

# Molecular Topological Profile (MOLTOP) - Simple and Strong Baseline for Molecular Graph Classification

Jakub Adamczyk<sup>a,\*</sup> and Wojciech Czech<sup>b</sup>

<sup>a</sup>Faculty of Computer Science, AGH University of Krakow, Krakow, Poland

<sup>b</sup>Faculty of Computer Science, AGH University of Krakow, Krakow, Poland

ORCID (Jakub Adamczyk): <https://orcid.org/0000-0003-4336-4288>, ORCID (Wojciech Czech): <https://orcid.org/0000-0002-1903-8098>

**Abstract.** We revisit the effectiveness of topological descriptors for molecular graph classification and design a simple, yet strong baseline. We demonstrate that a simple approach to feature engineering - employing histogram aggregation of edge descriptors and one-hot encoding for atomic numbers and bond types - when combined with a Random Forest classifier, can establish a strong baseline for Graph Neural Networks (GNNs). The novel algorithm, Molecular Topological Profile (MOLTOP), integrates Edge Betweenness Centrality, Adjusted Rand Index and SCAN Structural Similarity score. This approach proves to be remarkably competitive when compared to modern GNNs, while also being simple, fast, low-variance and hyperparameter-free. Our approach is rigorously tested on MoleculeNet datasets using fair evaluation protocol provided by Open Graph Benchmark. We additionally show out-of-domain generation capabilities on peptide classification task from Long Range Graph Benchmark. The evaluations across eleven benchmark datasets reveal MOLTOP's strong discriminative capabilities, surpassing the 1-WL test and even 3-WL test for some classes of graphs. Our conclusion is that descriptor-based baselines, such as the one we propose, are still crucial for accurately assessing advancements in the GNN domain.

## 1 Introduction

Graph classification has become a crucial type of supervised learning problem, increasingly relevant across various scientific domains. This surge in importance is largely attributed to the expanding quantity of structured datasets that represent pairwise relationships among various types of modeled entities. Graph classification algorithms are utilized in a variety of fields, particularly in chemoinformatics, where their application in Quantitative Structure-Activity Relationship (QSAR) modeling plays a critical role in predicting the functions of biochemically significant molecules [17]. Particularly, the prediction of ADME (Absorption, Distribution, Metabolism, Excretion) pharmacokinetic properties plays a pivotal role in supporting contemporary in-silico drug design [23].

Graph classification confronts a fundamental difficulty: measuring the dissimilarity between objects that are not situated in a metric space. Therefore, graphs, unlike more straightforward tabular or categorical data, require special methods to capture their complexity and relationships. Traditionally, this problem was solved by extracting isomorphism-invariant representations of graphs in the form of

feature vectors, also known as graph embeddings, descriptors, or fingerprints. Alternatively, explicit pairwise similarity measures, known as graph kernels, can be constructed to systematically compare graph substructures [35]. Both methods remain intrinsically unsupervised or task-independent, however domain-specific knowledge can be incorporated by careful feature engineering. Although graph descriptors have achieved success in various benchmark classification tasks, more recently, they are often surpassed by the more advanced graph representation learning models exemplified by Graph Neural Networks (GNNs). They learn task-specific representations and can take advantage of pre-training to reduce negative effects of limited training data [30, 46]. In developing a universal framework for graph classification, GNN models frequently incorporate descriptors, either as a method of input data augmentation or as supplementary global features in the readout layer [62, 21].

Given the computational expense of graph representation learning, the requirement for extensive training data, the challenge of transferring pre-trained knowledge to specialized prediction tasks, and the prevalence of domain-specific graph descriptors, comparing GNNs with traditional methods remains valuable. This is particularly true when descriptor-based methods serve as a baseline, indicating whether GNNs can learn additional, task-specific features. The studies [39] and [15] have identified significant obstacles hindering progress in the field of machine learning. These include challenges in effectively evaluating models, particularly issues related to non-replicable results and comparisons using inadequate baselines. Besides, the study presented in work [16] advocates for statistical rigor, when comparing classifiers across multiple datasets. More specifically, in the graph classification field, the authors of [22] describe problems with replicating GNN results caused by lack of strict separation between model selection and model evaluation step. Moreover, they show that under a fair comparison framework, simple structure-agnostic baselines can outperform GNN models such as GIN or GraphSAGE. In [44] the authors demonstrate that trivial 1-layer GCN can perform on par with complex GNNs such as DiffPool. The work [66] similarly notes the effectiveness of training-free vertex descriptors in link prediction tasks. In the realm of molecular graph classification it was shown that descriptor-based models, particularly those utilizing molecular fingerprints, not only yield better average prediction results than GNN models but also are computationally cheaper by an order of magnitude [49, 32]. The clear need of comparable prediction results and maintaining fair leaderboards led

\* Corresponding Author. Email: jadamczy@agh.edu.pl.

to the creation of benchmark datasets and related evaluation protocols such OGB [29], MoleculeNet [58] or TDC [31].

Motivated by research underscoring the value of robust baselines, and inspired by recent methods utilizing graph topology descriptors [10, 2], we propose Molecular Topological Profile (MOLTOP), a baseline method for molecular graph classification, utilizing both topological descriptors and simple atom and bond features. The resulting baseline, under the fair evaluation protocols offered by modern benchmarks, results in a surprisingly efficient and strong model, able to outperform contemporary GNNs. Our method is fast, scalable, robust in distinguishing graphs, non-parametric, and it exhibits low-variance in prediction tasks. Additionally, we present the studies verifying expressive power and feature importance of the proposed representation.

The code is available at <https://github.com/j-adamczyk/MOLTOP>, and supplementary material at <https://arxiv.org/abs/2407.12136>.

## 2 Related works

Graph descriptors, which generate isomorphism-invariant vectors representing graphs, exemplify the feature-engineering approach to graph classification. Descriptors are versatile in representing features at different levels – from granular to aggregated, local to global [13], and from purely structural aspects to those including multidimensional labels [36]. In practical applications, the descriptors from spectral graph theory [14, 51] or the ones using histogram aggregation of vertex/edge topological features [10, 2] have successfully rivaled more complex methods. The approach of generic graph descriptors was expanded by incorporating domain-specific representations, like molecular fingerprints, which have become widely used in predicting biochemical properties and molecular database search. Typical fingerprints are bit-vectors of a given size, built based on depth-first search explorations from each atom, and incorporating its 2D [19, 56] or 3D structure [5, 48]. The molecular property prediction based on molecular fingerprints can be highly competitive to GNNs, as shown in [32] and evident from OGB leaderboards [29].

Bypassing the need for manual feature engineering, GNNs provide an automated method for extracting task-specific graph features and transporting them directly to a trainable readout layer. Starting from early works introducing Graph Convolutional Network (GCN) [34] and GraphSAGE [27] the field of graph representation learning has evolved significantly, leading to the development of numerous models, as categorized by [41]. Some of these models, e.g., Graph Isomorphism Networks (GIN) [60] achieved state-of-the-art performance in benchmark graph classification tasks including molecular property prediction. GIN was designed to match the discriminative power of the Weisfeiler-Lehman isomorphism test, thereby offering additional insights into the representational capabilities of GNNs. Subsequently, in [62] the authors proposed a hybrid model D-MPNN, which combines edge-centered graph convolutions and molecular descriptors concatenated at the readout layer. That work represents a significant advancement in molecular graph classification, notable not only for its comprehensive and detailed analysis of model efficiency but also for its successful integration of the strengths of both GNNs and descriptors. In their work, [59] adopted the graph attention mechanism to develop the AttentiveFP model. This method is capable of utilizing atom and bond features, effectively extracting both local and global properties of a molecule.

