\ell, e^\ell, g^\ell) \quad (4)$$

The edge embedding is updated by concatenation of the edge feature with the node features at each end of the bond, with the global features. This is processed with the edge MLP and then summed with the skip connection, shown in 5.

$$\bar{e}_{uv}^\ell = \text{MLP}_{\text{edge}} \left([x_u^\ell | x_v^\ell | e_{uv}^\ell | g^\ell] \right) \quad (5)$$

The node embedding, shown in eq.6 concatenates the node features with the summed edge features of all edges connected (senders and receiver) and the global features before passing this vector through an MLP and finally adding the skip connection.

$$\bar{x}_i^\ell = \text{MLP}_{\text{node}} \left(\left[x_i^\ell \left| \sum_{(u,i) \in \mathcal{E}} \bar{e}_{ui}^\ell \right| \sum_{(i,v) \in \mathcal{E}} \bar{e}_{iv}^\ell \left| \sum_{(u,i) \in \mathcal{E}} x_u^\ell \right| g^\ell \right] \right) \quad (6)$$

The global node is concatenated with the sum of all node and edge features in the graph (eq. 7).

$$\bar{g}^\ell = \left[g^\ell \left| \sum_{j \in \mathcal{V}} \bar{x}_j^\ell \right| \sum_{(u,v) \in \mathcal{E}} \bar{e}_{uv}^\ell \right] \quad (7)$$

Where the final components are computed with skip-connections as:

$$x_i^{\ell+1} = \bar{x}_i^\ell + x_i^\ell; \quad e_{uv}^{\ell+1} = \bar{e}_{uv}^\ell + e_{uv}^\ell; \quad g^{\ell+1} = \bar{g}^\ell + g^\ell; \quad (8)$$

This is represented diagrammatically in Fig. 3.

A.2. Hyperparameter selection

We select hyperparameters for our fine-tuning as follows. We compute a hyperparameter sweep over the maximum number of epochs; the learning rate; the dropout rate and whether to use none, batch or layer normalization in the task head. Optionally, we sweep over the width and depth of the task head MLP. Each configuration is run on the same random seed. Following the instructions provided by TDC¹, we use the provided scaffold splits for our train/validation splits via the method `get_train_valid_split` and take the benchmark test split also provided by TDC. Then, for each dataset, we select the hyperparameters resulting from the model with the smallest validation loss and subsequently re-run this model on k random seeds. Here, we distinguish between two sweep configurations. In the first configuration, we only sweep over the learning rate $\in \{0.001, 0.0005, 0.0003, 0.0001, 5e^{-5}\}$ and set the number of epochs to 25, dropout to 0.1, hidden dimension to 1024 and the number of layers to 3. In the second configuration, we set the number of epochs to 25 and sweep over whether or not to use a skip connection, the learning rate $\in \{0.0005, 0.0003, 0.0001\}$, the hidden dimension $\in \{512, 1024, 2048\}$, the number of layers $\in \{3, 4\}$, dropout $\in \{0.0, 0.1\}$, the number of warmup epochs $\in \{0, 5\}$ and the learning rate schedule $\in \{\text{constant, linear, cosine}\}$.

¹Available at <https://tdcommons.ai/benchmark/overview/>

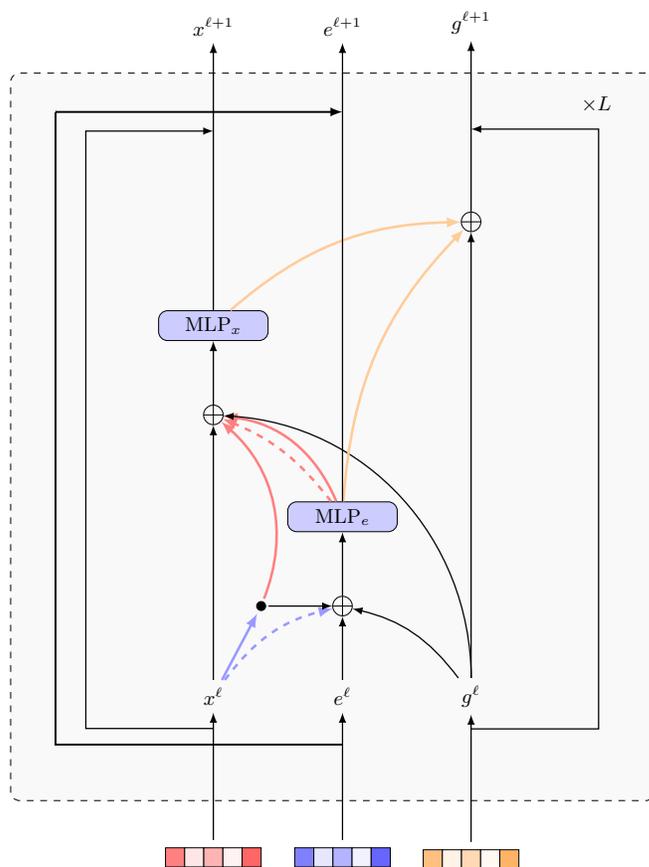


Figure 3. Example of the MPNN block architecture given in Eq.4. The edge update in Eq.5 gathers the nodes and edges before passing through the MLP first, then this output is used for the node update in Eq.6, gathering all connected node features and updated edge features. The global update in Eq.7 connects all nodes and edges, before finally the skip connections in Eq.8.

A.3. Experimentation with pooling methods

We evaluated three different pooling strategies when going from the node level to graph level representation and summarized our findings in A.3.

	Mean rank	MoIE	TOP1	TOP3
sum	3.9	16	5	10
mean	3.8	16	6	11
max	3.6	17	8	10

A.4. Ensemble Strategy

The strategy used for evaluating and ensembling the models is explained in the form of pseudo-code below.

Algorithm 1 Ensembling Strategy

```
Input: hyperparameters  $h_i$ , repetitions  $rep_i$ , fold  $fold_i$   
for each  $h_i$  in  $[hp_1, \dots, hp_{18}]$  do  
  for  $rep_i$  in  $range(num_{reps})$  do  
    select  $seed$   
    for  $fold_i$  in  $range(num_{folds})$  do  
      train a model on  $fold_i$   
      save best model based on val loss  
    end for  
    build ensemble of  $num_{folds}$  models  
    evaluate on ensemble  
    save mean and std of val and test scores  
  end for  
end for
```
