
Evaluating Self-Supervised Learning for Molecular Graph Embeddings

Hanchen Wang^{1*} Jean Kaddour^{2*} Shengchao Liu^{3,4*} Jian Tang^{3,5,6} Joan Lasenby¹ Qi Liu⁷
¹Cambridge ²UCL ³MILA ⁴UdeM ⁵HEC ⁶CIFAR ⁷HKU * Equal Contribution

Abstract

Graph Self-Supervised Learning (GSSL) provides a robust pathway for acquiring embeddings without expert labelling, a capability that carries profound implications for molecular graphs due to the staggering number of potential molecules and the high cost of obtaining labels. However, GSSL methods are designed not for optimisation within a specific domain but rather for transferability across a variety of downstream tasks. This broad applicability complicates their evaluation. Addressing this challenge, we present "Molecular Graph Representation Evaluation" (MOLGRAPHEVAL), generating detailed profiles of molecular graph embeddings with interpretable and diversified attributes. MOLGRAPHEVAL offers a suite of probing tasks grouped into three categories: (i) generic graph, (ii) molecular substructure, and (iii) embedding space properties. By leveraging MOLGRAPHEVAL to benchmark existing GSSL methods against both current downstream datasets and our suite of tasks, we uncover significant inconsistencies between inferences drawn solely from existing datasets and those derived from more nuanced probing. These findings suggest that current evaluation methodologies fail to capture the entirety of the landscape.

1 Introduction

Learning neural embeddings of molecular graphs has become of paramount importance in computer-aided drug discovery [1, 2]. For instance, a molecular property prediction (MPP) model can expedite and economise the design process by reducing the need for synthesising and measuring molecules. Thereby, such models can be immensely useful in the hit-to-lead and early lead optimisation phase of a drug discovery project [3]. However, obtaining labels of molecule properties is expensive and time-consuming, especially since the size of potential pharmacologically active molecules is estimated to be in the order of 10^{60} [4, 5].

Graph Self-Supervised Learning (GSSL) paves the way for learning molecular graph embeddings without human annotations that are transferable to various downstream datasets. Unfortunately, the evaluation of such general-purpose embeddings is fundamentally complex. Different proxy objectives will place different demands on them, and no single downstream dataset can be definitive. Moreover, many of the previously proposed GSSL works are disconnected in terms of the tasks they target and the datasets they use for evaluation, making direct comparison difficult.

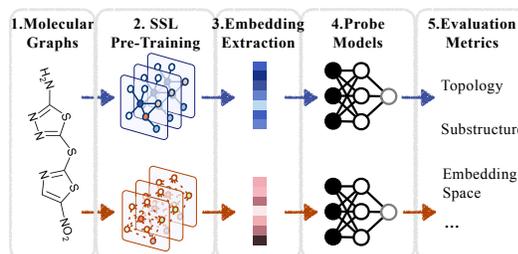


Figure 1: Overview of MOLGRAPHEVAL. Given molecular graphs, we train GNNs to predict SSL proxy objectives. We then extract embeddings of (possibly unseen) graphs using pre-trained models, which form the inputs for probe models, trained and evaluated on the designed metrics.

Contributions. Our goal is to unbiasedly evaluate molecular graph embeddings obtained by GSSL methods on existing downstream tasks and a new suite of probe tasks (Fig. 1). We summarised some key findings based on a total of **90,918 probe models** and **1,875 pre-trained GNNs**.

- 1) On MPP tasks, we observe every GSSL method can introduce substantial performance gains. Yet, there is a significant difference in the rank depending on whether we fine-tune the pre-trained network on the downstream dataset or not. Also, the pre-training configurations for obtaining the optimal embeddings or initialisation for fine-tuning are different; see - **Finding 2**.
- 2) Several discrepancies between MPP tasks and MOLGRAPHEVAL demonstrate how the latter complements GSSL evaluation with novel insights fostering future work:
 - While embeddings from a randomly initialised GNN perform poorly on MPP tasks, they sometimes outperform on topological properties (which can be useful in some molecular tasks), indicating that GSSL methods are not universally better, see - **Finding 4**.
 - While contrastive GSSL methods perform better than most other methods on MPP tasks, it might attribute to their superiority in identifying the crucial substructures, see - **Finding 8**.
 - In contrast to previous work [6], we find that feature distribution uniformity is not always a strong indicator for MPP performance. For instance, maximising the mutual information between multi-scale graph representations (INFOGRAPH) results in the most uniform distributions, yet, it ranks 7th among the 9 GSSLs on the downstream MPP tasks, see - **Finding 9**.

2 Related work

Graph SSL (GSSL) can be divided into contrastive and generative methods [7–9]. Contrastive GSSL [10–12] construct multiple views of the same graph via augmentations and then learn embeddings by aligning positive samples against negative ones. Generative GSSL [11, 13–15] yields embeddings by reconstructing input graphs. Zhu et al. [16] conduct an empirical analysis of contrastive GSSL methods and their components. In contrast, we investigate both generative and contrastive GSSL methods and propose a novel suite of tasks to probe the learned embeddings’ attributes.

Probe models and benchmarks on graphs. Probe models, trained exclusively on embedding vectors from pre-trained models, serve as an effective tool for evaluating the quality of learned embeddings [17]. Their effectiveness has been demonstrated across various domains such as language [18–23], vision [24–28], relational tables [29], and science [30–32]. While there exist benchmarks for graph learning [33–37], applying probe models to GSSL remains an unexplored frontier.

3 Preliminaries

Graph. A graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ consists of a set of nodes \mathcal{V} and edges \mathcal{E} . In molecular graphs, nodes are atoms, and edges are bonds. We use \mathbf{x}_u and \mathbf{x}_{uv} to denote the feature of node u and the bond feature between nodes $[u, v]$, respectively. For notation simplicity, we use an adjacency matrix $\mathbf{A} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ to represent the graph, where $\mathbf{A}[u, v] \neq 0$ if the nodes (u, v) are connected.

GNN. Graph neural networks (GNNs) give rise to learning molecular graph embeddings [14, 38–41]. A prototypical GNN relies on messaging passing [39], which updates atom-level embeddings based on their neighbourhoods. Given an input atom $\mathbf{h}_u^0 = \mathbf{x}_u$, we compute its embedding by:

$$\mathbf{m}_u^{t+1} = \sum_{v:\mathbf{A}[u,v] \neq 0} M_t(\mathbf{h}_u^t, \mathbf{h}_v^t, \mathbf{x}_{uv}) \quad \mathbf{h}_u^{t+1} = U_t(\mathbf{h}_u^t, \mathbf{m}_u^{t+1}) \quad (1)$$

where M_t and U_t are the “message” functions and “vertex update” functions, respectively. Repeating message passing for T steps, the embedding of each atom contains their T -hop neighbourhood information. A readout function R is then used to pool node-level embeddings for graph-level representations: $\hat{\mathbf{y}} = R(\{\mathbf{h}_u^T \mid u \in \mathcal{V}\})$. Following previous GSSL methods on molecular graphs, we adopt the Graph Isomorphism Network (GIN) [42] as the backbone model and incorporate edge features during message passing following [11].

Pre-Training. We inspect nine GSSL methods (1,875 configurations in total): EDGE PRED [13], INFOGRAPH [10], GPT-GNN [15], ATTRMASK [11], CONTEXT PRED [11], GROVER [43], GRAPHCL [12], JOAO [44], and GRAPH MVP [45]. We use all qualified molecules (around 0.33 million, *i.e.*, leave out the molecules that appeared in downstream datasets) from the GEOM dataset [46] to pre-train the GIN backbone. As many of these pre-training methods are not primarily designed for molecular graphs, we perform the grid search over the hyperparameter space and save

Table 1: **Evaluating GSSL methods on molecular property prediction tasks.** For each downstream dataset, we report the mean and standard deviation of the ROC-AUC scores over three random scaffold splits. The best and second best scores are marked **bold** and **bold**, respectively. The performance scores are based on the fixed pre-trained embeddings with linear probe models, we also report the average ROC-AUC scores with fine-tuned pre-trained GNN on MPP tasks (“Avg (FT)”). For each pre-training method, we report the highest scores in the table and their corresponding hyperparameter configurations in Tables 7 and 9 in Appendix B.

	BBBP	Tox21	ToxCast	Sider	ClinTox	MUV	HIV	Bace	Avg	Avg (FT)
# Molecules	2,039	7,831	8,575	1,427	1,478	93,087	41,127	1,513	–	–
# Tasks	1	12	617	27	2	17	1	1	–	–
RANDOM	50.7 \pm 2.5	64.9 \pm 0.5	53.2 \pm 0.3	53.2 \pm 1.1	63.1 \pm 2.3	62.1 \pm 1.3	66.1 \pm 0.7	63.4 \pm 1.8	59.60	66.16
EDGE PRED	54.2 \pm 1.0	66.2 \pm 0.2	54.4 \pm 0.1	56.1 \pm 0.1	65.4 \pm 5.0	59.5 \pm 0.9	73.6 \pm 0.4	71.4 \pm 1.2	62.59	68.16
ATTR MASK	62.7 \pm 2.7	65.7 \pm 0.8	56.1 \pm 0.2	58.3 \pm 1.5	61.9 \pm 6.4	60.9 \pm 1.8	65.5 \pm 1.4	64.8 \pm 2.6	61.99	69.20
GPT-GNN	62.0 \pm 0.9	64.9 \pm 0.7	55.4 \pm 0.2	55.3 \pm 0.8	55.0 \pm 5.1	61.2 \pm 1.5	71.2 \pm 1.5	61.0 \pm 1.2	60.74	67.58
INFOGRAPH	65.9 \pm 0.6	65.8 \pm 0.7	54.6 \pm 0.1	57.2 \pm 1.0	61.4 \pm 4.8	63.9 \pm 1.9	71.4 \pm 0.6	67.4 \pm 4.9	63.44	68.92
CONT. PRED	55.5 \pm 2.0	67.9 \pm 0.7	54.0 \pm 0.3	57.1 \pm 0.5	67.4 \pm 4.3	60.5 \pm 0.9	66.2 \pm 1.5	54.4 \pm 3.2	60.36	69.40
GROVER	67.0 \pm 0.3	63.9 \pm 0.3	53.6 \pm 0.4	59.9 \pm 1.7	65.0 \pm 6.4	62.7 \pm 1.4	67.8 \pm 1.0	69.0 \pm 4.7	63.62	69.97
GRAPHCL	64.7 \pm 1.7	69.1 \pm 0.5	56.2 \pm 0.2	59.5 \pm 0.9	60.8 \pm 3.0	60.6 \pm 1.8	72.5 \pm 1.4	77.0 \pm 1.7	65.04	70.33
JOAO	66.1 \pm 0.8	68.1 \pm 0.2	55.1 \pm 0.4	58.3 \pm 0.3	65.3 \pm 6.1	62.4 \pm 1.2	73.8 \pm 1.2	71.1 \pm 0.8	65.05	69.75
GRAPHMVP	69.2 \pm 1.8	63.8 \pm 0.3	55.5 \pm 0.3	58.6 \pm 0.4	58.7 \pm 1.9	63.8 \pm 1.3	68.6 \pm 1.0	73.3 \pm 4.7	63.92	70.06

the optimal settings. For these nine GSSL methods, we have pre-trained 1,875 GNNs with different configurations, as elaborated in Appendix B. We extract embeddings using the pre-trained weights, select the optimal hyperparameter sets based on their downstream MPP performance and use these optimal embeddings for further probing tasks.

Probe. We use probe models [18] to study whether self-supervised learned embeddings encode helpful structural information about graphs. Concretely, we extract embeddings from a pre-trained GNN and train a linear model to predict the probe tasks with node and graph embeddings as inputs. As the first work that designs probe methods on graph embeddings, we follow previous works on computer vision and natural language processing. We mainly compute and compare the quality of pre-trained embeddings using linear probe models. We have also experimented MLPs with one hidden layer as the probe models, as this architecture is utilised in some previous works. We observe similar findings with both probe architectures and reported the results of MLP probes in Appendix B.4. We use scaffold split to partition data into 80%/10%/10% for the training/validation/testing set. The training procedure runs for 100 epochs with a fixed learning rate of 0.001. We report the test results based on the best validation scores. To account for statistical significance, we average all experimental results over three independent runs. We find that different data splits are the primary cause for performance variations (\sim 2%), instead of initialising probe models with different random seeds ($<$ 0.01%).

4 Benchmarking GSSL on MPP

We first conduct a rigorous empirical investigation of the GSSL methods’ effectiveness in predicting the biochemical properties of molecules. Following previous work [11, 12], we consider eight molecular datasets consisting of 678 binary property prediction tasks [47, 48]. Unless explicitly stated otherwise, we extract the node/graph embeddings from the last GNN layer. We devise two settings: (i) fixed embeddings, where we train the probe models with fixed embeddings extracted from pre-trained GNNs; (ii) fine-tuned embeddings (“FT”), where we update weights of both the pre-trained GNNs and the probe models. Setting (i) follows the procedures in previous probing literature, while (ii) is widely utilised as the “pre-training, then fine-tuning” paradigm. We use Adam optimiser with no weight decay, set the batch size as 256, and apply identical pre-processing procedures for all experiments.

Findings. Table 1 notes the results, and we summarise the following findings, some of which contrast with those drawn from the concurrent study [49].

- 1) **All GSSL methods perform better than RANDOM.** By carefully exploring the pre-training hyperparameters, all GSSLs substantially improve the MPP tasks for both fixed and fine-tuned embeddings. Contrastive-based GSSL methods (*i.e.*, GRAPHCL, JOAO and GRAPHMVP) achieve the overall best performance. As [49] declares that molecular graph pretraining is ineffective; however, we find that their conclusions are based on a few selected finetuning datasets and fixed

pre-training hyperparameters. We further observe that such improvements will reduce when the number of molecules in downstream datasets increases. Specifically, for MUV [50], a dataset designed for validating virtual screening (used in drug discovery to find how likely molecules that bind to a drug target), the average performance gain brought by pre-training is -0.3%; while for BBBP, it is 12.3%. The number of molecules in BBBP is only 2% of the MUV's.

- 2) **Rankings differ between probing and fine-tuning.** The rank correlation between the fixed and fine-tuned embeddings is 0.77 (p-value=9e-4), indicating that we cannot utilise the rank of fixed embeddings as a definite indicator for fine-tuning performance, though they are positively correlated. Part of this observation has been spotted in a study on masked visual transformers [25]. In the context of molecular property prediction, embeddings pre-trained with JOAO achieve the best score with fixed scenarios but perform the fourth after end-to-end fine-tuning. The reason is unclear and should be investigated by future work.
- 3) **The optimal sets of pre-train hyperparameters for fixed and fine-tuned embeddings vary.** We observe that the optimal pre-training hyperparameters on fixed and fine-tuned embeddings differ. Only two out of nine GSSLs (INFOGRAPH and EDGE PRED) share the same set of optimal parameters, as detailed in Tables 7 and 9 in Appendix B. This suggests that probing the fixed embeddings might not truly reflect pre-trained models' performance on downstream MPP tasks, as it ignores the consequent improvements induced by fine-tuning. In Fig. 2, we visualised the hyperparameter space of the ATTRMASK pre-trainer, the local minima in the hyperparameter space distribute differently. As shown in Fig. 2 also Table 9, the best pre-training configuration for probing is "mask rate=0.85 and learning rate=1e-4", while in terms of fine-tuning scores, the optimal setting is "mask rate=0.50 and learning rate=5e-4". Also, it can be inferred that the optimal pre-training hyperparameters for different pre-training datasets vary; therefore, using reported hyperparameters without carefulness and concluding "graph pretraining is ineffective in molecular domain" is not convincing [49].