When operating in a low-resource learning regime, GNNs often struggle to build discriminative representations of graphs. The success of transfer learning in the field of Natural Language Process-

ing (NLP), coupled with the scarcity of training data in molecular property prediction, has inspired researchers to adopt different pre-training strategies tailored for GNNs. In their work, [30] introduced a comprehensive framework that employs pre-training techniques like context prediction or attribute masking. This approach enables the transfer of knowledge from large molecular datasets to general-purpose GNNs, enhancing classification accuracy on benchmark tasks. In parallel, the transformer-style architecture GROVER was introduced by [54], reporting notable advancements over existing state-of-the-art methods. It utilized the largest pre-training database at the time, comprising 10 million unlabeled molecules. Graph Contrastive Learning (GraphCL) [64] was introduced as another self-supervised learning method, leveraging parameterized graph augmentations and maximizing the mutual information between augmentations sharing the same semantics. GraphCL was further extended by enabling automatic, adaptive selection of augmentation parameters [65] (JOAO). New pre-training strategies leveraging 2D topological structures extracted by encoders and enriched by 3D views led to development efficient GraphMVP framework [42]. More recently, the GEM model [24] proposed incorporating 3D molecular properties, based on Merck molecular force field (MMFF) simulations. Combining those geometric features with the GIN model and pretraining on 20 million molecules from the ZINC database led to exceptional performance in a range of graph classification and regression tasks, although at a very high computational cost. The most recent work, [46] describes relative molecule self-attention transformer (R-MAT), which uses atom embeddings reflecting relative distance between atoms. R-MAT reports SOTA results of molecular benchmarks, but uses different datasets and data splits than other models, therefore it is difficult to compare to this approach.

In contrast to multiple works reporting high efficiency of pre-trained GNN models, many thorough ablation studies, such as [57], provide contrary results. They present important findings on why feature engineering combined with low parameter machine learning can still outperform complex models, and why the pre-training benefits can be diminished in practical property prediction setups.

## 3 Preliminaries

**Molecular graph.** Let  $G = (V, E)$  denote an undirected graph representing a molecule, where  $V$  and  $E$  are the sets of vertices (nodes, atoms) and edges (links, bonds), respectively. We also mark  $G = (A, X_n, X_e)$ , where  $A$  is the adjacency matrix,  $X_n$  is the node feature matrix and  $X_e$  is the edge feature matrix.

**Graph notation.** We denote single vertices as  $v$  or  $u$ , and edges as two element vertex sets  $e = \{u, v\}$ .  $\mathcal{N}(v)$  is the set of neighbors of a node  $v$ ,  $\deg(v)$  is the degree of a node  $v$ , i.e. the number of its neighbors,  $\deg(v) = |\mathcal{N}(v)|$ .

**Graph classification.** We consider the graph classification task, where we are given a dataset  $\mathbb{D} = (G^{(i)}, Y^{(i)}), i = 1, 2, \dots, N$ , of  $N$  graphs and their labels. Class (label) for a given graph  $Y^{(i)}$  is a boolean for single-task datasets. For multitask datasets, it is a binary vector of length  $T$  (for  $T$  binary classification tasks), and it can have missing labels.

## 4 Method

We propose Molecular Topological Profile (MOLTOP) as a baseline method for benchmarking against GNNs in molecular graph classification tasks. Baselines are simple and computationally cheap methods, expected to provide a reference point for more sophisticated

methods. While not a focus point of any paper, they are a necessary part of fair valuation of new algorithms, especially on new datasets.

For MOLTOP, given its role as a baseline method, simplicity and speed are just as crucial as classification accuracy. In order to achieve good performance on chemical data, we utilize both topological and molecular features. The method relies on extracting feature vectors from graphs independently, and using Random Forest to classify resulting tabular data. In contrast to previous baselines, either purely topological (e.g. LDP [10] and LTP [2]), or purely feature-based (“molecular fingerprint” from [22]), it incorporates both graph structure and atoms and bonds features, all of which are crucial in chemistry.

The first group of features we consider are vertex degree statistics, to directly summarize the basic topology of a 2-hop neighborhood around each node [10]. We denote the multiset of vertex neighbors’ degrees as  $DN(v) = \{\deg(u) | u \in \mathcal{N}(v)\}$ . For each atom, we then calculate the following statistics:  $\deg(v)$ ,  $\min(DN(v))$ ,  $\max(DN(v))$ ,  $\text{mean}(DN(v))$ ,  $\text{std}(DN(v))$ . In order to create graph-level features, they are compactly represented using histograms, a technique akin to the global readout in GNNs, but with higher expressivity than just simple mean or sum [33].

For molecular graphs, especially in medicinal chemistry, having a degree higher than 8 is very rare. Using the same number of bins for all features would result in a very large number of all-zero features for many molecules. Therefore, we propose to reduce the number of bins to 11 for  $\deg(v)$ ,  $\min(DN(v))$  and  $\max(DN(v))$ . This covers singular hydrogens, covalent bonds, and nearly all atoms with higher degrees than 8 (e.g. due to ionic or metallic bonding) in typical biochemistry.

Inspired by previous structural approaches and path-based molecular fingerprints, we add further topological descriptors to enhance this representation. We select features that work well for describing molecular fragments, and that should discriminate well between different scaffolds and functional groups. Concretely, we selected Edge Betweenness Centrality (EBC), Adjusted Rand Index (ARI) and SCAN Structural Similarity score. Each of those descriptors is computed for edges (bonds), but focuses on a different aspect of molecule structure. EBC considers global graph connectivity structure and its shortest path-based properties. ARI uses 3-hop subgraphs and neighborhood connectivity patterns. SCAN also considers local connectivity patterns, but is based on the notion of node clusters and outliers. Therefore, those features should provide complementary information in the feature vector. Another reason for utilizing edge features is that similar edge-focused approaches were successful in improving GNNs for molecular property prediction [62, 33, 53]. Each of those features is calculated for all bonds in the molecule and aggregated with a histogram.

**Edge betweenness centrality (EBC)** [26] is a centrality measure for edges, defined as a fraction of shortest paths going through that edge:

$$EBC(e) = \frac{2}{|V|(|V| - 1)} \sum_{u,v \in V} \frac{\sigma_{u,v}(e)}{\sigma_{u,v}}, \quad (1)$$

where  $\sigma_{u,v}$  is the total number of shortest paths between  $u$  and  $v$ , and  $\sigma_{u,v}(e)$  is the number of that paths going through  $e$ . The normalization factor before the sum ensures that the values lie in range  $[0, 1]$  and are unaffected by graph size. Information about the shortest paths in the graph is well known to be important in chemistry, being used e.g. in Wiener index and Hyper-Wiener index [18], and was also successfully incorporated into multiple GNNs [63, 53]. However, the histogram of centralities includes more information than the lengths

of shortest paths would, because it shows the actual distribution of critically important edges. If there are bonds with very high EBC values, it indicates the existence of bridge-like subgraphs, such as glycosidic bonds. It can also easily distinguish between linear and polycyclic scaffolds, since the ring-rich topologies will have smaller EBC values in general, while linear structures have many high-EBC bonds.

**Adjusted Rand Index (ARI)** [28] is a normalized measure of overlap between neighborhoods of two vertices  $u$  and  $v$ :

$$ARI(u, v) = \frac{2(ad - bc)}{(a + b)(b + d) + (a + c)(c + d)},$$

where  $a$  is the number of edges to other vertices that  $u$  and  $v$  have in common,  $b$  is the number of edges to other nodes for  $u$  that  $v$  does not have,  $c$  is the number of edges to other nodes for  $v$  that  $u$  does not have, and  $d$  is the number of edges to other nodes that neither  $u$  nor  $v$  has. Calculated for edge  $e = \{u, v\}$ , it provides information about the subgraphs of radius 3 (from neighbors of  $v$ , through edge  $e$ , to neighbors of  $u$ ).

Among various neighborhood overlap measures, ARI has a particularly strong statistical interpretation, being equivalent to Kohen’s  $\kappa$  defined on incident edge sets of  $u$  and  $v$  [28]. While this measure is typically used for link prediction, it can also be calculated for existing edges. This method has been used for identifying ‘incorrect’ links, where it surpassed other techniques [28], and a similar approach was also used in LTP [2]. Therefore, the histogram of ARI values should work well for existing edges, taking into consideration larger subgraphs than degree features and indicating the general connectivity patterns in a graph. In particular, it is capable of differentiating between star-like graphs (such as spiro compounds or functional groups containing atoms with high coordination number, e.g. phosphate groups) and polycyclic molecules characterized by grid-like subgraphs, like polycyclic aromatic hydrocarbons.

**SCAN** [61, 11], used for node clustering and graph sparsification, defines the structural similarity score for edges as:

$$SCAN(u, v) = \frac{|\mathcal{N}(u) \cap \mathcal{N}(v)| + 1}{\sqrt{(\deg(u) + 1)(\deg(v) + 1)}} \quad (2)$$

SCAN scores were designed to detect the edges critical for graph connectivity, and those corresponding to outliers. In molecules, the distribution of SCAN scores can easily distinguish between linear structures (e.g. alkanes with long carbon chains), where the scores are low in general, and well-connected, ring-rich molecules (e.g. steroids).

Molecular graph classification relies heavily on atom and bond features, meaning that baselines utilizing only graph topology are not expressive enough. In fact, purely feature-based “molecular fingerprint” baseline from [22], using only atom counts (i.e. counts of different chemical elements), can outperform some GNNs. Therefore, MOLTOP incorporates two such features: atomic numbers and bond types. They are the most apparent and consistently available features, universally employed by GNNs for analyzing molecules [30, 29, 59].

For atoms, we one-hot encode the atomic numbers up to 89 (with zero marking unknown types). We discard actinides and all further molecules, since they are all radioactive and extremely rarely used. For each chemical element, we compute its mean, standard deviation and sum (total count in the molecule) as graph-level features. We do the same for bonds, with 5 possible types (single, double, triple, aromatic, or miscellaneous). In principle, one could add further features the same way, if they are known to be important for a given problem, e.g. chirality.

All features are computed for each graph independently, with the same number of bins  $n_{bins}$  for all histograms. This results in  $11 \cdot n_{bins} + 90 \cdot 3 + 5 \cdot 3$ . This can result in many all-zero features, especially for atom features. We simply drop all such columns, based on the training data. Therefore, the number of features is often significantly reduced after this step.

The only hyperparameter of the feature extraction is the number of bins for histograms  $n_{bins}$ . In other works [10, 2, 33] this is either a hyperparameter, requiring tuning, or just an arbitrarily set number. For MOLTOP, we propose a data-driven solution instead, setting the number of bins equal to the median size of the molecules (i.e. number of atoms) in the dataset. This is motivated by the fact that molecule sizes for drug-like compounds typically follow a right-skewed, single-modal distribution with number of atoms very rarely exceeding 50 (see the supplementary material [1] for plots). This fact is often used in medicinal chemistry, e.g. by Lipinski’s rule of 5 [40]. Therefore, for the vast majority of data, it would bring little benefit to use a high number of bins. With this addition, MOLTOP does not require any hyperparameter tuning for feature extraction.

After feature extraction, we use Random Forest (RF) as a classifier. RF serves as an effective prediction model due to its low computational complexity and high scalability. Moreover, its performance is less sensitive to hyperparameter choices, unlike other commonly used classifiers such as SVMs or boosting methods [52]. It also natively supports multitask learning, which is common in molecular graph classification. The dimensionality of our representation is quite high, therefore we use larger number of trees and stronger regularization than default settings in Scikit-learn [50], to better incorporate new features and prevent overfitting. Based on the average results on validation sets of MoleculeNet (detailed in Section 5), MOLTOP uses 1000 trees, the entropy as splitting criterion, and minimum of 10 samples to perform a split. Those reasonable defaults make it a hyperparameter-free method, making it extremely easy to use and computationally cheap, which is important for baseline methods.

#### 4.1 Complexity analysis

The computational complexity of MOLTOP is the sum of complexities of its features, since they are computed independently. Vertex degree features have complexity  $O(|E|)$  [10]. Computing EBC has complexity  $O(|V||E|)$  [9]. Calculation of both ARI and SCAN Structural Similarity scores for all edges is pessimistically  $O(|V||E|)$ , but the expected complexity for molecular graphs is  $O(|E|)$  due to their sparsity (for proofs, see the supplementary material). The total complexity of feature extraction is thus  $O(|V||E|)$ .

## 5 Experiments and results

For the main evaluation of the proposed method, we selected 8 classification datasets from MoleculeNet benchmark [58] (described in detail in the supplementary material [1]), the most widely used molecular graph classification benchmark. For fair evaluation, we used deterministic scaffold split, with the splits provided by OGB [29]. This setting is much more challenging than random split, which does not enforce out-of-distribution generalization. In addition, 5 of those datasets are multitask, including massively multitask ToxCast dataset with 617 targets. In all cases, we follow the recommendation from [30], training 10 models with different random seeds, and we report mean and standard deviation of AUROC.

For implementation of MOLTOP feature extraction, we used PyTorch Geometric [25] and NetworkKit [3]. Those frameworks provide

efficient data structures and parallel processing. For Random Forest, we use Scikit-learn [50]. Since this implementation does not allow missing labels in the training set, we fill them with zeros. This is acceptable, since those tasks are already imbalanced, and such adjustment makes this even a bit more challenging.

### 5.1 Validation set experiments

During initial experiments, we verified our modelling choices by using average AUROC on validation sets. This setting was chosen, since we aimed to design a general baseline, that performs well on average for molecular classification. We started with only degree features with 50 bins (inspired by [2]), and added proposed improvements one by one. First, we validated that adding other topological descriptors, e.g. other centrality scores than EBC or other neighborhood overlap than ARI, gave results worse or similar to our proposed descriptors. Next, we confirmed that using all proposed statistics of atoms and bonds is crucial. Furthermore, we verified that using median molecule size as the number of bins gave results better or comparable to manual tuning. Lastly, we performed hyperparameter tuning of Random Forest, and resulting values were 1000 trees, entropy splitting criterion, and minimum of 10 samples to perform a split. Those align with our postulate to use more trees and stronger regularization.

We summarize the impact of adding described improvements in Table 1 (for more detailed tables see the supplementary material [1]). We report average AUROC and standard deviation for test sets of all 8 MoleculeNet datasets. All proposed changes improve the results, in particular the introduction of atom and bonds features in addition to pure topology. This shows that effective baselines for molecular data have to use both structure and domain-relevant features.