5 Molecular graph representation evaluation

The goal of GSSL for molecular graphs is to obtain embeddings that capture generic information about the molecule and its properties. However, there is no free lunch [51]: different training objectives optimise for different properties, and evaluating the extracted embeddings on only a handful of downstream datasets does not provide the whole picture (as we confirm empirically in Sec. 6). Also, from Sec. 4, we know the probing and fine-tuning performance are positively correlated, yet their optimal pre-training configurations are largely diverging. In the Appendix, we also provide results on the worst pre-training configurations, some of which cause negative transfer due to initialising the encoders into local bad minima. Investigations are required to understand what kind of property makes the pre-trained encoders differ.

To this end, we propose MOLGRAPHEVAL, which encompasses a variety of carefully-selected probe tasks, categorised into three classes: (i) **generic graph properties**, (ii) **molecular substructure properties** and (iii) **embedding space properties**. In the upcoming subsections, we explain the tasks in more detail and why they are essential for molecular graph embeddings.

5.1 Generic graph properties

Topological property statistics are often used as features in machine learning pipelines on graphs that do not rely on neural networks [52]. Based on their scale, they can be divided into {node-, pair-, and graph-} level statistics. For molecular graphs, topological metrics have been widely used as molecular descriptors in cheminformatics for decades [53–56], metrics at different scales will facilitate different tasks.

Node-level statistics accompany each node with a local topological measure, which could be used as features in node classification [52]. Concretely, node-level information such as degree [57] can reflect the reaction centres [58]; thus, it can aid in discovering chemical reactions [59].

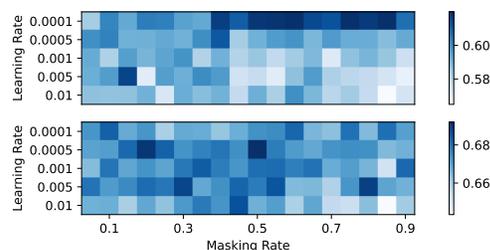


Figure 2: **Hyperparameter space of ATTRMASK** (mask rate \times learning rate), coloured by MPP test scores. Above: fixed; below: fine-tuned. Note the colour bars are different.

- **Node degree** (d_u) counts the number of edges incident to node u : $d_u = \sum_{v \in V} \mathbf{A}[u, v]$.
- **Centrality** (e_u) represents a node's importance. The eigenvector centrality is determined by a relation proportional to the average centrality of its neighbours: $e_u = (\sum_{v \in V} \mathbf{A}[u, v]e_v) / \lambda, \forall u \in V$.
- **Clustering coefficient** (c_u) measures how tightly clustered a node's neighbourhood is: $c_u = (|\{(v_1, v_2) \in \mathcal{E} : v_1, v_2 \in \mathcal{N}(u)\}|) / d_u^2$, i.e., the fraction of closed triangles in neighbourhood [60].

Graph-level statistics summarise global topology information and are helpful for graph classification tasks. For molecules, graph-level statistics can be used, e.g., to classify a molecule's solubility [61]. We briefly describe their intuitions; formal definitions can be found, e.g., in [52].

- **Diameter** is the maximum distance between the pair of vertices (i.e., longest path in a molecule).
- **Cycle basis** is a set of simple cycles that forms a basis of the graph cycle space. It is a minimal set that allows every even-degree subgraph to be expressed as a symmetric difference of basis cycles.
- **Connectivity** is the minimum number of elements (nodes or edges) that need to be removed to separate the remaining nodes into two or more isolated subgraphs.
- **Assortativity** measures the similarity of connections in the graph with respect to the node degree. It can be seen as the Pearson correlation coefficient of degrees between pairs of linked nodes.

Pair-level statistics quantify the relationships between nodes (atoms), which is vital in molecular modelling. For example, molecular docking techniques aim to predict the best matching binding mode of a ligand to a macro-molecular partner [62]. For predicting such binding compatibility, connectivity and distance awareness (how close a pair of atoms can be) are important. In our implementation, we randomly select a fixed number (i.e., 10) of atom pairs from each molecular graph. These pairs are then categorised based on their originating molecule, ensuring all pairs from a single molecule are designated to a singular split: either train, validation, or test.

- **Link prediction** tests whether two nodes are connected or not, given their embeddings and inner products. Based on the principle of *homophily*, it is expected that embeddings of connected nodes are more similar compared to disconnected pairs:

$$\mathbf{S}_{\text{Link}}[u, v, \mathbf{x}_u^T \mathbf{x}_v] = \mathbb{1}_{\mathcal{N}(u)}(v) \quad (2)$$

- **Jaccard coefficient** seeks to quantify the overlap between neighbourhoods while minimising the biases induced by node degrees [63]:

$$\mathbf{S}_{\text{Jaccard}}[u, v] = |\mathcal{N}(u) \cap \mathcal{N}(v)| / |\mathcal{N}(u) \cup \mathcal{N}(v)| \quad (3)$$

- **Katz index** is a global overlap statistic defined by the number of paths between a pair of nodes:

$$\mathbf{S}_{\text{Katz}}[u, v] = \sum_{i=1}^{\infty} \beta^i \mathbf{A}^i[u, v] \quad (4)$$

where $\beta \in \mathbb{R}^+$ determines the weight between short and long paths. $\beta < 1$ reduces the weight of long paths, in implementations we set $\beta = 1$ to give all paths equal importance.

5.2 Molecular substructure properties

Table 2: ROC-AUC scores of classifiers predicting molecular properties.

LINEAR REGRESSION	RANDOM FOREST	XGBOOST	RANDOM (FIX/FT)	JOAO (FIX)	GRAPHCL (FT)
59.91	61.95	62.31	59.60 / 66.16	65.05	70.33

Molecular substructures often serve as reliable indicators of biochemical properties [64–67]. For instance, molecules with benzene rings typically share consistent physical properties, such as solubility, as well as chemical characteristics like aromaticity [68].

Substructures. We investigate 24 substructures from three groups: **rings** (Benzene, Beta lactams, ..., Thiophene); **functional groups** (Amides, Amidine, ..., Urea); and **redox active sites** (Allylic). We provide chemical knowledge on how they relate with molecular properties in Appendix D.

How predictive are substructures? To demonstrate that substructures are quite predictive of molecular properties, we utilise counts of substructures within a molecular graph as the input for

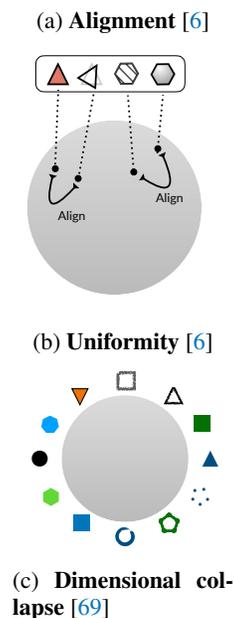
classic ML methods (linear regression, random forest, and XGBoost) to predict the molecular properties on eight downstream datasets. Table 2 shows the results. For ease of comparison, we add the performance of RANDOM, GRAPHCL (FIX), and JOAO (FT) from Table 1. Notably, even basic models trained exclusively on substructures yield performance akin to the randomly trained GNN baseline. This underscores the profound correlation between substructures and MPP task efficacy.

5.3 Embedding space properties

Beyond MPP metrics, we assess domain-agnostic properties of the embedding space produced by pre-trained graph encoders. These properties correlate positively with downstream generalisation [6]. Hence, they can serve as proxies for embedding degeneration, especially in label-scarce scenarios. We explore three such properties in MOLGRAPHEVAL:

- **Alignment** quantifies how similar produced embeddings are for similar samples [6]. Ideally, two samples with the same (or very similar) semantics should be mapped to nearby features, thus mostly invariant to unneeded noise factors. To examine alignment, we construct positive and negative molecule pairs. Positive pairs in a dataset are those that share identical molecular properties, whereas negative pairs are those that differ in their properties. A better alignment represents a better nearest neighbourhood formulation, which has been especially useful for tasks such as compound potency prediction [70].
- **Uniformity** measures how uniformly the embeddings are distributed on the unit hypersphere [6]. A more uniformly distributed embedding space is expected to be with better generalisation under some mild assumptions.
- **Dimensional collapse** refers to the problem of embedding vectors spanning a lower-dimensional subspace instead of the entire available embedding space [69, 71]. Following Jing et al. [69], one simple way to test the occurrence of dimension collapse is to inspect the number of non-zero singular values of a matrix stacking the embedding vectors $\mathbf{Z} = \{\mathbf{z}_i\}_{i=1}^N$.

Figure 3: Embedding space property.



6 Results

6.1 Generic graph properties

For node- and graph-level topological properties, we benchmark GSSL methods on all the molecular graphs from each dataset; for pair-level metrics, we bootstrap 10k node pairs from each dataset for evaluation. The reported test scores are averaged over three runs. We plot the distribution of these metrics in Appendix E.

Findings. Table 3 shows the results, and we summarise these findings.

- 4) **RANDOM outperforms almost every GSSL method on node- and pair-level metrics**, therefore incorporating the randomised features would bring substantial advantages for tasks that are based on local geometry [72]. We observe that among six node- and pair-level graph metrics, randomised embeddings perform the best in four. We further verify this via T-SNE embeddings in Fig. 4, where randomised embeddings can form more interpretable clusters w.r.t. “Atom Nodes”. In Appendix C, we provide a numerical analysis to show that one of the reasons might be the choice of the GNN layer’s initialisation and summation message-passing function. We further empirically experiment with different initialisation strategies, finding that the discriminative power of randomised embeddings on these local metrics disappears.

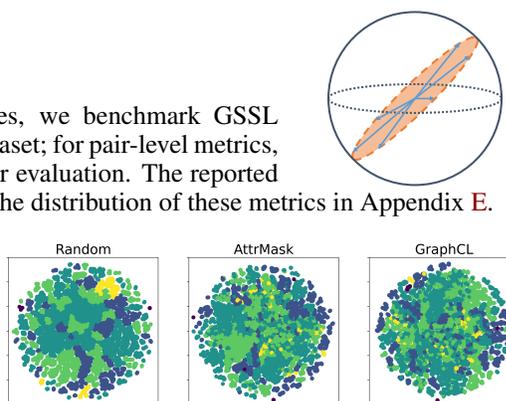


Figure 4: **RANDOM generates more interpretable latent space:** T-SNE visualisation of node embeddings produced by RANDOM, ATTR-MASK, and GRAPHCL on BBBP. Each dot represents an atom from a molecular graph, coloured by node degrees. We notice that RANDOM embeddings form more coherent clusters, consistent with their performance on node-level tasks. While this behaviour is not observed in other atom-level metrics such as node centrality (Fig. 8).

- 5) **RANDOM falls short of predicting graph-level metrics**, which is in alignment with the fact that all GSSL methods outperform RANDOM on the graph-level MPP tasks. However, ranking top on recovering graph topological metrics does not guarantee the embeddings have the best generalisation on downstream MPP tasks.
- 6) **Performance in topological metrics aligns with pre-training objectives.** EDGE PRED and ATTR MASK set the pre-training objectives to predict the adjacency matrix and masked nodes and edges, respectively; the embeddings extracted from these methods compose more information on the node- and pair-level topological metrics. Also, incorporating 3D geometry in the pre-training (*i.e.*, GRAPH MVP) helps retain pair-level topological properties. It aligns with the fact that geometry helps improve target discovery [73–75], as atom pair interactions play a major role.

Table 3: **Benchmarking topological properties.** We report the mean square error or the cross entropy on eight datasets (*i.e.*, smaller is better). For each topological metric, the best and second-best scores are marked **bold** and bold, respectively. We omit such annotations for the “Clustering coefficient (*i.e.*, Cluster)” metric as all test errors are similar. We also report the average ranks of each GSSL method on these metrics (“R”), grouped by the three abstract levels of topology. We use the pre-training configurations that achieve the best MPP test performance with fixed embeddings, as reported in Tables 1 and 9.

	Node			Pair				Graph					
	Degree	Cent.	Cluster	R	Link	Jaccard	Katz	R	Diameter	Conn.	Cycle	Assort.	R
RANDOM	0.001	0.008	0.003	<u>1.5</u>	0.078	0.012	0.017	2.8	177.924	0.087	2.933	0.029	8.6
EDGE PRED	0.031	0.009	0.003	3	0.067	0.014	0.016	2.2	159.825	0.073	2.596	0.026	6.5
ATTR MASK	0.009	0.009	0.003	2.7	0.082	0.015	0.020	4.7	110.793	0.062	2.207	0.019	2
GPT-GNN	0.123	0.009	0.003	4.7	0.014	0.021	0.029	5.3	111.688	0.074	2.854	0.026	6.5
INFOGRAPH	0.054	0.009	0.004	5	0.088	0.019	0.021	6	84.339	0.066	2.100	0.029	6
CONT.PRED	0.164	0.010	0.004	8	0.014	0.021	0.045	5.7	138.304	0.067	2.150	0.027	5.3
GROVER	0.120	0.012	0.004	8.2	0.114	0.047	0.059	10	78.352	0.064	2.058	0.021	3.3
GRAPHCL	0.060	0.010	0.004	6.7	0.084	0.028	0.026	7.3	90.336	0.066	2.287	0.026	6.1
JOAO	0.067	0.010	0.004	7	0.089	0.041	0.025	8	95.335	0.063	2.352	0.024	5.3
GRAPH MVP	0.199	0.010	0.004	8.3	0.077	0.014	0.017	3	109.198	0.065	2.372	0.030	5.5

6.2 Substructure properties

We use Cramér’s V statistics to identify the five substructures mostly associated with downstream biochemical properties, as reported in Table 5, detailed in Table 14. We probe the pre-trained embeddings to predict these substructures. We provide the complete results and plot the distributions of all substructures in Appendix D.

Findings. Table 4 shows the results, where we **bold** the best and underline the worst scores of each substructure. We report the Spearman rank correlation and p-values between the performance of recognising substructures and predicting molecular properties. We highlight the following findings.

- 7) **Substructure detection performance correlates well with MPP performance.** Pre-trained embeddings notably surpass the RANDOM in both substructure detection and Multiple Property Prediction (MPP)

Table 4: **Benchmarking substructure properties.** We report the mean square errors on the test split averaged over the eight downstream MPP datasets.

	allylic	amide	benzene	ether	halogen
RANDOM	0.959	<u>16.917</u>	<u>1.100</u>	<u>2.024</u>	<u>1.127</u>
EDGE PRED	0.780	14.173	0.797	1.608	0.939
ATTR MASK	0.926	14.703	0.976	1.742	0.501
GPT-GNN	0.872	15.629	0.783	1.912	0.341
INFOGRAPH	0.740	6.747	0.583	1.128	0.706
CONT.PRED	<u>1.040</u>	16.636	0.980	1.787	1.075
GROVER	0.715	6.576	0.558	0.957	0.298
GRAPHCL	0.652	7.598	0.525	1.077	0.319
JOAO	0.654	7.926	0.531	1.071	0.310
GRAPH MVP	0.905	6.992	0.649	1.037	0.311
Corr.	0.830	0.770	0.915	0.879	0.806
p-value	3e-3	9e-3	2e-4	8e-4	5e-3

Table 5: **Cramér’s V between molecular substructure counts and binary properties**, averaged over 678 property prediction tasks or eight downstream datasets.

	allylic	benzene	amide	ether	halogen
Task	0.1144	0.1630	0.0881	0.1034	0.1721
Dataset	0.1024	0.1227	0.1336	0.1083	0.1086

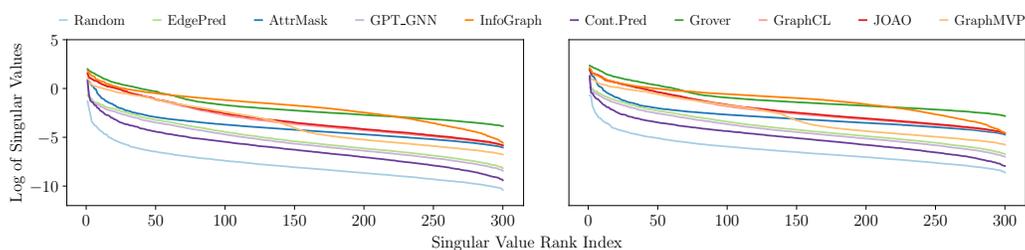


Figure 5: **Dimensional collapse analysis on BBBP**. Left: graph; Right: node.

tasks, with the sole exception of continuous prediction in the “allyli” category. The superior performance of GSSLs in predicting molecular properties could be attributed to their capacity for substructure awareness. This observation is further corroborated by the high positive rank correlation and substantial statistical significance, as indicated by all p-values being less than 1%. This evidence suggests that integrating substructure awareness into GSSL methods could potentially enhance the accuracy of molecular property predictions.