**Table 1:** The results of model improvements.

Model	Avg. AUROC $\uparrow$
Degree features, 50 bins	63.8 $\pm$ 0.6
Add topological edge features	65.5 $\pm$ 0.9
Add atoms and bonds features	69.0 $\pm$ 0.8
Median bins	69.4 $\pm$ 0.9
Reduce degrees bins, drop constant features	70.4 $\pm$ 0.7
Use tuned Random Forest hyperparameters	72.5 $\pm$ 0.5

### 5.2 MoleculeNet classification

We compared MOLTOP to 18 other graph classification methods on MoleculeNet benchmark, with results in Table 2. We compare it to methods from three groups: general-purpose GNNs, GNNs designed specifically for molecular data, and graph classification baselines. This way, we verify not only that MOLTOP improves upon previous baselines, but also achieves strong performance in comparison to sophisticated, domain-specific models.

We include 8 general-purpose GNNs: GIN, GCN and GraphSAGE from [30], both with and without context prediction pretraining, as well as recent models based on contrastive learning, GraphCL [64] and JOAO [65]. For GNNs designed specifically for molecular property prediction, we include multiple recent models utilizing different approaches to incorporating molecular features: D-MPNN [62], AttentiveFP [59], GROVER [54] (large variant), GraphMVP [43] (regular and contrastive variants), and GEM [24]. We also compare to four other baselines: purely topological LDP [10] and LTP [2], purely feature-based “molecular fingerprint” from [22] (which uses atom counts as features), and ECFP, the molecular fingerprint commonly

**Table 2:** Classification results on MoleculeNet. “Pretr” denotes if the model is pretrained. The best result for each dataset is bolded, and also for each model group (general-purpose GNNs, molecular GNNs, baselines), the best model and its average AUROC and rank are bolded.

Model	Pretr.	BACE	BBBP	HIV	ClinTox	MUV	SIDER	Tox21	ToxCast	Avg. AUROC $\uparrow$	Avg. rank $\downarrow$
GIN	No	70.1 $\pm$ 5.4	65.8 $\pm$ 4.5	75.3 $\pm$ 1.9	71.8 $\pm$ 2.5	58.0 $\pm$ 4.4	57.3 $\pm$ 1.6	74.0 $\pm$ 0.8	63.4 $\pm$ 0.6	67 $\pm$ 2.7	14.4
<b>GIN</b>	Yes	84.5 $\pm$ 0.7	68.7 $\pm$ 1.3	79.9 $\pm$ 0.7	81.3 $\pm$ 2.1	72.6 $\pm$ 1.5	62.7 $\pm$ 0.8	78.1 $\pm$ 0.6	65.7 $\pm$ 0.6	<b>74.2 <math>\pm</math> 1</b>	<b>4.3</b>
GCN	No	73.6 $\pm$ 3.0	64.9 $\pm$ 3.0	75.7 $\pm$ 1.1	73.2 $\pm$ 1.4	65.8 $\pm$ 4.5	60.0 $\pm$ 1.0	74.9 $\pm$ 0.8	63.3 $\pm$ 0.9	68.9 $\pm$ 2	12.3
GCN	Yes	82.3 $\pm$ 3.4	70.6 $\pm$ 1.6	78.2 $\pm$ 0.6	79.4 $\pm$ 1.8	63.6 $\pm$ 1.7	62.4 $\pm$ 0.5	75.8 $\pm$ 0.3	65.3 $\pm$ 0.1	72.2 $\pm$ 1.3	6.1
GraphSAGE	No	72.5 $\pm$ 1.9	69.6 $\pm$ 1.9	74.4 $\pm$ 0.7	72.7 $\pm$ 1.4	59.2 $\pm$ 4.4	60.4 $\pm$ 1.0	74.7 $\pm$ 0.7	63.3 $\pm$ 0.5	68.4 $\pm$ 1.6	12.1
GraphSAGE	Yes	80.7 $\pm$ 0.9	63.9 $\pm$ 2.1	76.2 $\pm$ 1.1	78.4 $\pm$ 2.0	60.7 $\pm$ 2.0	60.7 $\pm$ 0.5	76.8 $\pm$ 0.3	64.9 $\pm$ 0.2	70.3 $\pm$ 1.1	9.6
GraphCL	Yes	68.7 $\pm$ 7.8	67.5 $\pm$ 3.3	75.0 $\pm$ 0.4	78.9 $\pm$ 4.2	77.1 $\pm$ 1.0	60.1 $\pm$ 1.3	75.0 $\pm$ 0.3	62.8 $\pm$ 0.2	70.6 $\pm$ 2.3	11.1
JOAO	Yes	72.9 $\pm$ 2.0	66.0 $\pm$ 0.6	76.6 $\pm$ 0.5	66.3 $\pm$ 3.9	77.0 $\pm$ 2.2	60.7 $\pm$ 1.0	74.4 $\pm$ 0.7	62.7 $\pm$ 0.6	69.6 $\pm$ 1.4	11.4
D-MPNN	No	80.9 $\pm$ 0.6	71.0 $\pm$ 0.3	77.1 $\pm$ 0.5	<b>90.6 <math>\pm</math> 0.6</b>	78.6 $\pm$ 1.4	57.0 $\pm$ 0.7	75.9 $\pm$ 0.7	65.5 $\pm$ 0.3	74.6 $\pm$ 0.6	5.9
AttentiveFP	No	78.4 $\pm$ 2.2	64.3 $\pm$ 1.8	75.7 $\pm$ 1.4	84.7 $\pm$ 0.3	76.6 $\pm$ 1.5	60.6 $\pm$ 3.2	76.1 $\pm$ 0.5	63.7 $\pm$ 0.2	72.5 $\pm$ 1.4	9
GROVER	Yes	81.0 $\pm$ 1.4	69.5 $\pm$ 0.1	68.2 $\pm$ 1.1	76.2 $\pm$ 3.7	67.3 $\pm$ 1.8	65.4 $\pm$ 0.1	73.5 $\pm$ 0.1	65.3 $\pm$ 0.5	70.8 $\pm$ 1.1	9.1
GraphMVP	Yes	76.8 $\pm$ 1.1	68.5 $\pm$ 0.2	74.8 $\pm$ 1.4	79.0 $\pm$ 2.5	75.0 $\pm$ 1.4	62.3 $\pm$ 1.6	74.5 $\pm$ 0.4	62.7 $\pm$ 0.1	71.7 $\pm$ 1.1	10.3
GraphMVP-C	Yes	81.2 $\pm$ 0.9	72.4 $\pm$ 1.6	77.0 $\pm$ 1.2	77.5 $\pm$ 4.2	75.0 $\pm$ 1.0	63.9 $\pm$ 1.2	74.4 $\pm$ 0.2	63.1 $\pm$ 0.4	73.1 $\pm$ 1.3	7.1
<b>GEM</b>	Yes	<b>85.6 <math>\pm</math> 1.1</b>	<b>72.4 <math>\pm</math> 0.4</b>	80.6 $\pm$ 0.9	90.1 $\pm$ 1.3	<b>81.7 <math>\pm</math> 0.5</b>	<b>67.2 <math>\pm</math> 0.4</b>	<b>78.1 <math>\pm</math> 0.1</b>	<b>69.2 <math>\pm</math> 0.4</b>	<b>78.1 <math>\pm</math> 0.6</b>	<b>1.3</b>
ECFP	No	83.8 $\pm$ 0.4	68.6 $\pm$ 0.5	76.3 $\pm$ 0.6	71.7 $\pm$ 1.6	66.9 $\pm$ 1.3	67.1 $\pm$ 0.4	72.8 $\pm$ 0.2	60.4 $\pm$ 0.4	71 $\pm$ 0.7	9.9
“molecular fingerprint”	No	71.5 $\pm$ 0.2	68.3 $\pm$ 0.3	65.5 $\pm$ 0.6	65.5 $\pm$ 0.9	49.8 $\pm$ 0.0	59.0 $\pm$ 0.2	63.5 $\pm$ 0.2	57.5 $\pm$ 0.1	62.6 $\pm$ 0.3	17.3
LDP	No	80.5 $\pm$ 0.3	63.3 $\pm$ 0.4	72.1 $\pm$ 0.4	58.2 $\pm$ 2.1	50.0 $\pm$ 0.9	59.8 $\pm$ 0.5	66.7 $\pm$ 0.2	59.5 $\pm$ 0.4	63.8 $\pm$ 0.3	16
LTP	No	80.7 $\pm$ 0.3	65.6 $\pm$ 0.3	73.0 $\pm$ 0.7	61.7 $\pm$ 1.7	53.2 $\pm$ 1.7	60.8 $\pm$ 0.5	67.7 $\pm$ 0.5	60.0 $\pm$ 0.4	65.3 $\pm$ 0.5	13.8
<b>MOLTOP</b>	No	82.9 $\pm$ 0.2	68.9 $\pm$ 0.2	<b>80.8 <math>\pm</math> 0.3</b>	73.6 $\pm$ 0.7	66.7 $\pm$ 1.9	66.0 $\pm$ 0.5	76.3 $\pm$ 0.2	64.4 $\pm$ 0.3	<b>72.5 <math>\pm</math> 0.5</b>	<b>6</b>

used as a baseline for GNNs [62] (using default settings). For those baselines, we use Random Forest with 500 trees as a classifier, which follows [2] and is a common setting in chemoinformatics.

Following best practices for statistical comparison of classifiers from [16], we report average model rank across datasets, in addition to average AUROC. This metric is less influenced by outliers among scores, and therefore better measures how the model really performs on average. In particular, the ClinTox dataset often gives very unstable results [57], and the average rank should be less susceptible to this problem.

The main observation is that MOLTOP, under this fair comparison protocol, outperforms the majority of models on average, often by a large margin. In terms of average rank, it exceeds all GNNs without pretraining except for D-MPNN, which has almost identical average rank. It also has results better than most pretrained GNNs, even including recent, complex models like JOAO, GROVER and GraphMVP. This is particularly significant, since MOLTOP does not utilize any external knowledge like those models, nor did it require very costly pretraining on massive datasets. Our results are also notably stable, with low standard deviations, indicating the robustness of this approach.

MOLTOP does not require any pretraining, and requires only around 50 minutes for the entire benchmark (with massively multitask ToxCast taking the majority of the time). In addition, it has very low standard deviations, indicating stable and robust behavior. This shows that fair comparison, using strong baselines, remains important even in the era of large pretrained models.

Outperforming GNNs can be explained by the global nature of features used by MOLTOP. Those models, while sophisticated, still rely on an inherently local message-passing paradigm, and especially without pretraining it is hard for them to fully understand molecular relations on limited data.

The only models that have better average rank than MOLTOP are pretrained GIN, D-MPNN, and GEM. However, using Wilcoxon signed-rank test (recommended by [16]) with  $\alpha = 0.05$ , we determined that difference with GIN and D-MPNN performance is not statistically significant (p-values 0.547 and 0.742). Only GEM outperforms MOLTOP significantly (p-value 0.016), but we note that it has an enormous computational cost, including fine-tuning and even inference, since it requires generation of multiple conformers and Merck molecular force field (MMFF) optimization. Those operations can easily take minutes per molecule, can often fail for molecules

with complicated geometries (e.g. highly rotatable bonds), and are simply impossible in many cases, e.g. for compounds with disconnected components like salts. Therefore, MOLTOP always achieves results better or as good as GNNs, except for GEM, which has major practical downsides.

MOLTOP also improves upon other baselines by a large margin. Previous approaches like LDP, LTP and “molecular fingerprint” of [22] often fail to beat almost any GNNs, and thus are unsuitable for molecular data. Notably, we even outperform ECFP4 fingerprint, often used to compare against GNNs. This shows that improving upon existing baselines remains important for fair comparison.

We present additional comparisons with graph kernels in the supplementary material [1]. We omit them here, because due to OOM errors they couldn’t be computed on HIV and MUV datasets, meaning that we can’t directly compare their average AUROC and rank to other models in Table 2.

### 5.3 HIV leaderboard results

We further evaluate MOLTOP on the HIV dataset featured in OGB leaderboard [29], comparing it to various cutting-edge models that do not provide results on the whole MoleculeNet benchmark. The results are shown in Table 3. While this is the same HIV dataset as used before, here we are not allowed to use the validation data (due to leaderboard rules), even when no hyperparameters are tuned.

Since there are currently 34 models on the leaderboard, here we present a few selected ones. MOLTOP achieves 14-th rank, outperforming well-known PNA [12] and DeeperGCN [38], and coming very close to Graphormer [63], GSAT [47] and CIN [7]. It is also narrowly better than very powerful Directional GSN [8]. If we lift the limitation of not using the validation set, which is quite arti-

**Table 3:** Selected results on HIV leaderboard in OGB. MOLTOP results are marked in bold.

Method	Rank $\downarrow$	Test AUROC $\uparrow$	Valid AUROC $\uparrow$
CIN	9	80.94 $\pm$ 0.57	82.77 $\pm$ 0.99
GSAT	10	80.67 $\pm$ 0.95	83.47 $\pm$ 0.31
Graphormer	14	80.51 $\pm$ 0.53	83.10 $\pm$ 0.89
<b>MOLTOP</b>	<b>13</b>	<b>80.42<math>\pm</math>0.25</b>	<b>80.33<math>\pm</math>0.54</b>
P-WL	15	80.39 $\pm$ 0.40	82.79 $\pm$ 0.59
Directional GSN	15	80.39 $\pm$ 0.90	84.73 $\pm$ 0.96
PNA	18	79.05 $\pm$ 1.32	85.19 $\pm$ 0.99
DeeperGCN	21	78.58 $\pm$ 1.17	84.27 $\pm$ 0.63

cial for hyperparameter-free MOLTOP, it gets 80.8% AUROC and outperforms both GSAT and Graphormer.

#### 5.4 Peptides classification

In order to further evaluate the out-of-distribution generalization abilities of MOLTOP, we utilize the peptides-func dataset from LRGB benchmark [20], concerning peptide function classification. The characteristics of this data are very different from MoleculeNet, with peptides being much bigger molecules, with larger diameter and long-range dependencies. We do not perform any tuning, requiring the hyperparameter-free baseline to perform reasonably well even on this very different domain.

In Table 4, we compare to the results from [20]. Remarkably, MOLTOP outperforms all GNNs, including graph transformers specifically designed for this task, e.g. SAN with RWSE embeddings. At the same time, it is much more stable, with very low standard deviation. This shows that it indeed works very well as a baseline, even for novel datasets and molecular domains.