- 8) **Motif-based and Contrastive-based GSSL methods have better substructure awareness.** GROVER, GRAPHCL, JOAO, and JOAOV2 perform consistently better than most other GSSL methods on substructure detection also in the MPP tasks. Note that the optimal pre-training configurations for GROVER is “Motif”-based loss¹, as reported in Table 9.

6.3 Embedding space properties

Findings. Based on the above results, we summarise the following findings. For alignment, we randomly select 10k positive and negative pairs of molecular graphs from BBBP and Tox21 datasets, calculate the cosine similarity and plot the histogram in Fig. 6. We choose ATTRMASK and GRAPHCL to represent generative and contrastive GSSL methods, respectively. Table 16 presents uniformity values as defined in [6]; Fig. 5 plots the magnitude of the singular values in the logarithm scale provided in Table 15.

Table 6: **Benchmarking embedding space properties.** We report the rank correlations and p-values.

	Node Embed	Graph Embed	Uniformity
Correlation	0.806	0.927	0.842
p-value	4e-3	6e-3	2e-3

- 9) **Compared with the RANDOM initialised GNN, GSSL methods give rise to better alignments, promote more uniform features and lift the spectrum.**

GSSL embeddings form distinguishable distributions for positive/negative pairs, while the RANDOM embeddings do not (a phenomenon often referred to as over-smoothing [76], see Fig. 6). All the GSSL methods have better uniformly distributed embeddings on all datasets (in Table 16). However, we found that a better alignment is not necessary to achieve better generalisation for the domain of the molecular graph (Table 16). The singular values of stacked GSSL embeddings (both node and graph) are larger than RANDOM’s by multiple magnitudes; also, we observe that the magnitude of the spectrum positively correlates with the downstream MPP performance (in Tables 6 and 15).

7 Conclusion

In this work, we challenged common practices in evaluating graph self-supervised learned embeddings of molecular graphs. First, we extensively searched the optimal hyperparameters and evaluated GSSL methods on common molecular property prediction tasks in an unbiased and controlled manner. Next, we presented MOLGRAPHEVAL, which is a diverse collection of probe tasks divided into three categories: (i) topological properties, (ii) substructure properties, and (iii) embedding space properties. Then we evaluated GSSL methods on MOLGRAPHEVAL and found surprising insights not revealed by the evaluation of MPP tasks alone.

¹Here the concepts of “Motif” and “Substructure” are identical, under the context of molecular graphs.

The purpose of this work is to complement current evaluation practices with probe tasks and metrics that reveal novel insights, rather than arguing about which combinations of pre-training tasks yield the best downstream performance. Also, as our primary focus is the pre-trained GNN encoders, we leave the investigations of comparing probing and fine-tuning embeddings in the future. Our empirical findings suggest that there are many open questions on how to learn robust molecular graph embeddings without labels and a better understanding of these, along with a new methodology for solving some of the issues mentioned earlier (*e.g.*, dimensional collapse), are yet to come. Nevertheless, we are optimistic that the tasks proposed in this paper will benefit the GSSL research community to tackle these challenges and applied scientists in fields like drug discovery to yield additional insights that can help their problem.

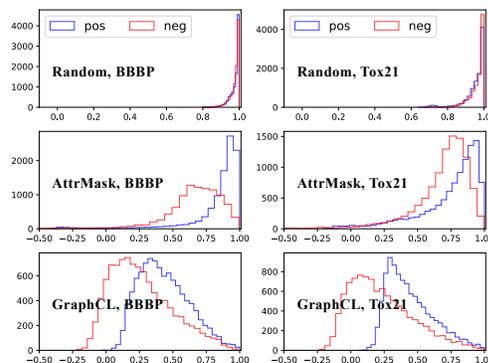


Figure 6: Alignment in the embedding space.

Acknowledge

We thank Le Song, Anima Anandkumar, Matthew Welborn for valuable discussions.

References

- [1] Eugene N Muratov, Jürgen Bajorath, Robert P Sheridan, Igor V Tetko, Dmitry Filimonov, et al. Qsar without borders. *Chemical Society Reviews*, 49(11):3525–3564, 2020. 1
- [2] Hanchen Wang, Tianfan Fu, Yuanqi Du, Wenhao Gao, Kexin Huang, Ziming Liu, Payal Chandak, Shengchao Liu, Peter Van Katwyk, Andreea Deac, et al. Scientific discovery in the age of artificial intelligence. *Nature*, 620(7972):47–60, 2023. 1
- [3] Megan Stanley, John F Bronskill, Krzysztof Maziarz, Hubert Misztela, et al. FS-mol: A few-shot learning dataset of molecules. In *NeurIPS Datasets and Benchmarks Track*, 2021. 1
- [4] RS Bohacek, C McMartin, and WC Guida. The art and practice of structure-based drug design: a molecular modeling perspective. *Medicinal research reviews*, 16(1):3–50, 1996. 1
- [5] Yuyang Wang, Jianren Wang, Zhonglin Cao, and Amir Barati Farimani. Molecular contrastive learning of representations via graph neural networks. *Nature Machine Intelligence*, 2022. 1
- [6] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *ICML*, 2020. 2, 6, 8, 32
- [7] Yixin Liu, Ming Jin, Shirui Pan, Chuan Zhou, Yu Zheng, Feng Xia, and Philip Yu. Graph self-supervised learning: A survey. *IEEE TKDE*, 2022. 2
- [8] Yaochen Xie, Zhao Xu, Jingtun Zhang, Zhengyang Wang, and Shuiwang Ji. Self-supervised learning of graph neural networks: A unified review. *IEEE TPAMI*, 2022.
- [9] Xiao Liu, Fanjin Zhang, Zhenyu Hou, Li Mian, Zhaoyu Wang, et al. Self-supervised learning: Generative or contrastive. *IEEE TKDE*, 2021. 2
- [10] Fan-Yun Sun, Jordan Hoffmann, et al. Infograph: Unsupervised and semi-supervised graph-level representation learning via mutual information maximization. In *ICLR*, 2020. 2
- [11] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, et al. Strategies for pre-training graph neural networks. In *ICLR*, 2020. 2, 3
- [12] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, et al. Graph contrastive learning with augmentations. In *NeurIPS*, 2020. 2, 3
- [13] William Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *NIPS*, 2017. 2
- [14] Shengchao Liu, Mehmet Furkan Demirel, and Yingyu Liang. N-gram graph: Simple unsupervised representation for graphs, with applications to molecules. In *NeurIPS*, 2018. 2
- [15] Ziniu Hu, Yuxiao Dong, Kuansan Wang, Kai-Wei Chang, and Yizhou Sun. Gpt-gnn: Generative pre-training of graph neural networks. In *KDD*, 2020. 2
- [16] Yanqiao Zhu, Yichen Xu, Qiang Liu, and Shu Wu. An empirical study of graph contrastive learning. In *NeurIPS Datasets and Benchmarks Track*, 2021. 2
- [17] Guillaume Alain and Yoshua Bengio. Understanding intermediate layers using linear classifier probes. In *ICLR Workshop*, 2017. 2
- [18] Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. Linguistic knowledge and transferability of contextual representations. In *NAACL*, 2019. 2, 3
- [19] John Hewitt and Christopher D. Manning. A structural probe for finding syntax in word representations. In *NAACL*, 2019.
- [20] Ian Tenney et al. BERT rediscovers the classical NLP pipeline. In *ACL*, 2019.
- [21] Ganesh Jawahar, Benoît Sagot, and Djamel Seddah. What does BERT learn about the structure of language? In *ACL*, 2019.

- [22] Nora Kassner and Hinrich Schütze. Negated and misprimed probes for pretrained language models: Birds can talk, but cannot fly. In *ACL*, 2020.
- [23] Lisa Anne Hendricks, John Mellor, Rosalia Schneider, Jean-Baptiste Alayrac, et al. Decoupling the role of data, attention, and losses in multimodal transformers. *TACL*, 2021. 2
- [24] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, et al. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 2
- [25] Kaiming He, Xinlei Chen, Saining Xie, et al. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. 4
- [26] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *ICCV*, 2021.
- [27] Chunyuan Li, Jianwei Yang, Pengchuan Zhang, Mei Gao, Bin Xiao, et al. Efficient self-supervised vision transformers for representation learning. In *ICLR*, 2022.
- [28] Hanchen Wang, Qi Liu, Xiangyu Yue, Joan Lasenby, and Matthew J. Kusner. Unsupervised point cloud pre-training via occlusion completion. In *ICCV*, 2021. 2
- [29] Shengchao Liu, David Vazquez, Jian Tang, and Pierre-André Noël. Flaky performances when pretraining on relational databases. *arXiv:2211.05213*, 2022. 2
- [30] Roshan Rao, Nicholas Bhattacharya, Neil Thomas, Yan Duan, Xi Chen, et al. Evaluating protein transfer learning with tape. In *NeurIPS*, 2019. 2
- [31] Alexander Rives, Joshua Meier, Tom Sercu, et al. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *PNAS*, 2021.
- [32] Ahmed Elnaggar, Michael Heinzinger, Christian Dallago, Ghalia Rihawi, Yu Wang, et al. Prottrans: towards cracking the language of life’s code through self-supervised deep learning and high performance computing. *IEEE TPAMI*, 2021. 2
- [33] Weihua Hu, Matthias Fey, Hongyu Ren, Maho Nakata, Yuxiao Dong, and Jure Leskovec. Ogb-lsc: A large-scale challenge for machine learning on graphs. In *NeurIPS Dataset and Benchmark Track*, 2021. 2
- [34] Qinkai Zheng, Xu Zou, Yuxiao Dong, Yukuo Cen, Da Yin, Jiarong Xu, Yang Yang, and Jie Tang. Graph robustness benchmark: Benchmarking the adversarial robustness of graph machine learning. In *NeurIPS Dataset and Benchmark Track*, 2021.
- [35] Yuanqi Du, Shiyu Wang, Xiaojie Guo, Hengning Cao, Shujie Hu, Junji Jiang, Aishwarya Varala, Abhinav Angirekula, and Liang Zhao. Graphgt: Machine learning datasets for graph generation and transformation. In *NeurIPS Dataset and Benchmark Track*, 2021.
- [36] Shurui Gui, Xiner Li, Limei Wang, and Shuiwang Ji. Good: A graph out-of-distribution benchmark. In *NeurIPS Dataset and Benchmark Track*, 2022.
- [37] Yijian Qin, Ziwei Zhang, Xin Wang, Zeyang Zhang, and Wenwu Zhu. Nas-bench-graph: Benchmarking graph neural architecture search. In *NeurIPS Dataset and Benchmark Track*, 2022. 2
- [38] David Duvenaud, Dougal Maclaurin, Jorge A.-Iparraguirre, Rafael Gómez-Bombarelli, et al. Convolutional networks on graphs for learning molecular fingerprints. In *NIPS*, 2015. 2
- [39] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *ICML*, 2017. 2
- [40] Kevin Yang, Kyle Swanson, Wengong Jin, Connor Coley, et al. Analyzing learned molecular representations for property prediction. *Journal of Chemical Information and Modeling*, 2019.
- [41] Gabriele Corso, Luca Cavalleri, Dominique Beaini, Pietro Liò, and Petar Veličković. Principal neighbourhood aggregation for graph nets. In *NeurIPS*, 2020. 2

- [42] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *ICLR*, 2019. 2
- [43] Yu Rong, Yatao Bian, Tingyang Xu, Weiyang Xie, Ying Wei, et al. Self-supervised graph transformer on large-scale molecular data. In *NeurIPS*, 2020. 2
- [44] Yuning You, Tianlong Chen, Yang Shen, and Zhangyang Wang. Graph contrastive learning automated. In *ICML*, 2021. 2
- [45] Shengchao Liu, Hanchen Wang, Weiyang Liu, Joan Lasenby, Hongyu Guo, et al. Pre-training molecular graph representation with 3d geometry. In *ICLR*, 2022. 2
- [46] Simon Axelrod and Rafael Gomez-B. Geom: Energy-annotated molecular conformations for property prediction and molecular generation. *Scientific Data*, 2022. 2, 17
- [47] Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, et al. Moleculenet: a benchmark for molecular machine learning. *Chemical Science*, 2018. 3
- [48] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, et al. Open graph benchmark: Datasets for machine learning on graphs. In *NeurIPS*, 2021. 3
- [49] Ruoxi Sun, Hanjun Dai, and Adams Wei Yu. Does GNN pretraining help molecular representation? In *NeurIPS*, 2022. 3, 4, 17
- [50] Martin Vogt, Dagmar Stumpfe, Hanna Geppert, and Jürgen Bajorath. Scaffold hopping using two-dimensional fingerprints: True potential, black magic, or a hopeless endeavor? guidelines for virtual screening. *Journal of Medicinal Chemistry*, 53(15):5707–5715, 2010. 4
- [51] David H Wolpert and William G Macready. No free lunch theorems for optimization. *IEEE transactions on evolutionary computation*, 1(1):67–82, 1997. 4
- [52] William Hamilton. *Graph Representation Learning*. Morgan & Claypool Publishers, 2020. 4, 5
- [53] James Devillers and Alexandru T Balaban. *Topological Indices and Related Descriptors in QSAR and QSPAR*. CRC Press, 2000. 4
- [54] Dejun Jiang, Zhenxing Wu, Chang-Yu Hsieh, Guangyong Chen, Ben Liao, et al. Could graph neural networks learn better molecular representation for drug discovery? a comparison study of descriptor-based and graph-based models. *Journal of Cheminformatics*, 2021.
- [55] Mati Karelson. *Molecular descriptors in QSAR/QSPR*. Wiley-Interscience, 2000.
- [56] Frank Emmert-Streib. *Statistical modelling of molecular descriptors in QSAR/QSPR*. John Wiley & Sons, 2012. 4
- [57] Tomislav Došlic, Boris Furtula, Ante Graovac, Ivan Gutman, Sirous Moradi, and Zahra Yarahmadi. On vertex-degree-based molecular structure descriptors. *Communications in Mathematical and in Computer Chemistry*, 66(2):613–626, 2011. 4
- [58] Ramil I Nugmanov, Ravil N Mukhametgaleev, Tagir Akhmetshin, Timur R Gimadiev, Valentina A Afonina, et al. Cgrtools: Python library for molecule, reaction, and condensed graph of reaction processing. *Journal of Chemical Information and Modeling*, 2019. 4
- [59] William Bort, Igor Baskin, Timur Gimadiev, Artem Mukanov, et al. Discovery of novel chemical reactions by deep generative recurrent neural network. *Scientific Reports*, 2021. 4
- [60] Duncan J Watts and Steven H Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 1998. 5
- [61] Benjamin Sanchez-Lengeling, Loïc M Roch, José Darío Perea, Stefan Langner, Christoph J Brabec, et al. A bayesian approach to predict solubility parameters. *Advanced Theory and Simulations*, 2(1):1800069, 2019. 5

- [62] Veronica Salmaso and Stefano Moro. Bridging molecular docking to molecular dynamics in exploring ligand-protein recognition process: An overview. *Frontiers in pharmacology*, 2018. 5
- [63] Linyuan Lü and Tao Zhou. Link prediction in complex networks: A survey. *Physica A: statistical mechanics and its applications*, 2011. 5
- [64] Gerta Rücker and Christoph Rücker. Substructure, subgraph, and walk counts as measures of the complexity of graphs and molecules. *Journal of Chemical Information and Computer Sciences*, 41(6):1457–1462, 2001. 5
- [65] Youngchun Kwon, Dongseon Lee, Youn-Suk Choi, Kyoham Shin, and Seokho Kang. Compressed graph representation for scalable molecular graph generation. *Journal of Cheminformatics*, 2020.
- [66] Ryuichiro Hataya, Hideki Nakayama, and Kazuki Yoshizoe. Graph energy-based model for substructure preserving molecular design. *arXiv:2102.04600*, 2021.
- [67] Xian-bin Ye, Quanlong Guan, Weiqi Luo, Liangda Fang, Zhao-Rong Lai, et al. Molecular substructure graph attention network for molecular property identification in drug discovery. *Pattern Recognition*, 128:108659, 2022. 5
- [68] John E McMurry. *Organic chemistry with biological applications*. Cengage Learning, 2014. 5
- [69] Li Jing, Pascal Vincent, Yann LeCun, and Yuandong Tian. Understanding dimensional collapse in contrastive self-supervised learning. In *ICLR*, 2022. 6
- [70] Tiago Janela and Jürgen Bajorath. Simple nearest-neighbour analysis meets the accuracy of compound potency predictions using complex machine learning models. *Nature Machine Intelligence*, pages 1–10, 2022. 6
- [71] Tianyu Hua, Wenxiao Wang, Zihui Xue, Sucheng Ren, Yue Wang, et al. On feature decorrelation in self-supervised learning. In *ICCV*, 2021. 6
- [72] Tianjin Huang, Tianlong Chen, Meng Fang, Vlado Menkovski, Jiaxu Zhao, Lu Yin, Yulong Pei, Decebal Constantin Mocanu, Zhangyang Wang, Mykola Pechenizkiy, et al. You can have better graph neural networks by not training weights at all: Finding untrained graph tickets. In *Learning on Graphs Conference*, 2021. 6
- [73] Hannes Stärk, Dominique Beaini, Gabriele Corso, Prudencio Tossou, Christian Dallago, Stephan Günnemann, and Pietro Liò. 3d infomax improves gns for molecular property prediction. In *ICML*, 2022. 7
- [74] Hannes Stärk, Octavian Ganea, Lagnajit Pattanaik, Regina Barzilay, and Tommi Jaakkola. Equibind: Geometric deep learning for drug binding structure prediction. In *ICML*, 2022.
- [75] Xingang Peng, Shitong Luo, Jiaqi Guan, Qi Xie, Jian Peng, and Jianzhu Ma. Pocket2mol: Efficient molecular sampling based on 3d protein pockets. In *ICML*, 2022. 7
- [76] Deli Chen, Yankai Lin, Wei Li, Peng Li, et al. Measuring and relieving the over-smoothing problem for graph neural networks from the topological view. In *AAAI*, 2020. 8
- [77] Thomas A Halgren. Merck molecular force field. i. basis, form, scope, parameterization, and performance of mmff94. *Journal of computational chemistry*, 17(5-6):490–519, 1996. 17
- [78] Xiaomin Fang, Lihang Liu, Jieqiong Lei, Donglong He, Shanzhuo Zhang, Jingbo Zhou, Fan Wang, Hua Wu, and Haifeng Wang. Geometry-enhanced molecular representation learning for property prediction. *Nature Machine Intelligence*, 4(2):127–134, 2022. 17
- [79] AG Donchev, VD Ozrin, MV Subbotin, OV Tarasov, and VI Tarasov. A quantum mechanical polarizable force field for biomolecular interactions. *PNAS*, 2005. 17

- [80] Michael D Beachy, David Chasman, Robert B Murphy, Thomas A Halgren, and Richard A Friesner. Accurate ab initio quantum chemical determination of the relative energetics of peptide conformations and assessment of empirical force fields. *JACS*, 1997.
- [81] Ilana Y Kanal, John A Keith, and Geoffrey R Hutchison. A sobering assessment of small-molecule force field methods for low energy conformer predictions. *International Journal of Quantum Chemistry*, 118(5):e25512, 2018.
- [82] Rui Jiao, Jiaqi Han, Wenbing Huang, Yu Rong, and Yang Liu. 3d equivariant molecular graph pretraining. *arXiv:2207.08824*, 2022.
- [83] Gengmo Zhou, Zhifeng Gao, Qiankun Ding, Hang Zheng, et al. Uni-mol: A universal 3d molecular representation learning framework. *chemRxiv*, 2022. 17
- [84] Shengchao Liu, Hongyu Guo, and Jian Tang. Molecular geometry pretraining with se (3)-invariant denoising distance matching. *arXiv:2206.13602*, 2022. 17
- [85] Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole Von Lilienfeld. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, 2014. 17
- [86] Raphael JL Townshend, Martin Vögele, Patricia Suriana, Alexander Derry, et al. Atom3d: Tasks on molecules in three dimensions. In *NeurIPS Dataset and Benchmark Track*, 2021. 17
- [87] Yixin Liu, Ming Jin, Shirui Pan, Chuan Zhou, Yu Zheng, Feng Xia, and S Yu Philip. Graph self-supervised learning: A survey. *IEEE TKDE*, 2022. 17
- [88] Xiao Luo, Wei Ju, Meng Qu, Yiyang Gu, Chong Chen, Minghua Deng, Xian-Sheng Hua, and Ming Zhang. Clear: Cluster-enhanced contrast for self-supervised graph representation learning. *IEEE TNNLS*, 2022. 18
- [89] Shuai Lin, Chen Liu, Pan Zhou, Zi-Yuan Hu, Shuojia Wang, Ruihui Zhao, Yefeng Zheng, Liang Lin, Eric Xing, and Xiaodan Liang. Prototypical graph contrastive learning. *IEEE TNNLS*, 2022. 18
- [90] Haifeng Li, Jun Cao, Jiawei Zhu, Qinyao Luo, Silu He, and Xuying Wang. Augmentation-free graph contrastive learning of invariant-discriminative representations. *IEEE TNNLS*, 2023. 18
- [91] Jintang Li, Ruofan Wu, Wangbin Sun, Liang Chen, Sheng Tian, Liang Zhu, Changhua Meng, Zibin Zheng, and Weiqiang Wang. Maskgae: Masked graph modeling meets graph autoencoders. *arXiv:2205.10053*, 2022. 18
- [92] Qiaoyu Tan, Ninghao Liu, Xiao Huang, Rui Chen, Soo-Hyun Choi, and Xia Hu. Mgae: Masked autoencoders for self-supervised learning on graphs. *arXiv:2201.02534*, 2022.
- [93] Zhenyu Hou, Xiao Liu, Yukuo Cen, Yuxiao Dong, Hongxia Yang, Chunjie Wang, and Jie Tang. Graphmae: Self-supervised masked graph autoencoders. In *KDD*, 2022. 18
- [94] Lirong Wu, Haitao Lin, Cheng Tan, Zhangyang Gao, and Stan Z Li. Self-supervised learning on graphs: Contrastive, generative, or predictive. *IEEE TKDE*, 2021. 18
- [95] Mohammad Sadegh Akhondzadeh, Vijay Lingam, and Aleksandar Bojchevski. Probing graph representations. In *AISTATS*, 2023. 18
- [96] Shengchao Liu, Weitao Du, Yanjing Li, Zhuoxinran Li, Zhiling Zheng, Chenru Duan, Zhiming Ma, Omar Yaghi, Anima Anandkumar, Christian Borgs, et al. Symmetry-informed geometric representation for molecules, proteins, and crystalline materials. *arXiv:2306.09375*, 2023. 18
- [97] Zhengdao Chen, Lei Chen, Soledad Villar, and Joan Bruna. Can graph neural networks count substructures? In *NeurIPS*, 2020. 21
- [98] Markus Zopf. 1-wl expressiveness is (almost) all you need. *arXiv:2202.10156*, 2022. 21

- [99] Alan D McNaught, Andrew Wilkinson, et al. *Compendium of chemical terminology*, volume 1669. Blackwell Science Oxford, 1997. 21
- [100] Evan J Horn, Brandon R Rosen, Yong Chen, Jiaze Tang, Ke Chen, Martin D Eastgate, and Phil S Baran. Scalable and sustainable electrochemical allylic c–h oxidation. *Nature*, 533(7601):77–81, 2016. 21
- [101] Akihiko Nakamura and Masahisa Nakada. Allylic oxidations in natural product synthesis. *Synthesis*, 45(11):1421–1451, 2013.
- [102] Liela Bayeh, Phong Q Le, and Uttam K Tambar. Catalytic allylic oxidation of internal alkenes to a multifunctional chiral building block. *Nature*, 547(7662):196–200, 2017. 21
- [103] Yousaf Ali, Shafida A Hamid, and Umer Rashid. Biomedical applications of aromatic azo compounds. *Mini reviews in medicinal chemistry*, 18(18):1548–1558, 2018. 22
- [104] Esteban Lanzarotti, Lucas A Defelipe, Marcelo A Marti, and Adrián G Turjanski. Aromatic clusters in protein–protein and protein–drug complexes. *Journal of cheminformatics*, 12(1):1–9, 2020. 22
- [105] Ana R Gomes, Carla L Varela, Elisiário J Tavares-da Silva, and Fernanda MF Roleira. Epoxide containing molecules: A good or a bad drug design approach. *European Journal of Medicinal Chemistry*, 201:112327, 2020. 22
- [106] Min Tian, Ying Peng, and Jiang Zheng. Metabolic activation and hepatotoxicity of furan-containing compounds. *Drug Metabolism and Disposition*, 50(5):655–670, 2022. 22
- [107] Jeffrey B Sperry and Dennis L Wright. Furans, thiophenes and related heterocycles in drug discovery. *Current opinion in drug discovery & development*, 8(6):723–740, 2005. 22
- [108] Franciszek Saczewski and Łukasz Balewski. Biological activities of guanidine compounds. *Expert opinion on therapeutic patents*, 19(10):1417–1448, 2009. 22
- [109] Marcelo Z Hernandez, Suellen Melo T Cavalcanti, Diogo Rodrigo M Moreira, Walter Filgueira de Azevedo Junior, and Ana Cristina Lima Leite. Halogen atoms in the modern medicinal chemistry: hints for the drug design. *Current drug targets*, 11(3):303–314, 2010. 22
- [110] Angeliki P Kourounakis, Dimitrios Xanthopoulos, and Ariadni Tzara. Morpholine as a privileged structure: a review on the medicinal chemistry and pharmacological activity of morpholine containing bioactive molecules. *Medicinal Research Reviews*, 40(2):709–752, 2020. 23
- [111] Tao Wang, Philipp M Stein, Hongwei Shi, Chao Hu, Matthias Rudolph, and A Stephen K Hashmi. Hydroxylamine-mediated c–c amination via an aza-hock rearrangement. *Nature Communications*, 12(1):1–11, 2021. 23
- [112] Saloni Kakkar and Balasubramanian Narasimhan. A comprehensive review on biological activities of oxazole derivatives. *BMC chemistry*, 13(1):1–24, 2019. 23
- [113] Yoshio Hamada. *Role of pyridines in medicinal chemistry and design of BACE1 inhibitors possessing a pyridine scaffold*. InTech Rijeka, 2018. 23
- [114] Yong Ling, Zhi-You Hao, Dong Liang, Chun-Lei Zhang, Yan-Fei Liu, and Yan Wang. The expanding role of pyridine and dihydropyridine scaffolds in drug design. *Drug Design, Development and Therapy*, 15:4289, 2021. 23
- [115] Constantinos G Neochoritis, Ting Zhao, and Alexander Domling. Tetrazoles via multicomponent reactions. *Chemical reviews*, 119(3):1970–2042, 2019. 24
- [116] Mahesh T Chhabria, Shivani Patel, Palmi Modi, and Pathik S Brahmshatriya. Thiazole: A review on chemistry, synthesis and therapeutic importance of its derivatives. *Current topics in medicinal chemistry*, 16(26):2841–2862, 2016. 24
- [117] Rashmi Shah and Prabhakar Kumar Verma. Therapeutic importance of synthetic thiophene. *Chemistry Central Journal*, 12(1):1–22, 2018. 24

Appendix

Table of Contents

A Overview	17
B Pre-Training	17
B.1 GEOM dataset	17
B.2 GSSL methods	17
B.3 Hyperparameters search	18
B.4 Probe models	19
B.5 Computation efficiency	19
B.6 Practical guides for future research	20
C Randomised embeddings	20
D Substructure	21
D.1 Discussions on substructure counting	21
D.2 Detailed description and performance	21
D.3 Cramer's V	25
D.4 Distribution	25
E More results on metrics	28
E.1 Topological metric	28
E.2 Spectrum	30
E.3 Uniformity	32
F Extending to more GSSL methods and datasets	32
G Potential ethical and social implications	32

A Overview

Automated MOLGRAPHEVAL pipeline. In the codebase of MOLGRAPHEVAL, we provide setup scripts (in “env/”) for both the docker and conda virtual environment. The end-to-end benchmarking pipeline consists of three modules (in Fig. 7):

- Pre-training the GNN models (other GNN model classes such as MPNN, GCN, GAT are implemented in the codebase);
- Extracting the node/graph/pair-level embeddings from the pre-trained or the randomly initialised GNN models;
- Probing the quality of embeddings with the proposed metrics.

We have meticulously packaged the components of MOLGRAPHEVAL for ease of access and potential extension. The pre-trained methods are housed in “src/pretrainers”, model libraries in “src/models”, pre-training and downstream datasets in “src/datasets”, and probing metrics in “src/validation”. This modular design allows flexibility and extensibility to incorporate new model architectures and datasets. All configurations, specified in YAML files and processed by argument parsers in “src/config”, are managed by scripts (with templates provided in “script”). Furthermore, we have implemented loggers to chronicle training/validation/testing curves during both pre-training and probing. For added convenience, templates for visualising embeddings and analysing datasets are also available.

We open-sourced MOLGRAPHEVAL in <https://github.com/hansen7/MolGraphEval>.

B Pre-Training

We next describe the additional details of the pre-training used in the MOLGRAPHEVAL benchmark.