**Table 4:** Results on peptides-func dataset. Best result is bolded.

Method	Test average precision $\uparrow$
GCN	59.30 $\pm$ 0.23
GCNII	55.43 $\pm$ 0.78
GINE	54.98 $\pm$ 0.79
GatedGCN	58.64 $\pm$ 0.77
GatedGCN+RWSE	60.69 $\pm$ 0.35
Transformer+LapPE	63.26 $\pm$ 1.26
SAN+LapPE	63.84 $\pm$ 1.21
SAN+RWSE	64.39 $\pm$ 0.75
MOLTOP	<b>64.59 <math>\pm</math> 0.05</b>

#### 5.5 Time efficiency benchmark

While MOLTOP has very low feature extraction complexity, as outlined in Section 4.1, we also measure wall time for both feature extraction and RF training, summarized in Table 5. On four smallest datasets, it requires less than ten seconds for both, and at most about a minute for a further three datasets. In particular, feature extraction takes only 35 seconds on over 15 thousands of peptides, which are very large molecules by the standard of molecular graph classification, which mostly concerns small, drug-like compounds. On MUV, which is by far the largest in terms of the number of molecules, feature extraction still takes only about 2.5 minutes. In general, we note that MOLTOP feature extraction is embarrassingly parallel, and can process almost arbitrary number of molecules, given enough CPUs.

**Table 5:** MOLTOP timings.

Dataset	# molecules	# tasks	Feature extraction [s]	Training [s]
BACE	1513	1	2	1
BBBP	2039	1	3	1
HIV	41127	1	57	6
ClinTox	1478	17	2	1
MUV	93087	2	144	38
SIDER	1427	27	2	2
Tox21	7831	12	11	4
ToxCast	8575	617	12	232
Peptides-func	15535	10	35	5

Overall, the ToxCast takes the most time, but for the training part, due to being massively multitask. This is most likely the artifact of the implementation, since such dataset are rare and Scikit-learn implementation of RF is not particularly optimized for those cases. In

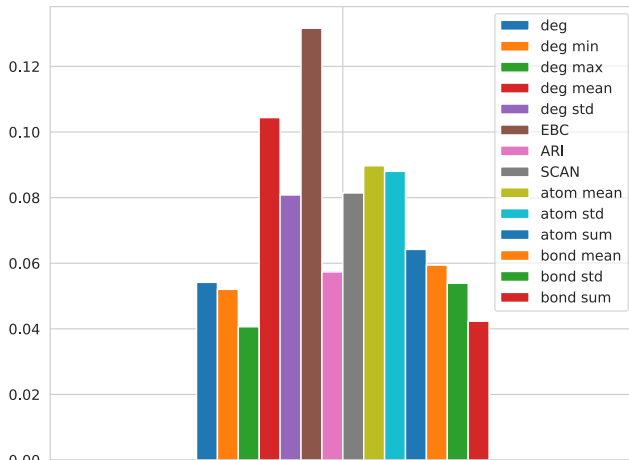
fact, this is the only dataset in molecular property prediction that we are aware of with such huge number of tasks.

For comparison with GNNs, we focus on peptides-func dataset, for which [20] provides wall times. Computing LapPE or RWSE embeddings alone, which are necessary for feature augmentation to get reasonable performance of GNNs on this dataset (due to long-range dependencies), takes about a minute. This does not even take into consideration the training time of GNN model itself.

Finally, since MOLTOP is hyperparameter-free, it does not require time for tuning for new datasets. This is especially advantageous in comparison to GNNs, which require extensive tuning of at least learning rate and regularization parameters for new datasets. This increases their training cost multiple times, while MOLTOP can just be used as-is.

#### 5.6 Feature importance analysis

We additionally validate the importance of features leveraging Random Forest average decrease in entropy. This metric is effective in identifying the features that are most useful for the model. The importance of a feature is the sum of importances of its histograms bins, since each bin is treated as a separate feature for the classifier. Next, we average values obtained from 10 classifiers on each dataset, based on different random seeds. To aggregate this information for the entire benchmark, we further average the importances for all 8 datasets. This is shown in Fig. 1.



**Figure 1:** Average MOLTOP feature importances.

The main outcome is that all features are useful, as there are none with very low importance. The most influential feature is EBC, which highlights that global information is particularly important for molecular data, and validates our initial belief. The least useful feature is the maximal degree of neighbors, which is expected, as node degrees are typically very low in chemistry. For additional ablation studies, see the supplementary material [1].

#### 5.7 Expressivity experiments

Lastly, we analyzed the expressive power of MOLTOP topological features in distinguishing graphs. It is typically represented using a hierarchy of  $k$ -dimensional Weisfeiler-Lehman isomorphism tests [37, 55], but can also be verified by using particular classes of graphs, which are known to be hard to distinguish for computational methods. Typical GNNs are at most as powerful as 1-WL test [60], but

with specific extensions, often utilizing topological descriptors in forms of shortest paths or subgraphs counting, GNNs become more powerful [63, 8].

We verified the discriminative power of MOLTOP feature vectors using *graph8c* and *sr25* datasets [6]. Our topological features are all integers after histogram aggregation, so we deem two graphs to be different if they have different values of any feature. We perform paired comparisons of graphs this way, where the number of pairs is 61M for *graph8c* and 105 for *sr25*. We report number of errors, i.e. undistinguished pairs, in Table 6.

MOLTOP achieves very good results, showing high power in distinguishing graphs. It outperforms all message-passing GNNs, probably due to usage of features that incorporate more global information. Additionally, it performs almost as well as PPGN [45] and GNNML3 [6], which are provably as powerful as 2-FWL test, equivalent to 3-WL test. It also achieves perfect result on *sr25*, which consists of strongly regular graphs. This is particularly exceptional, as they are 3-WL equivalent [4], which means that MOLTOP can distinguish graphs for which even 3-WL test fails.

We provide examples of graphs distinguishable by MOLTOP, but not e.g. by 1-WL test, in the supplementary material [1].

**Table 6:** The number of undistinguished pairs of graphs in *graph8c* and *sr25*.

Model	graph8c ↓	sr25 ↓
MLP	293K	105
GCN	4775	105
GAT	1828	105
GIN	386	105
ChebNet	44	105
PPGN	0	105
GNNML1	333	105
GNNML3	0	105
MOLTOP	3	0

## 6 Conclusion

We presented a new type of molecular graph embedding, which leverages local and global structural information aggregated from vertex and edge descriptors, as well as basic semantics of bonds and atoms. Combined with low parameter classification using Random Forests, it forms a robust baseline algorithm for molecular property prediction called MOLTOP. The key advantages of MOLTOP are: low computational cost, no hyperparameter tuning required, and high discriminative power, which surpasses 1-WL isomorphism test. Based on fair evaluation protocols and deterministic scaffold splits, we show that MOLTOP is surprisingly competitive with GNNs, including out-of-generalization applications to new datasets. With additional verification of results using Wilcoxon signed-rank test, we show that our proposed model is better or as good as all baselines and GNNs, except for GEM model, which uses computationally expensive and error-prone 3D molecular modelling.

In the future work, we plan to experiment with incorporating additional features, and adapt this approach to e.g. materials chemistry. We also want to more thoroughly analyze the theoretical aspects of feature descriptors and their discriminative abilities in terms of WL hierarchy.

We conclude that strong baselines, such as MOLTOP, are still important to gain deep insights into advances of GNN pre-training and assessing benefits of incorporating spatial or structural information, especially in the experimental setups with limited computational budgets.

## Acknowledgements

Research was supported by the funds assigned by Polish Ministry of Science and Higher Education to AGH University of Krakow, and by the grant from “Excellence Initiative - Research University” (IDUB) for the AGH University of Krakow. We gratefully acknowledge Poland’s high-performance Infrastructure PLGrid ACK Cyfronet AGH for providing computer facilities and support within computational grant. We would like to thank Alexandra Elbakyan for her work and support for accessibility of science.