B.1 GEOM dataset

We avoid using the updated version of GEOM (‘New drug-like molecules’ and ‘MoleculeNet’) to remove the overlap between pre-training and downstream datasets. Compared with other molecular datasets that contain 3D conformation structures, GEOM [46] has the following advantages:

- **Preciseness.** Compared with toolkits like RDKit or MMFF [77] (used in studies such as ChemRL-GEM [78]), DFT-based calculations (used in GEOM) will provide more precise computation results on the 3D molecular conformation structures [79–83]. Such errors in the molecular geometries have been proven harmful to property predictions (Appendix B of [84])
- **Comprehensive.** GEOM provides a more comprehensive collection in comparison with other quantum chemistry-based datasets (*e.g.*, QM9 [85], Atom3D [86]) in terms of quantity and diversity. In comparison with ZINC15 and SAVI used in this concurrent study [49], GEOM provides accurate 3D conformation structures of molecules, which allows to compare more GSSL methods such as GraphMVP.

B.2 GSSL methods

Self-Supervised Learning (SSL) generally bifurcates into contrastive and generative methodologies, each characterised by its unique supervised signal as highlighted by [87]. Contrastive SSL operates by contrasting representations at the inter-data level, while generative SSL emphasises intra-data level reconstruction. Both strategies have been the subject of comprehensive research.

Contrastive GSSL creates multi-layered views of each graph, each capturing different granularity levels, from nodes and subgraphs to the complete graph. It aligns representations for views originating from the same data and distinguishes those from unrelated datasets, targeting a unified

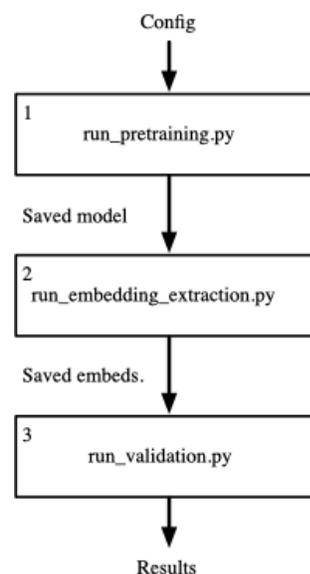


Figure 7: MOLGRAPHEVAL.

embedding space. The effectiveness of various approaches largely hinges on their view design. For example, INFOGRAPH contrasts node and graph views, while GRAPHCL and JOAO delve into graph-level transformations. Further avenues to improve Contrastive GSSL include:

- CLEAR [88] captures graph structure at both global and local scopes, enhancing semantic information granularity and consistency between multiple views;
- PGCL [89] addresses sampling bias by clustering graphs into groups represented by prototype vectors, focusing on the dataset’s global semantics;
- iGCL [90] employs a Siamese architecture to generate positive samples sans data augmentation. The proposed ID loss eschews negative sampling while promoting feature-wise discriminability.

Generative GSSL zeroes in on reconstructing the graph structure, striving to derive representations that capture the core characteristics of the data. Noteworthy examples in this category include EDGE PRED and ATTR MASK, which predict adjacency matrices and mask tokens, respectively. Meanwhile, GPT-GNN employs an auto-regressive approach tailored for holistic graph reconstruction. In line with this methodology, masked graph autoencoders, as cited in [91–93], have garnered considerable attention.

GROVER-Motif leverages domain-specific knowledge to extract motifs from molecules, assigning SSL the role of predicting motif presence. Diverging from the paradigms of contrastive and generative GSSL, recent explorations like [94] categorise this approach as predictive GSSL. In this framework, the supervisory signals are derived from self-generated labels.

There is a limited body of work dedicated to understanding GSSL methods. In a notable study, Akhondzadeh et al. [95] explored the use of probing tasks to measure and contrast the richness of graph representations derived from various models. A significant revelation from this study is that transformer-based GNNs capture chemically pertinent information more effectively than message-passing GNNs. Further integrating the data from 3D structures, Liu et al. [96] presented Geom3D — a comprehensive framework for benchmarking geometric representation learning techniques applicable to molecules, proteins, and materials. This framework encompasses 16 cutting-edge geometric models and evaluates their efficacy across 46 diverse scientific challenges, spanning small molecules, proteins, and crystalline substances. An innovative aspect of Geom3D is its approach to categorising geometric models into three groups: invariant, spherical frame equivariant, and vector frame equivariant.

Table 7: Hyperparameters search space of GSSL .

METHOD	HYPERPARAMETERS	# MODELS
EDGE PRED	LEARNING RATE	15
ATTR MASK	MASK RATE, LEARNING RATE	300
GPT-GNN	LEARNING RATE	15
INFOGRAPH	LEARNING RATE	15
GROVER	LEARNING RATE, “CONTEXTURAL” OR “MOTIF”-BASED LOSS	30
CONT.PRED	LEARNING RATE, CONTEXT SIZE, # NEGATIVE SAMPLES	300
GRAPHCL	LEARNING RATE, AUG STRENGTH, AUG PROB	360
JOAO	LEARNING RATE, GAMMA, LOSS VERSION (V1 OR V2)	300
GRAPHMVP	LEARNING RATE, TEMPERATURE, ALPHA2, # CONFORMER	540
Total		1875

B.3 Hyperparameters search

We search the optimal hyperparameters of pre-training methods, details are summarised in Tables 7 to 9. We select the best hyperparameter of each GSSL method based on their averaged score on downstream datasets (in Table 1, linear models).

We provide details in calculating the **1875** GNN configurations in Table 7. The number of probe models is calculated as follows: 1875 * 8 (MPP datasets) * 2 (Fix, FT) * 3 (Seed) + 9 (GSSL, Optimal) * 10 (Topological Metrics) * 3 (Seed) + 9 (GSSL, Optimal) * 24 (Substructure) * 3 (Seed) = **90918**. It takes over 4 terabytes to save these pre-trained models.

B.4 Probe models

Ideally, probe models should be neither too simple to capture the representation’s information, nor too powerful to learn precise property prediction themselves. If overly powerful, the probe’s performance might not accurately reflect the information embedded in the representations. **In light of these complexities, we select a linear model as our probe, aligning with the choice prevalent in most probe studies.**

Table 8: Range of Grid Search on Hyperparameters Space.

HYPERPARAMETERS	RANGE
LEARNING RATE, ALL BUT GRAPHMVP	[0.01, 0.005, 0.001, 0.0005, 0.0001]
LEARNING RATE, GRAPHMVP	[0.001, 0.0005, 0.0001]
MASK RATE	[0.05, 0.10, ..., 0.95]
CONTEXT SIZE	[2, 3, 4, 5]
# NEGATIVE SAMPLES	[1, 2, 3, 4, 5]
AUG STRENGTH	[0.2, 0.4, 0.6, 0.8]
AUG PROBABILITY	[0.1, 0.2, ..., 1.0]
GAMMA	[0.1, 0.2, ..., 1.0]
TEMPERATURE	[0.1, 0.2, 0.5, 1, 2]
ALPHA2	[0.1, 1, 10]
# CONFORMER	[1, 5, 10, 20]

Table 9: Optimal hyperparameters based on **linear probing and finetuning scores** on MPP tasks.

METHOD	OPTIMAL HYPERPARAMETERS (left: probing / right: fine-tuning)
EDGE PRED	LEARNING RATE=1e-2/1e-2
ATTR MASK	LEARNING RATE=1e-4/5e-4, MASK RATE=0.85/0.50
GPT-GNN	LEARNING RATE=1e-2/1e-4
INFOGRAPH	LEARNING RATE=1e-4/1e-4
GROVER	LEARNING RATE=1e-4/1e-3, “MOTIF”/“CONTEXTUAL”-BASED LOSS
CONT. PRED	LEARNING RATE=1e-3/5e-3, CONTEXT SIZE=1/1, # NEGATIVE SAMPLES=5/1
GRAPHCL	LEARNING RATE=1e-3/1e-3, AUG STRENGTH=0.2/0.6, AUG PROB=0.8/0.5,
JOAO	LEARNING RATE=1e-3/1e-3, GAMMA=0.9/0.6, “V1”/“V1”-VERSION LOSS
GRAPHMVP	LEARNING RATE=5e-4/5e-4, ALPHA2=0.1/10.0, TEMPERATURE=0.1/0.2, # CONFORMER=5/5

B.5 Computation efficiency

We present the number of trainable parameters alongside the average training time per epoch (utilising a single A100 GPU) for each GSSL method in Table 10.

Table 10: Computational efficiency of GSSL methods

METHOD	NUMBER OF PARAMETERS (Million)	TRAINING TIME (Second)
EDGE PRED	7.462	101
ATTRMASK	7.606	38
GPT-GNN	7.606	972
INFOGRAPH	7.823	40
GROVER	7.566	39
CONT.PRED	12.00	202
GRAPHCL	8.186	65
JOAO	8.186	382
GRAPHMVP	15.84	119

B.6 Practical guides for future research

In summary, for prospective advancements in Graph SSL, the following practical guidelines should be considered:

- Develop novel pretext tasks that emphasise beneficial invariances and geometry, as unveiled by MOLGRAPHEVAL probes. This includes tasks that enhance substructure modelling and preserve local topology.
- Relying solely on either probing or fine-tuning may not provide a comprehensive understanding. It's noteworthy that weak probing performance doesn't necessarily correlate with subpar fine-tuning outcomes (*e.g.*, GRAPHMVP). The method of choice should be based on the specific downstream task, whether it's property prediction, generation, optimisation, or interaction modelling.
- Innovate superior data augmentation techniques specifically for molecular graphs to produce valuable views for contrastive learning. Probes can serve as an instrumental means to assess the quality of these augmentations.
- The prevailing notion of achieving a more uniform embedding space is helpful doesn't always hold true in the context of molecular graphs.
- Additionally, probes can be harnessed as a potent tool to scrutinise the impacts of varied negative sampling and augmentation strategies, exemplified by the comparison between GRAPHCL and JOAO.

C Randomised embeddings

How GNN models are initialised. We first analyse how weights in the GNNs are initialised (PyTorch and PyG).

- **Edge Embedding Layers** uses 'xavier uniform', essentially samples from uniform distribution
- **GNN Layers** in fact only have MLP weights (see [PyG Doc](#)), same initialisation as Linear layers.
- **Linear Layers** samples from uniform distribution for both weight and bias ([PyTorch Doc](#))

Since all the weights (Edge embedding layers, GNN layers, and Linear layers) in the GINs are extracted from some uniform distribution of some positive ranges. As the GIN layer essentially consists of multiplications and additions, the expected statistics of the node embeddings from randomised GINs are proportional to the number of connected neighbours (*i.e.*, node degrees). Therefore, the randomised embeddings form discriminative clusters in Fig. 4. As for other node-level metrics (in Fig. 8), we don't observe a good clustering formed from randomised embeddings.

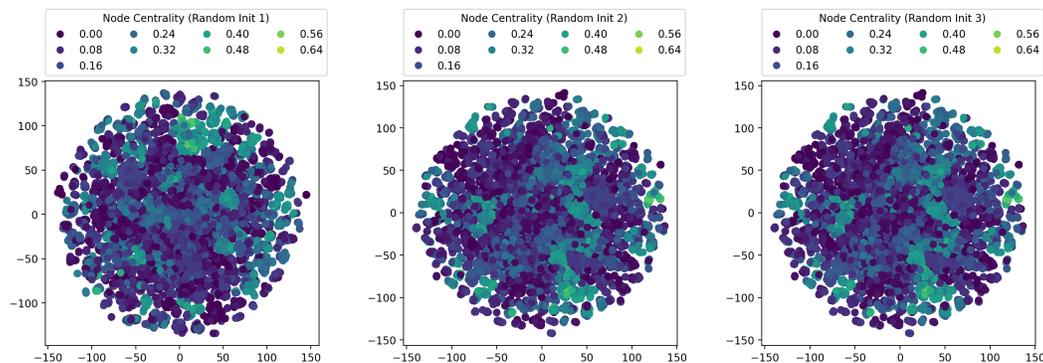


Figure 8: T-SNE visualisation of node centrality produced by RANDOM (three different seeds).

D Substructure

D.1 Discussions on substructure counting

A recent study [97] asserts that message-passing neural networks (MPNNs), including Graph Isomorphism Networks (GIN), struggle with the exact counting of certain induced-subgraph structures. While this observation does not directly indicate that Graph Self-Supervised Learning (GSSL) aids in precise subgraph counting, we find that the knowledge acquired from GSSL pre-training can be construed as substructure awareness.

Conversely, another study [98] demonstrates that in the context of molecular graphs, the theoretical expressiveness limitations of MPNNs, as described by the Weisfeiler-Lehman Isomorphism test, do not impair generalisation performance in real-world datasets. As evidenced in Table 6, GIN models, especially those that are pre-trained, exhibit substantial proficiency in identifying substructures.

In conclusion, the question of whether the expressiveness of the backbone model is a limiting factor in molecular domains remains a topic of ongoing debate. The investigation of the backbone model operates independently of the GSSL pretraining analysis undertaken in our study. As such, we propose deferring this line of inquiry for future exploration.

D.2 Detailed description and performance

In Table 11, we provide the descriptions of the molecular substructures (mainly from documents on the [rdkit.Chem.Fragments](#), textbooks [99] and Wikipedia). We also listed some molecular properties that are affected by these substructures. Table 12 and Table 13 report the detailed performance of substructure property prediction.

Table 11: Descriptions of substructure and which properties their existence would affect. Molecules containing X substructure are usually named as “X-Compounds”, “X derivatives” and “X moieties”.

Substructure	Description & Affected molecular properties
allylic	Allylic oxidations have featured in hundreds of chemical syntheses, due to their particular electrochemical properties [100–102].
amide	An amidine is a compound with the general formula $RC(=O)NR'R''$, where R, R', and R'' represent organic groups or hydrogen atoms. It has significant impacts on the mechanical, acid-base, and solubility properties of molecules [wikipage].

amidine	<p>Amidines are organic compounds with the functional group $RC(NR)NR_2$, where the R groups can be the same or different. They are the imine derivatives of amides ($RC(O)NR_2$). Amidines are much more basic than amides and are among the strongest uncharged/unionized bases.</p> <p>Several drug or drug candidates feature amidine substituents. Examples include the antiprotozoal Imidocarb, the insecticide amitraz, the anthelmintic tribendimidine, and xylamidine, an antagonist at the 5HT_{2A} receptor [wiki page].</p>
Azo	<p>Azo compounds are compounds bearing the functional group diazenyl $R-N=N-R'$, in which R and R' can be either aryl or alkyl. Certain azo compounds are known to have antibiotic, antiviral, antifungal, antineoplastic, and cytotoxic properties [103].</p>
benzene	<p>Benzene (aromatic rings) is an organic chemical compound with the molecular formula C_6H_6. Aromatic rings are important residues for biological interactions and appear to a large extent as part of protein–drug and protein–protein interactions. They are relevant for both protein stability and molecular recognition processes due to their natural occurrence in aromatic aminoacids (Trp, Phe, Tyr and His) as well as in designed drugs since they are believed to contribute to optimising both affinity and specificity of drug-like molecules [104].</p>
epoxide	<p>An epoxide is a cyclic ether with a three-atom ring. Epoxide-containing molecules have therapeutic value. The main therapeutic interest is as anticancer agents. The main mechanisms are enzyme inhibition, induction of cell cycle arrest, apoptosis. Other therapeutic interests are for heart failure, infections, gastrointestinal diseases [105].</p>
ether	<p>Ethers are a class of organic compounds that contain an ether group—an oxygen atom connected to two alkyl or aryl groups. They have the general formula $R-O-R'$, where R and R' represent the alkyl or aryl groups. The C-O bonds that comprise simple ethers are strong. They are unreactive toward all but the strongest bases. Although generally of low chemical reactivity, they are more reactive than alkanes. Some important reactions include cleavage, peroxide formation, Lewis bases, and alpha-halogenation [wiki page].</p>
furan	<p>Furan is a heterocyclic organic compound, consisting of a five-membered aromatic ring with four carbon atoms and one oxygen atom. The furan ring present in the chemical structures may be one of the domineering factors to bring about the toxic response resulting from the generation of reactive epoxide or cis-enedial intermediates which are of the potential to react with biomacromolecules [106, 107].</p>
guanido	<p>Guanidine is the compound with the formula $HNC(NH_2)_2$. It is a colourless solid that dissolves in polar solvents. It is a strong base that is used in the production of plastics and explosives. Most guanidine derivatives are in fact salts containing the conjugate acid [wiki page].</p> <p>Guanidine-containing derivatives constitute a very important class of therapeutic agents suitable for the treatment of a wide spectrum of diseases [108].</p>
halogen	<p>The halogens are a group in the periodic table consisting of five or six chemically related elements: fluorine (F), chlorine (Cl), bromine (Br), iodine (I), and astatine (At). Halogens are highly reactive, and as such can be harmful or lethal to biological organisms in sufficient quantities. This high reactivity is due to the high electronegativity of the atoms due to their high effective nuclear charge [wiki page].</p> <p>A significant number of drugs and drug candidates in clinical development are halogenated structures. For a long time, insertion of halogen atoms on hit or lead compounds was predominantly performed to exploit their steric effects, through the ability of these bulk atoms to occupy the binding site of molecular targets [109].</p>

imidazole	Imidazole is an organic compound with the formula C ₃ N ₂ H ₄ . It is a white or colourless solid that is soluble in water, producing a mildly alkaline solution. This ring system is present in important biological building blocks, such as histidine and the related hormone histamine. Many drugs contain an imidazole ring, such as certain antifungal drugs, the nitroimidazole series of antibiotics, and the sedative midazolam [wikipedia].
imide	In organic chemistry, an imide is a functional group consisting of two acyl groups bound to nitrogen. Being highly polar, imides exhibit good solubility in polar media. The N–H center for imides derived from ammonia is acidic and can participate in hydrogen bonding. Unlike the structurally related acid anhydrides, they resist hydrolysis and some can even be recrystallized from boiling water [wikipedia]. Immunomodulatory imide drugs (IMiDs) are a class of immunomodulatory drugs (drugs that adjust immune responses) containing an imide group.
lactam	A beta-lactam (β -lactam) ring is a four-membered lactam. The β -lactam ring is part of the core structure of several antibiotic families, the principal ones being the penicillins, cephalosporins, carbapenems, and monobactams, which are, therefore, also called β -lactam antibiotics [wikipedia].
morpholine	Morpholine is an organic chemical compound having the chemical formula O(CH ₂ CH ₂) ₂ NH. Morpholine is a heterocycle featured in numerous approved and experimental drugs as well as bioactive molecules. It is often employed in the field of medicinal chemistry for its advantageous physicochemical, biological, and metabolic properties, as well as its facile synthetic routes [110].
NO (hydroxylamine)	Hydroxylamine is an organic compound with the formula NH ₂ OH. The material is a white crystalline, hygroscopic compound. Hydroxylamines and their derivatives are powerful aminating reagents, which are often used for arene C–H and X–H aminations (X=O, N, S, P) as well as Schmidt-type reaction ⁴⁶ , serving as alternative ways to introduce amino groups on various chemical skeletons [111].
oxazole	Oxazoles is a doubly unsaturated 5-membered ring having one oxygen atom at position 1 and a nitrogen at position 3 separated by a carbon in-between. Substitution pattern in oxazole derivatives play a pivotal role in delineating the biological activities like antimicrobial, anticancer, antitubercular anti-inflammatory, antidiabetic, antiobesity and antioxidant etc [112].
piperidine	Piperidine is an organic compound with the molecular formula (CH ₂) ₅ NH. This heterocyclic amine consists of a six-membered ring containing five methylene bridges (–CH ₂ –) and one amine bridge (–NH–). Piperidine and its derivatives are ubiquitous building blocks in pharmaceuticals [26] and fine chemicals. The piperidine structure is found in, for example: Icaridin, SSRIs, stimulants and nootropics, SERM etc. Piperidine is also commonly used in chemical degradation reactions, such as the sequencing of DNA in the cleavage of particular modified nucleotides. Piperidine is also commonly used as a base for the deprotection of Fmoc-amino acids used in solid-phase peptide synthesis [wikipedia].
piperazine	Piperazine is an organic compound that consists of a six-membered ring containing two nitrogen atoms at opposite positions in the ring. Many currently notable drugs contain a piperazine ring as part of their molecular structure (“Substituted piperazine”). Examples include: Antianginals, Antidepressants, Antihistamines etc [wikipedia].
pyridine	Pyridine is a basic heterocyclic organic compound with the chemical formula C ₅ H ₅ N. Pyridine moieties are often used in drugs because of their characteristics such as basicity, water solubility, stability, and hydrogen bond-forming ability, and their small molecular size [113]. Pyridine-based ring systems are one of the most extensively used heterocycles in the field of drug design, primarily due to their profound effect on pharmacological activity, which has led to the discovery of numerous broad-spectrum therapeutic agents [114].

tetrazole	Tetrazoles are a class of synthetic organic heterocyclic compound, consisting of a 5-member ring of four nitrogen atoms and one carbon atom. Tetrazole derivatives are a prime class of heterocycles, very important to medicinal chemistry and drug design due to not only their bioisosterism to carboxylic acid and amide moieties but also to their metabolic stability and other beneficial physicochemical properties [115].
thiazole	Thiazole, or 1,3-thiazole, is a heterocyclic compound that contains both sulfur and nitrogen. The versatility of thiazole nucleus demonstrated by the fact that it is an essential part of penicillin nucleus and some of its derivatives which have shown antimicrobial (sulfazole), antiretroviral (ritonavir), antifungal (abafungin), antihistaminic and antithyroid activities [116].
thiophene	Thiophene is a heterocyclic compound with the formula C ₄ H ₄ S. In medicine, thiophene derivatives shows antimicrobial, analgesic and anti-inflammatory, antihypertensive, and antitumor activity while they are also used as inhibitors of corrosion of metals or in the fabrication of light-emitting diodes in material science [117].
urea	Urea, also known as carbamide, is an organic compound with chemical formula CO(NH ₂) ₂ . This amide has two -NH ₂ groups joined by a carbonyl (C=O) functional group. It has similar effects as amide groups.

Table 12: Substructure detection, part I. We **bold** the best and underline the worst scores.

	allylic	amide	amidine	azo	benzene	epoxide	ether	furan	guanido	halogen	imidazole	imide
RANDOM	0.959	<u>16.917</u>	<u>0.054</u>	<u>0.020</u>	<u>1.100</u>	<u>0.024</u>	<u>2.024</u>	<u>0.036</u>	<u>0.126</u>	<u>1.127</u>	<u>0.080</u>	<u>0.062</u>
EDGE PRED	0.780	14.173	0.046	0.018	0.797	0.021	1.608	0.033	0.098	0.939	0.074	0.031
ATTR MASK	0.926	14.703	0.047	0.019	0.976	0.022	1.742	0.028	0.112	0.501	0.077	0.029
GPT-GNN	0.872	15.629	0.044	0.017	0.783	0.021	1.912	0.023	0.117	0.341	0.077	0.037
INFOGRAPH	0.740	6.747	0.050	0.019	0.583	0.022	1.128	0.021	0.086	0.706	0.062	0.038
CONT.PRED	<u>1.040</u>	16.636	0.053	0.020	0.980	0.023	1.787	0.034	0.126	1.075	0.078	0.033
GROVER	0.715	6.576	0.025	0.008	0.558	0.023	0.957	0.008	0.064	0.298	0.069	0.021
GRAPHCL	0.652	7.598	0.039	0.016	0.525	0.023	1.077	0.012	0.080	0.319	0.051	0.026
JOAO	0.654	7.926	0.043	0.015	0.531	0.023	1.071	0.013	0.085	0.310	0.048	0.026
GRAPHMVP	0.905	6.992	0.043	0.017	0.649	0.019	1.037	0.019	0.084	0.311	0.060	0.027
SSL Worse (#)	1	0	0	0	0	0	0	0	0	0	0	0

Table 13: Substructure detection, part II. We **bold** the best and underline the worst scores.

	lactam	morpholine	NO	oxazole	piperidine	piperzine	pyridine	tetrazole	thiazole	thiophene	urea
RANDOM	<u>0.018</u>	<u>0.031</u>	0.022	<u>0.009</u>	0.212	0.058	<u>0.176</u>	0.014	<u>0.040</u>	<u>0.052</u>	0.045
EDGE PRED	0.016	0.020	0.022	0.008	0.182	0.048	0.158	0.014	0.039	0.046	0.041
ATTR MASK	0.016	0.028	0.022	0.008	0.192	0.053	0.174	0.014	0.038	0.044	0.044
GPT-GNN	0.012	0.021	0.022	0.008	0.177	0.049	0.128	0.014	0.039	0.041	0.037
INFOGRAPH	0.012	0.026	0.024	0.007	0.185	0.048	0.127	<u>0.016</u>	0.029	0.036	<u>0.046</u>
CONT.PRED	0.016	<u>0.031</u>	0.022	0.008	<u>0.214</u>	<u>0.059</u>	0.170	0.014	0.039	0.047	0.045
GROVER	0.014	0.022	<u>0.028</u>	0.004	0.158	<u>0.043</u>	0.096	0.005	0.018	0.018	0.014
GRAPHCL	0.014	0.019	0.021	0.006	0.169	0.031	0.084	0.009	0.023	0.018	0.023
JOAO	0.014	0.021	0.020	0.006	0.168	0.035	0.082	0.009	0.022	0.018	0.025
GRAPHMVP	0.011	0.021	0.022	0.006	0.155	0.038	0.129	0.009	0.027	0.025	0.028
SSL Worse (#)	0	0	2	0	1	1	0	1	0	0	1

D.3 Cramer's V

Cramér's V quantifies the strength of the association between the molecular substructure counts (*i.e.*, chemical fragments) and their biochemical properties. It is defined as:

$$V = \sqrt{\chi^2 / (n \cdot \min(k - 1, r - 1))} = \sqrt{\chi^2 / n} \quad (r \equiv 2) \quad (5)$$

where n is the sample size, k and r are the total number of substructure counts and property categories (binary), respectively. The Chi-squared statistics χ^2 is then calculated as:

$$\chi^2 = \sum_{i,j} (n_{(i,j)} - n_{(i,\cdot)} \cdot n_{(\cdot,j)} / n)^2 / (n_{(i,\cdot)} \cdot n_{(\cdot,j)} / n) \quad (6)$$

where $n_{(i,j)}$ is the total occurrence for the pair of (i, j) . Here i is the specific count of a certain substructure, and j represents the certain outcome of a molecular biochemical property. Cramér's V value ranges from 0 to 1, representing the associated strength between two categorical variables.

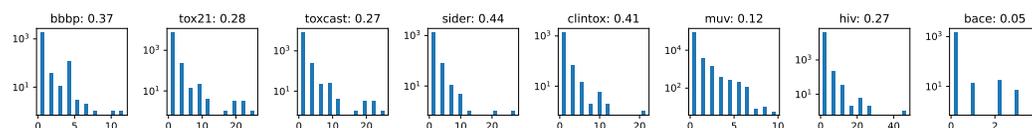
Table 14: Cramér's V between molecular substructure counts and downstream properties.

Pre-training	BBBP	Tox21	ToxCast	Sider	ClinTox	MUV	HIV	Bace	Avg (Task)	Avg (Data)
allylic	0.1602	0.1345	0.1156	0.1276	0.0935	0.0413	0.0280	0.1186	0.1144	0.1024
amide	0.2692	0.0490	0.0858	0.1841	0.1326	0.0235	0.0689	0.2556	0.0881	0.1336
amidine	0.0360	0.0291	0.0412	0.0323	0.0158	0.0117	0.0396	0.1328	0.0399	0.0423
azo	0.0400	0.0399	0.0393	0.0277	0.0123	0.0007	0.2082	-	0.0384	0.0526
benzene	0.1476	0.1632	0.1691	0.1149	0.1112	0.0289	0.1374	0.1091	0.1630	0.1227
epoxide	0.0273	0.0481	0.0449	0.0300	0.0049	0.0005	0.0086	-	0.0437	0.0235
ether	0.2314	0.0694	0.1060	0.1069	0.1023	0.0185	0.0498	0.1821	0.1034	0.1083
furan	0.0635	0.0257	0.0387	0.0227	0.0061	0.0311	0.0148	0.0135	0.0375	0.0270
guanido	0.0765	0.0201	0.0509	0.0715	0.0286	0.0057	0.0094	0.1088	0.0499	0.0464
halogen	0.1488	0.0849	0.1827	0.0773	0.0908	0.0143	0.0347	0.2353	0.1721	0.1086
imidazole	0.0601	0.0427	0.0492	0.0460	0.1212	0.0102	0.0398	0.1280	0.0483	0.0622
imide	0.0951	0.0246	0.0401	0.0428	0.0518	0.0094	0.0188	-	0.0392	0.0404
lactam	0.4263	0.0184	0.0116	0.0646	0.0543	0.0006	0.0048	-	0.0182	0.0830
morpholine	0.0512	0.0126	0.0343	0.0268	0.0425	0.0068	0.0101	0.0668	0.0329	0.0314
N_O	0.0438	0.0195	0.0467	0.0391	0.0709	0.0195	0.0144	0.0537	0.0452	0.0385
oxazole	0.0126	0.0184	0.0321	0.0359	0.0123	0.0079	0.0080	0.0364	0.0312	0.0205
piperidine	0.1450	0.0305	0.0844	0.0575	0.0418	0.0079	0.0226	0.0935	0.0803	0.0604
piperzine	0.0509	0.0214	0.0421	0.0776	0.0648	0.0111	0.0192	0.0063	0.0424	0.0367
pyridine	0.0598	0.0402	0.0549	0.0338	0.0833	0.0129	0.0300	0.1747	0.0529	0.0612
tetrazole	0.1161	0.0158	0.0251	0.0300	0.0286	0.0083	0.0123	0.0334	0.0247	0.0337
thiazole	0.1389	0.0521	0.0345	0.0445	0.0183	0.0118	0.0173	0.0539	0.0348	0.0464
thiophene	0.0356	0.0467	0.0472	0.0315	0.0113	0.0166	0.0081	0.0438	0.0456	0.0301
urea	0.0790	0.0236	0.0506	0.0471	0.0268	0.0079	0.0329	0.0516	0.0489	0.0399

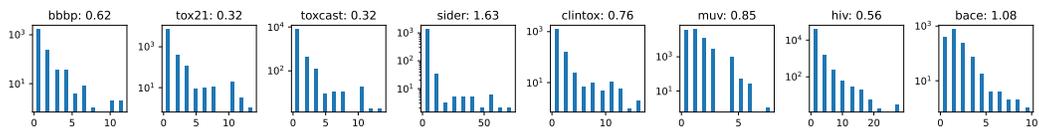
D.4 Distribution

We plot the distribution of 24 molecular substructures, the y-axis (*i.e.*, counts) is log scale. Alongside each substructure name, we includes its average counts (*i.e.*, ground truth). While certain substructures, such as amidine and azo, are infrequently found in the molecules from the eight downstream datasets, we note that others like Amide and Benzene are prevalent. These substructures are recognised for their significant relevance to molecular properties, as detailed in Table 4.