## References

- [1] J. Adamczyk and W. Czech. Supplementary material for “Molecular Topological Profile (MOLTOP) - Simple and Strong Baseline for Molecular Graph Classification”. URL <https://arxiv.org/abs/2407.12136>.
- [2] J. Adamczyk and W. Czech. Strengthening Structural Baselines for Graph Classification Using Local Topological Profile. In *International Conference on Computational Science*, pages 597–611. Springer, 2023.
- [3] E. Angriman, A. van der Grinten, M. Hamann, H. Meyerhenke, and M. Penschuck. Algorithms for Large-scale Network Analysis and the NetowKit Toolkit. In *Algorithms for Big Data: DFG Priority Program 1736*, pages 3–20. Springer Nature Switzerland Cham, 2023.
- [4] V. Arvind, F. Fuhlbrück, J. Köbler, and O. Verbitsky. On Weisfeiler-Leman Invariance: Subgraph Counts and Related Graph Properties. *Journal of Computer and System Sciences*, 113:42–59, 2020.
- [5] S. D. Axen, X.-P. Huang, E. L. Cáceres, L. Gendele, B. L. Roth, and M. J. Keiser. A Simple Representation of Three-Dimensional Molecular Structure. *Journal of Medicinal Chemistry*, 60(17):7393–7409, 2017.
- [6] M. Balcilar, P. Héroux, B. Gauzere, P. Vasseur, S. Adam, and P. Honeine. Breaking the Limits of Message Passing Graph Neural Networks. In *International Conference on Machine Learning*, pages 599–608. PMLR, 2021.
- [7] C. Bodnar, F. Frasca, N. Otter, Y. Wang, P. Lio, G. F. Montufar, and M. Bronstein. Weisfeiler and Lehman Go Cellular: CW Networks. *Advances in Neural Information Processing Systems*, 34:2625–2640, 2021.
- [8] G. Bouritsas, F. Frasca, S. Zafeiriou, and M. M. Bronstein. Improving Graph Neural Network Expressivity via Subgraph Isomorphism Counting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):657–668, 2022.
- [9] U. Brandes. A Faster Algorithm for Betweenness Centrality. *Journal of Mathematical Sociology*, 25(2):163–177, 2001.
- [10] C. Cai and Y. Wang. A simple yet effective baseline for non-attributed graph classification. *arXiv preprint arXiv:1811.03508*, 2018.
- [11] Y. Chen, H. Ye, S. Vedula, A. Bronstein, R. Dreslinski, T. Mudge, and N. Talati. Demystifying Graph Sparsification Algorithms in Graph Properties Preservation. In *50th International Conference on Very Large Databases (VLDB 2024)*. ACM, 2024.
- [12] G. Corso, L. Cavalleri, D. Beaini, P. Liò, and P. Veličković. Principal Neighbourhood Aggregation for Graph Nets. *Advances in Neural Information Processing Systems*, 33:13260–13271, 2020.
- [13] W. Czech. Invariants of distance k-graphs for graph embedding. *Pattern Recognition Letters*, 33(15):1968–1979, 2012.
- [14] N. de Lara and E. Pineau. A Simple Baseline Algorithm for Graph Classification. *Relational Representation Learning Workshop, NIPS 2018*, 2018.
- [15] M. Dehghani, Y. Tay, A. A. Gritsenko, Z. Zhao, N. Houlsby, F. Diaz, D. Metzler, and O. Vinyals. The benchmark lottery. *arXiv preprint arXiv:2107.07002*, 2021.
- [16] J. Demšar. Statistical Comparisons of Classifiers over Multiple Data Sets. *The Journal of Machine Learning Research*, 7:1–30, 2006.
- [17] J. Deng, Z. Yang, H. Wang, I. Ojima, D. Samaras, and F. Wang. A systematic study of key elements underlying molecular property prediction. *Nature Communications*, 14(1):6395, 2023.
- [18] M. V. Diudea and I. Gutman. Wiener-Type Topological Indices. *Croatia Chemica Acta*, 71(1):21–51, 1998.
- [19] J. L. Durant, B. A. Leland, D. R. Henry, and J. G. Nourse. Reoptimization of MDL keys for use in drug discovery. *Journal of Chemical Information and Computer Sciences*, 42(6):1273–1280, 2002.
- [20] V. P. Dwivedi, L. Rampášek, M. Galkin, A. Parviz, G. Wolf, A. T. Luu, and D. Beaini. Long Range Graph Benchmark. *Advances in Neural Information Processing Systems*, 35:22326–22340, 2022.