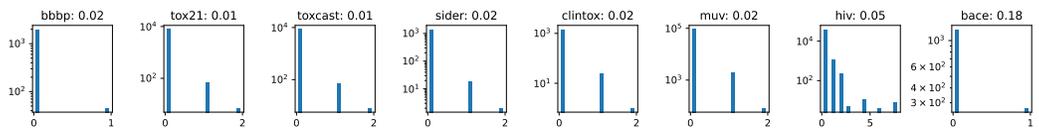
Allylic Oxide.



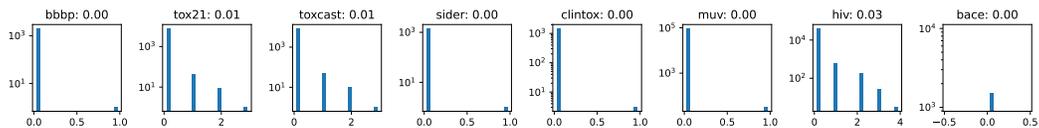
Amide.



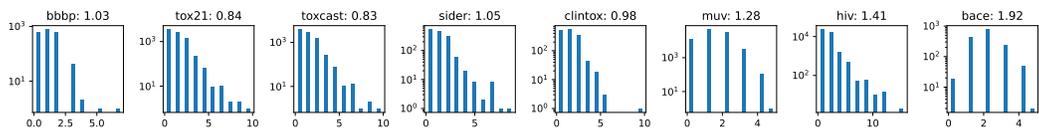
Amidine.



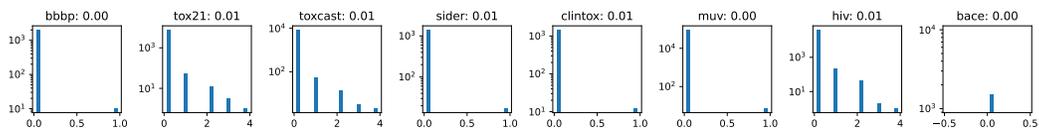
AZO.



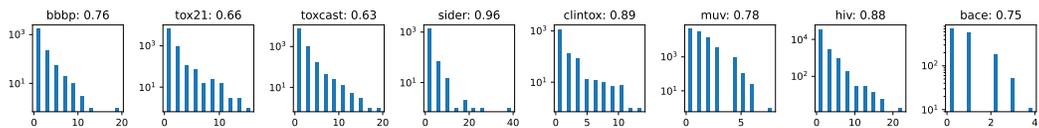
Benzene.



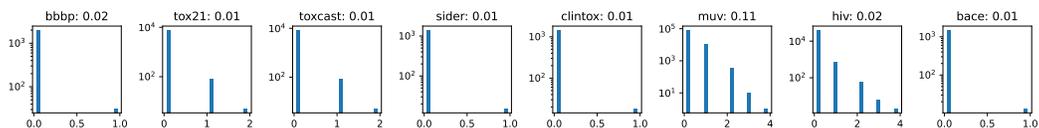
Epoxide.



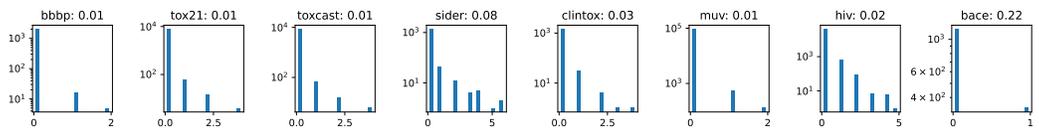
Ether.



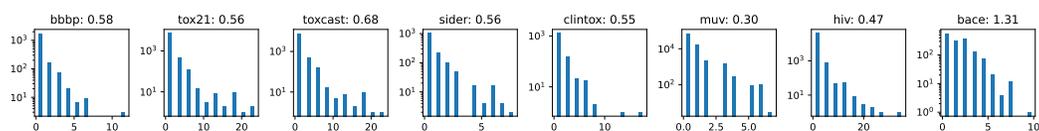
Furan.



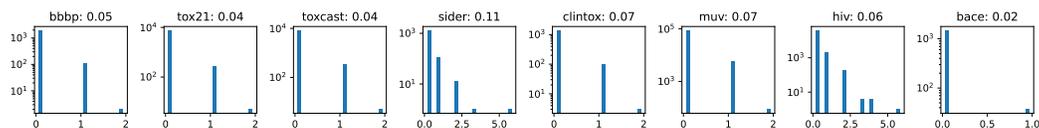
Guanido.



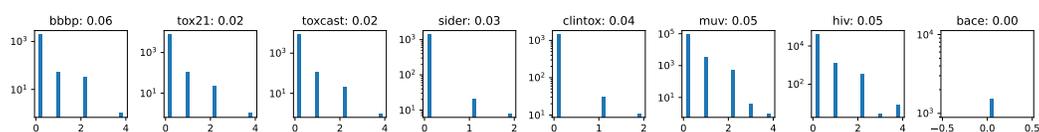
Halogen.



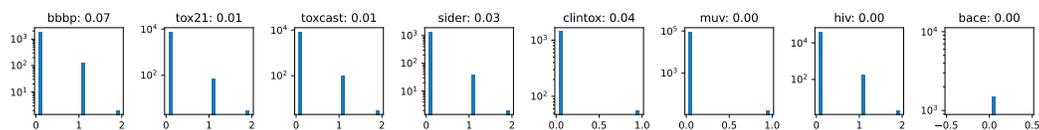
Imidazole.



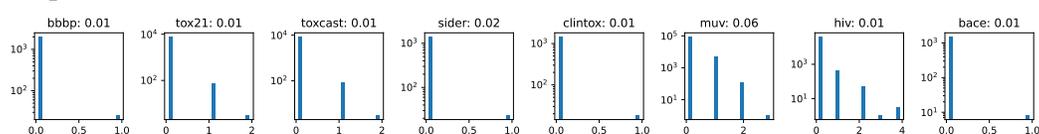
Imide.



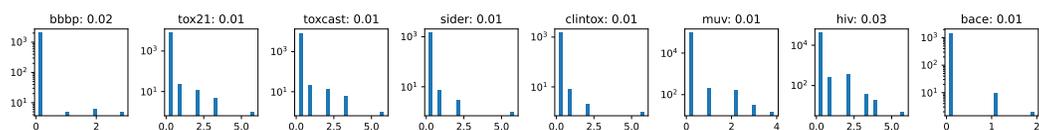
Lactam.



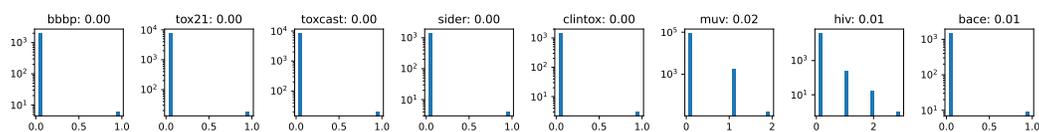
Morpholine.



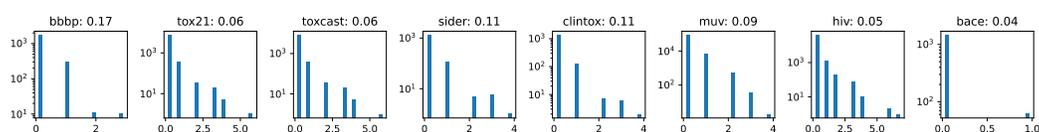
N_O.



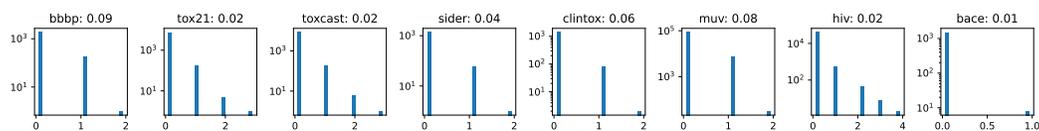
Oxazole.



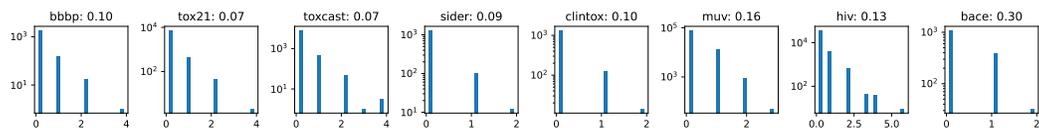
Piperidine.



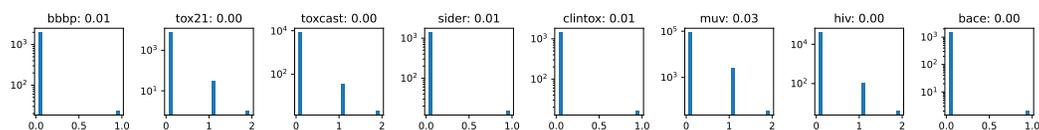
Piperzine.



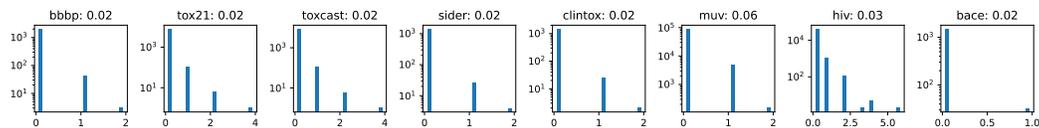
Pyridine.



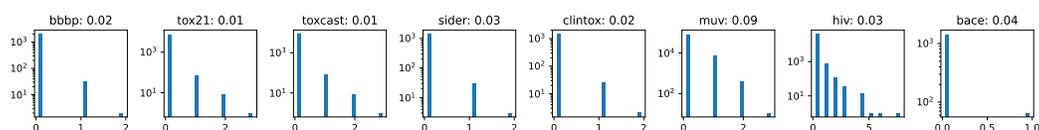
Tetrazole.



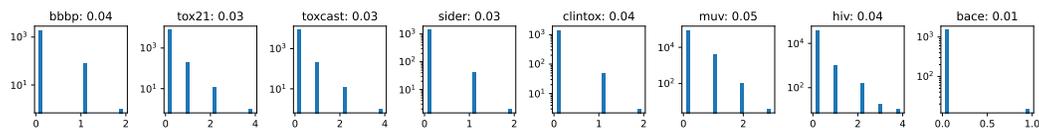
Thiazole.



Thiophene.



Urea.

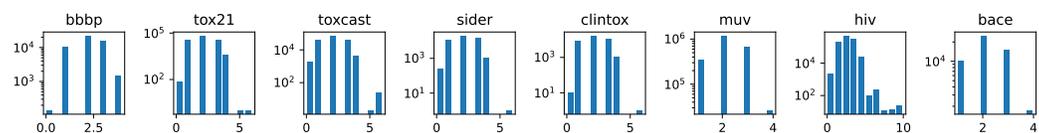


E More results on metrics

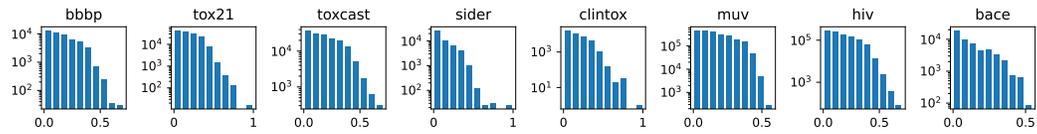
E.1 Topological metric

We have depicted the distribution of structural metrics and substructures with respect to the downstream datasets using histograms. Please note that the vertical axes may occasionally be represented on a logarithmic scale.

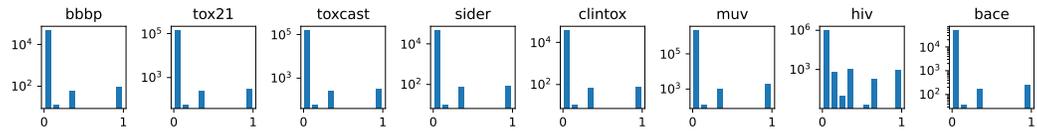
Node Degree.



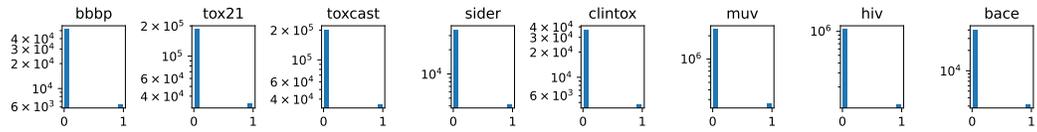
Node Centrality.



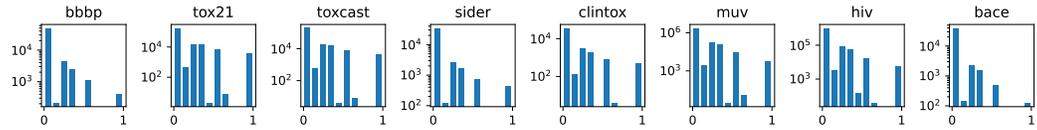
Node Clustering Coefficient.



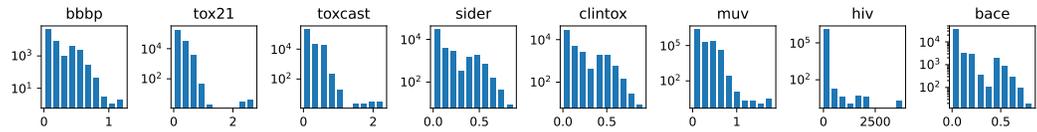
Link Prediction.



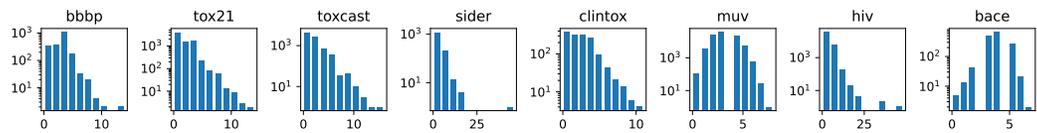
Jaccard Coefficient.



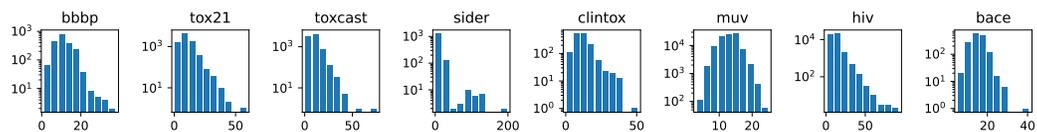
Katz Index.



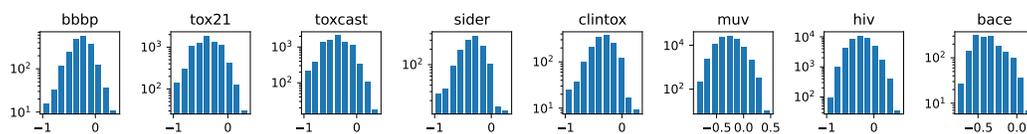
Cycle Basis.



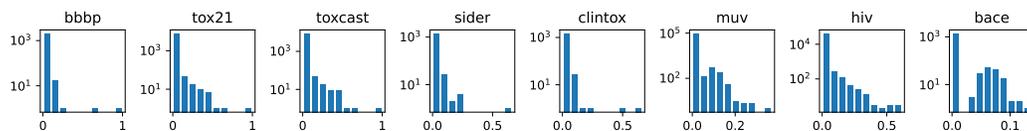
Graph Diameter.



Aassortativity Coefficient.



Average Clustering Coefficient.



E.2 Spectrum

Our observations suggest a positive correlation between the magnitudes of singular values in pre-trained embedding spaces and their performance in the downstream MPP tasks.

Table 15: Benchmarking the magnitude of the singular values (node/graph) in the spectrum, BBBP.

Dimension	1	50	100	200	300
RANDOM	1.96e-01/4.81e-02	8.14e-06/3.35e-07	1.08e-06/3.87e-08	9.10e-08/2.20e-09	2.44e-09/4.08e-11
EDGE PRED	1.25e+00/2.12e-01	5.30e-03/5.96e-04	3.92e-04/3.56e-05	1.31e-05/7.10e-07	1.97e-07/6.79e-09
ATTR MASK	9.97e+01/7.57e+00	9.71e-03/1.16e-03	2.06e-03/2.02e-04	2.95e-04/1.96e-05	1.86e-05/9.00e-07
GPT-GNN	7.78e+01/2.46e+01	4.21e-03/6.04e-04	6.26e-04/7.11e-05	5.32e-05/4.40e-06	1.29e-06/6.50e-08
INFOGRAPH	1.29e+02/6.56e+01	9.52e-01/3.23e-01	2.68e-01/6.54e-02	2.43e-02/3.38e-03	2.80e-05/2.69e-06
CONT. PRED	1.81e+01/5.89e+00	3.40e-04/3.72e-05	4.36e-05/3.42e-06	1.65e-06/9.08e-08	1.14e-08/3.92e-10
GROVER	2.21e+02/9.86e+01	2.18e+00/5.11e-01	1.13e-01/2.04e-02	1.68e-02/1.97e-03	1.51e-03/1.36e-04
GRAPHCL	7.63e+01/3.40e+01	4.26e-01/7.27e-02	1.45e-02/1.76e-03	6.84e-04/5.11e-05	2.42e-05/1.37e-06
JOAO	8.12e+01/3.58e+01	4.16e-01/7.89e-02	2.18e-02/2.67e-03	8.63e-04/6.74e-05	2.64e-05/1.52e-06
GRAPHMVP	2.07e+01/1.17e+01	2.57e-01/8.22e-02	1.94e-02/4.00e-03	4.39e-05/5.75e-06	1.76e-06/1.71e-07
Correlation	0.661/0.781	0.806/0.927	0.879/0.903	0.697/0.770	0.794/0.794

We provide more visualisations of the spectrum of the embedding space from different datasets.

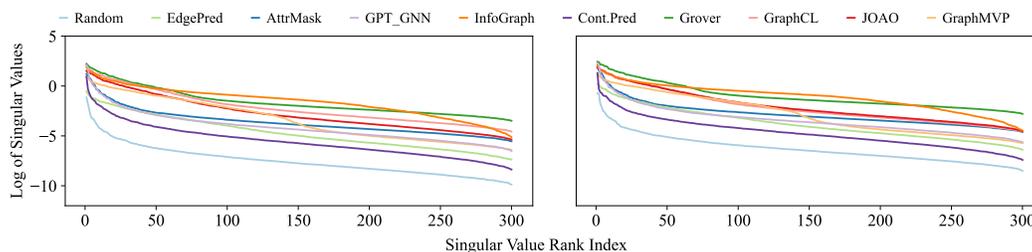


Figure 9: Spectrum of the GSSL embedding space on Tox21 dataset, Left: Node; Right: Graph.

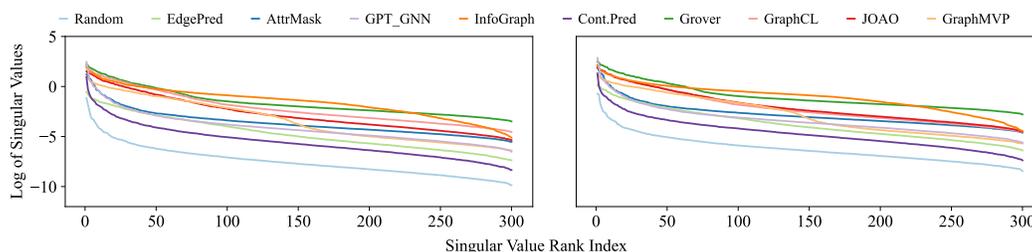


Figure 10: Spectrum of the GSSL embedding space on Toxcast dataset, Left: Node; Right: Graph.

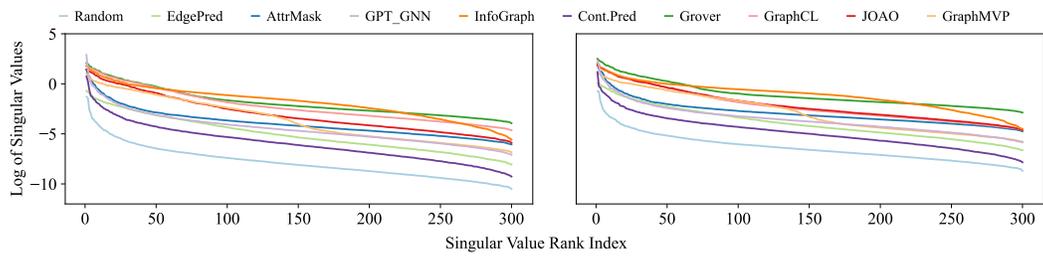


Figure 11: Spectrum of the GSSL embedding space on Sider dataset, Left: Node; Right: Graph.

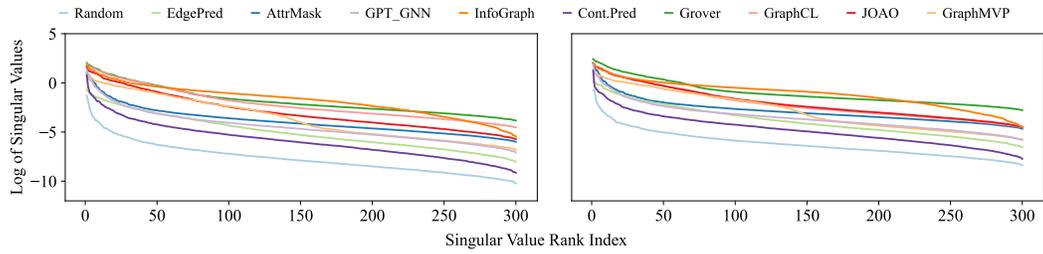


Figure 12: Spectrum of the GSSL embedding space on Clintox dataset, Left: Node; Right: Graph.

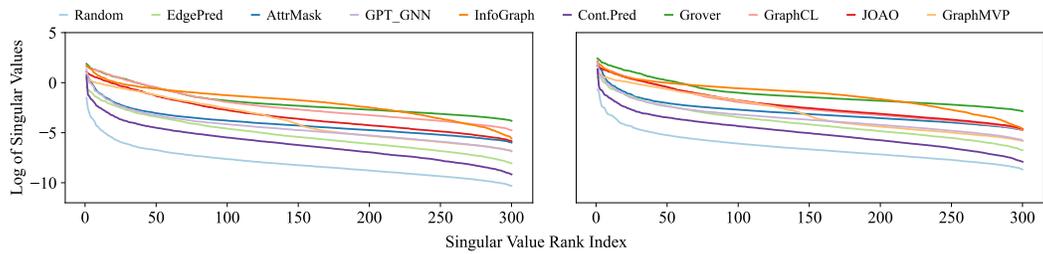


Figure 13: Spectrum of the GSSL embedding space on MUV dataset, Left: Node; Right: Graph.

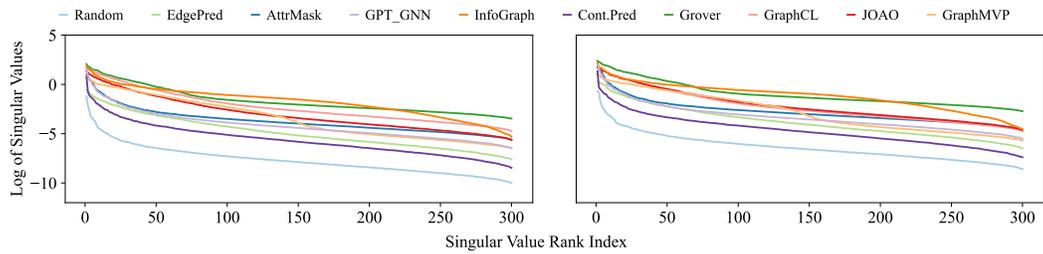


Figure 14: Spectrum of the GSSL embedding space on HIV dataset, Left: Node; Right: Graph.

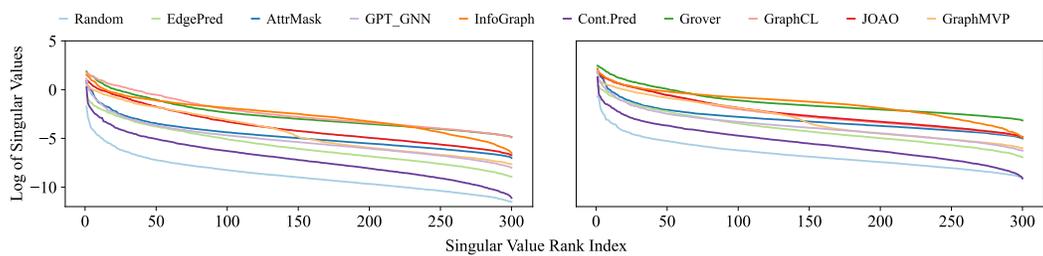


Figure 15: Spectrum of the GSSL embedding space on BACE dataset, Left: Node; Right: Graph.

E.3 Uniformity

We provide more comprehensive results on the uniformity metric of the embedding space in Table 16. The upper bound for uniformity is zero, while the lower bound, as derived from Corollary 1 in the Appendix of [6], is both data and hyperparameter-dependent.

Table 16: Evaluating GSSL methods on uniformity.

Dataset	BBBP	Tox21	ToxCast	Sider	ClinTox	MUV	HIV	Bace
RANDOM	-0.214	-0.315	-0.313	-0.282	-0.242	-0.128	-0.245	-0.079
EDGE PRED	-2.319	-2.969	-2.909	-2.424	-2.384	-2.634	-2.804	-1.653
ATTRMASK	-8.332	-9.348	-9.560	-7.786	-8.399	-9.241	-9.979	-6.389
GPT-GNN	-5.714	-5.701	-5.773	-5.550	-5.708	-5.279	-5.068	-4.977
INFOGRAPH	-10.065	-10.925	-11.310	-7.623	-9.831	-15.337	-13.506	-9.620
CONT.PRED	-2.636	-3.074	-3.209	-2.748	-2.780	-2.390	-2.865	-1.708
GROVER	-10.208	-12.142	-12.356	-11.907	-9.975	-15.772	-14.244	-10.512
GRAPHCL	-10.010	-11.073	-11.458	-10.155	-9.845	-13.390	-12.530	-8.513
JOAO	-10.000	-10.955	-11.360	-10.012	-9.846	-13.450	-12.505	-8.523
GRAPHMVP	-8.864	-9.770	-9.863	-8.782	-8.825	-12.405	-11.301	-7.540
Correlation	0.842	-0.600	0.097	-0.210	0.309	0.169	0.821	0.285
p-value	0.002	0.067	0.789	0.559	0.384	0.641	0.004	0.425

F Extending to more GSSL methods and datasets

We demonstrate that the proposed MOLGRAPHEVAL benchmark can readily be extended to incorporate more pre-training methods, such as GRAPHMAE and GRCL, as well as larger datasets, like the 2M ZINC15. It’s pertinent to note that the downstream datasets can also be employed for pre-training. Importantly, we observe that thorough optimisation of pre-training hyperparameters can serve as a crucial factor for performance improvements in pre-training, notwithstanding the advancements in the design of pre-training tasks.

Table 17: Evaluating GSSL methods on molecular property prediction tasks, on 2M ZINC15.

For each downstream dataset, we report the mean and standard deviation of the ROC-AUC scores over three random scaffold splits. The performance scores are based on the fixed pre-trained embeddings with linear probe models, we also report the average ROC-AUC scores with fine-tuned pre-trained GNN on MPP tasks (“Avg (FT)”).

	BBBP	Tox21	ToxCast	Sider	ClinTox	MUV	HIV	Bace	Avg	Avg (FT)
# Molecules	2,039	7,831	8,575	1,427	1,478	93,087	41,127	1,513	-	-
# Tasks	1	12	617	27	2	17	1	1	-	-
RANDOM	50.7 \pm 2.5	64.9 \pm 0.5	53.2 \pm 0.3	53.2 \pm 1.1	63.1 \pm 2.3	62.1 \pm 1.3	66.1 \pm 0.7	63.4 \pm 1.8	59.60	66.16
ATTRMASK	49.8 \pm 0.6	66.7 \pm 0.3	52.9 \pm 0.4	53.8 \pm 1.7	62.2 \pm 2.9	52.8 \pm 1.7	69.0 \pm 1.4	66.6 \pm 4.9	59.22	69.49
GRAPHCL	64.9 \pm 0.5	70.5 \pm 0.6	56.1 \pm 0.2	58.0 \pm 1.4	63.4 \pm 3.1	61.2 \pm 1.6	75.6 \pm 0.9	70.9 \pm 3.8	65.07	70.09
GRAPHMAE	58.6 \pm 2.3	64.4 \pm 0.6	55.3 \pm 0.1	57.0 \pm 1.8	67.4 \pm 1.9	57.3 \pm 1.5	71.6 \pm 1.2	51.6 \pm 3.6	60.41	68.91
GRCL	61.8 \pm 2.1	67.5 \pm 0.4	54.3 \pm 0.4	55.8 \pm 1.3	62.0 \pm 2.4	61.9 \pm 1.4	68.8 \pm 0.9	68.6 \pm 1.5	62.58	67.64

G Potential ethical and social implications

There are a few potential ethical considerations that could arise from the molecular graph representation learning methods discussed in :

Bias and fairness: We recognise that the molecular datasets used for pre-training and benchmarking may inadvertently encode biases or lack diversity, which could lead to models that unfairly under-

Table 18: **Benchmarking topological properties, on 2M ZINC15.** We report the mean square error or the cross entropy on eight datasets (*i.e.*, smaller is better).

	Node			Pair			Graph			
	Degree	Cent.	Cluster	Link	Jaccord	Katz	Diameter	Conn.	Cycle	Assort.
RANDOM	0.001	0.008	0.003	0.078	0.012	0.017	177.924	0.087	2.933	0.029
ATTRMASK	0.058	0.010	0.003	0.082	0.013	0.020	142.832	0.071	3.232	0.027
GRAPHCL	0.063	0.009	0.003	0.079	0.018	0.022	87.426	0.064	1.716	0.017
GRAPHMAE	0.032	0.011	0.004	6.885	0.215	0.038	123.677	0.066	3.094	0.025
RGCL	0.356	0.011	0.003	0.075	0.013	0.017	124.869	0.067	3.528	0.029

perform on certain data subsets. We have made efforts to use representative datasets, but further testing on diverse molecules and monitoring for fairness issues remain important future directions.

Relatedly, an inherent risk with data-driven methods is "stereotyping" along chemical lines, where models reinforce historical assumptions. We believe conscious testing on new classes of molecules can mitigate this. The goal is to build broad molecular understanding.

Dual use: As with many advances in chemistry ML, dual use concerns exist. While we aim to accelerate beneficial therapeutics, we strongly discourage misuse of these methods to design harmful substances. We support efforts toward responsible AI practices.

Environmental impact: Training large molecular graph models consumes significant computational resources and energy, efforts should be made to minimise the environmental impact, *e.g.*, by using efficient methods and carbon-neutral compute.

Privacy: While molecular graph data is less sensitive than data about individuals, any efforts to link models back to original private datasets should be handled carefully.

Overall, we make efforts to align research practices with principles of sharing benefits, avoiding harm, mitigating biases, and practising transparency. We welcome feedback from the community as we continue working to address these ethical dimensions. Our goal is innovation that aligns with broad social good.