- [21] V. P. Dwivedi, C. K. Joshi, A. T. Luu, T. Laurent, Y. Bengio, and X. Bresson. Benchmarking Graph Neural Networks. *Journal of Machine Learning Research*, 24(43):1–48, 2023.
- [22] F. Errica, M. Podda, D. Bacciu, and A. Micheli. A Fair Comparison of Graph Neural Networks for Graph Classification. In *Proceedings of the 8th International Conference on Learning Representations (ICLR)*, 2020.
- [23] C. Fang, Y. Wang, R. Grater, S. Kapadnis, C. Black, P. Trapa, and S. Sciabola. Prospective Validation of Machine Learning Algorithms for Absorption, Distribution, Metabolism, and Excretion Prediction: An Industrial Perspective. *Journal of Chemical Information and Modeling*, 2023.
- [24] X. Fang, L. Liu, J. Lei, D. He, S. Zhang, J. Zhou, F. Wang, H. Wu, and H. Wang. Geometry-enhanced molecular representation learning for property prediction. *Nature Machine Intelligence*, 4(2):127–134, 2022.
- [25] M. Fey and J. E. Lenssen. Fast Graph Representation Learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*, 2019.
- [26] M. Girvan and M. E. Newman. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences*, 99(12):7821–7826, 2002.
- [27] W. Hamilton, Z. Ying, and J. Leskovec. Inductive Representation Learning on Large Graphs. *Advances in Neural Information Processing Systems*, 30, 2017.
- [28] M. Hoffman, D. Steinley, and M. J. Brusco. A Note on Using the Adjusted Rand Index for Link Prediction in Networks. *Social Networks*, 42:72–79, 2015.
- [29] W. Hu, M. Fey, M. Zitnik, Y. Dong, H. Ren, B. Liu, M. Catasta, and J. Leskovec. Open Graph Benchmark: Datasets for Machine Learning on Graphs. *Advances in Neural Information Processing Systems*, 33:22118–22133, 2020.
- [30] W. Hu, B. Liu, J. Gomes, M. Zitnik, P. Liang, V. Pande, and J. Leskovec. Strategies for Pre-training Graph Neural Networks. In *International Conference on Learning Representations*, 2020.
- [31] K. Huang, T. Fu, W. Gao, Y. Zhao, Y. Roohani, J. Leskovec, C. W. Coley, C. Xiao, J. Sun, and M. Zitnik. Therapeutics Data Commons: Machine Learning Datasets and Tasks for Drug Discovery and Development. *Proceedings of Neural Information Processing Systems, NeurIPS Datasets and Benchmarks*, 2021.
- [32] D. Jiang, Z. Wu, C.-Y. Hsieh, G. Chen, B. Liao, Z. Wang, C. Shen, D. Cao, J. Wu, and T. Hou. Could graph neural networks learn better molecular representation for drug discovery? a comparison study of descriptor-based and graph-based models. *Journal of Cheminformatics*, 13(1):1–23, 2021.
- [33] S. Kearnes, K. McCloskey, M. Berndl, V. Pande, and P. Riley. Molecular Graph Convolutions: Moving Beyond Fingerprints. *Journal of Computer-Aided Molecular Design*, 30:595–608, 2016.
- [34] T. N. Kipf and M. Welling. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations (ICLR)*, 2017.
- [35] N. M. Kriege, F. D. Johansson, and C. Morris. A Survey on Graph Kernels. *Applied Network Science*, 5(1):1–42, 2020.
- [36] R. Lazarz and M. Idzik. Relation Order Histograms as a Network Embedding Tool. In *Computational Science–ICCS 2021: 21st International Conference, Krakow, Poland, June 16–18, 2021, Proceedings, Part II 21*, pages 224–237. Springer, 2021.
- [37] A. Leman and B. Weisfeiler. A reduction of a graph to a canonical form and an algebra arising during this reduction. *Nauchno-Tekhnicheskaya Informatsiya*, 2(9):12–16, 1968.
- [38] G. Li, C. Xiong, A. Thabet, and B. Ghanem. DeeperGCN: All You Need to Train Deeper GCNs. *arXiv preprint arXiv:2006.07739*, 2020.
- [39] T. Liao, R. Taori, I. D. Raji, and L. Schmidt. Are We Learning Yet? A Meta Review of Evaluation Failures Across Machine Learning. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [40] C. A. Lipinski, F. Lombardo, B. W. Dominy, and P. J. Feeney. Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Advanced Drug Delivery Reviews*, 64:4–17, 2012.
- [41] R. Liu, S. Cantürk, F. Wenkel, S. McGuire, X. Wang, A. Little, L. O’Bray, M. Perlmutter, B. Rieck, M. Hirn, et al. Taxonomy of Benchmarks in Graph Representation Learning. In *Learning on Graphs Conference*, pages 6–1. PMLR, 2022.
- [42] S. Liu, H. Wang, W. Liu, J. Lasenby, H. Guo, and J. Tang. Pre-training Molecular Graph Representation with 3D Geometry. *arXiv preprint arXiv:2110.07728*, 2021.
- [43] S. Liu, H. Wang, W. Liu, J. Lasenby, H. Guo, and J. Tang. Pre-training Molecular Graph Representation with 3D Geometry. In *International Conference on Learning Representations*, 2022.
- [44] E. Luzhnica, B. Day, and P. Liò. On Graph Classification Networks, Datasets and Baselines. *arXiv preprint arXiv:1905.04682*, 2019.
- [45] H. Maron, H. Ben-Hamu, H. Serviansky, and Y. Lipman. Provably Powerful Graph Networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- [46] Ł. Maziarka, D. Majchrowski, T. Danel, P. Gaiński, J. Tabor, I. Podolak, P. Morkisz, and S. Jastrzębski. Relative Molecule Self-Attention Transformer. *Journal of Cheminformatics*, 16(1):3, 2024.
- [47] S. Miao, M. Liu, and P. Li. Interpretable and Generalizable Graph Learning via Stochastic Attention Mechanism. In *International Conference on Machine Learning*, pages 15524–15543. PMLR, 2022.
- [48] R. Nilakantan, N. Bauman, J. S. Dixon, and R. Venkataraghavan. Topological torsion: a new molecular descriptor for SAR applications. Comparison with other descriptors. *Journal of Chemical Information and Computer Sciences*, 27(2):82–85, 1987.
- [49] A. Pappu and B. Paige. Making Graph Neural Networks Worth It for Low-Data Molecular Machine Learning. In *Machine Learning for Molecules Workshop @ NeurIPS 2020*, 2020.
- [50] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [51] E. Pineau. Using Laplacian Spectrum as Graph Feature Representation. *arXiv preprint arXiv:1912.00735*, 2019.
- [52] P. Probst, A.-L. Boulesteix, and B. Bischl. Tunability: Importance of Hyperparameters of Machine Learning Algorithms. *The Journal of Machine Learning Research*, 20(1):1934–1965, 2019.
- [53] L. Rampásek, M. Galkin, V. P. Dwivedi, A. T. Luu, G. Wolf, and D. Beaini. Recipe for a General, Powerful, Scalable Graph Transformer. *Advances in Neural Information Processing Systems*, 35:14501–14515, 2022.
- [54] Y. Rong, Y. Bian, T. Xu, W. Xie, Y. Wei, W. Huang, and J. Huang. Self-Supervised Graph Transformer on Large-Scale Molecular Data. *Advances in Neural Information Processing Systems*, 33:12559–12571, 2020.
- [55] R. Sato. A Survey on The Expressive Power of Graph Neural Networks. *arXiv preprint arXiv:2003.04078*, 2020.
- [56] N. Stiefl, I. A. Watson, K. Baumann, and A. Zaliani. ERG: 2D Pharmacophore Descriptions for Scaffold Hopping. *Journal of Chemical Information and Modeling*, 46(1):208–220, 2006.
- [57] R. Sun, H. Dai, and A. W. Yu. Does GNN Pretraining Help Molecular Representation? *Advances in Neural Information Processing Systems*, 35:12096–12109, 2022.
- [58] Z. Wu, B. Ramsundar, E. N. Feinberg, J. Gomes, C. Geniesse, A. S. Pappu, K. Leswing, and V. Pande. MoleculeNet: A Benchmark for Molecular Machine Learning. *Chemical Science*, 9(2):513–530, 2018.
- [59] Z. Xiong, D. Wang, X. Liu, F. Zhong, X. Wan, X. Li, Z. Li, X. Luo, K. Chen, H. Jiang, et al. Pushing the Boundaries of Molecular Representation for Drug Discovery with the Graph Attention Mechanism. *Journal of Medicinal Chemistry*, 63(16):8749–8760, 2019.
- [60] K. Xu, W. Hu, J. Leskovec, and S. Jegelka. How Powerful are Graph Neural Networks? In *International Conference on Learning Representations*, 2019.
- [61] X. Xu, N. Yuruk, Z. Feng, and T. A. Schweiger. SCAN: A Structural Clustering Algorithm for Networks. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 824–833, 2007.
- [62] K. Yang, K. Swanson, W. Jin, C. Coley, P. Eiden, H. Gao, A. Guzman-Perez, T. Hopper, B. Kelley, M. Mathea, et al. Analyzing Learned Molecular Representations for Property Prediction. *Journal of Chemical Information and Modeling*, 59(8):3370–3388, 2019.
- [63] C. Ying, T. Cai, S. Luo, S. Zheng, G. Ke, D. He, Y. Shen, and T.-Y. Liu. Do Transformers Really Perform Bad for Graph Representation? *Advances in Neural Information Processing Systems*, 34:28877–28888, 2021.
- [64] Y. You, T. Chen, Y. Sui, T. Chen, Z. Wang, and Y. Shen. Graph Contrastive Learning with Augmentations. *Advances in Neural Information Processing Systems*, 33:5812–5823, 2020.
- [65] Y. You, T. Chen, Y. Shen, and Z. Wang. Graph Contrastive Learning Automated. In *International Conference on Machine Learning*, pages 12121–12132. PMLR, 2021.
- [66] W. Zhang, Z. Sheng, M. Yang, Y. Li, Y. Shen, Z. Yang, and B. Cui. NAFS: A Simple yet Tough-to-beat Baseline for Graph Representation Learning. In *International Conference on Machine Learning*, pages 26467–26483. PMLR, 2022.