
Will More Expressive Graph Neural Networks do Better on Generative Tasks?

Xiandong Zou

Imperial College London, UK
xz1320@ic.ac.uk

Xiangyu Zhao

Imperial College London, UK
x.zhao22@imperial.ac.uk

Pietro Liò

University of Cambridge, UK
pietro.lio@cl.cam.ac.uk

Yiren Zhao

Imperial College London, UK
a.zhao@imperial.ac.uk

Abstract

Graph generation poses a significant challenge as it involves predicting a complete graph with multiple nodes and edges based on simply a given label. This task also carries fundamental importance to numerous real-world applications, including de-novo drug and molecular design. In recent years, several successful methods have emerged in the field of graph generation. However, these approaches suffer from two significant shortcomings: (1) the underlying Graph Neural Network (GNN) architectures used in these methods are often underexplored; and (2) these methods are often evaluated on only a limited number of metrics. To fill this gap, we investigate the expressiveness of GNNs under the context of the molecular graph generation task, by replacing the underlying GNNs of graph generative models with more expressive GNNs. Specifically, we analyse the performance of six GNNs in two different generative frameworks—autoregressive generation models, such as GCPN and GraphAF, and one-shot generation models, such as GraphEBM—on six different molecular generative objectives on the ZINC-250k dataset. Through our extensive experiments, we demonstrate that advanced GNNs can indeed improve the performance of GCPN, GraphAF, and GraphEBM on molecular generation tasks, but GNN expressiveness is not a necessary condition for a good GNN-based generative model. Moreover, we show that GCPN and GraphAF with advanced GNNs can achieve state-of-the-art results across 17 other non-GNN-based graph generative approaches, such as variational autoencoders and Bayesian optimisation models, on the proposed molecular generative objectives (DRD2, Median1, Median2), which are important metrics for de-novo molecular design.

1 Introduction

Graph generation has always been viewed as a challenging task as it involves the prediction of a complete graph comprising multiple nodes and edges based on a given label. This task, however, holds paramount importance in a wide array of real-world applications, such as de-novo molecule design [1]. In de-novo molecule generation, the chemical space is discrete by nature, and the entire search space is huge, which is estimated to be between 10^{23} and 10^{60} [2]. The generation of novel and valid molecular graphs with desired physical, chemical and biological property objectives should also be considered, these together with the large search space makes it a difficult task, since *these property objectives are highly complex and non-differentiable* [3].

Recently, there has been significant progress in molecular graph generation with Graph Neural Network (GNN)-based deep generative models, such as the autoregressive generative models: Graph Convolutional Policy Network (GCPN) [3] and Flow-based Autoregressive Model (GraphAF) [4] and the one-shot generative model: Graph Generation with Energy-based Model (GraphEBM) [5].

Those models all use the Relational Graph Convolutional Network (R-GCN) [6] as their inner graph representation model. Meanwhile, several researchers have focused on improving GNN expressiveness by introducing architectural changes [7–9], leading to the discovery of diverse forms of GNNs that excel in graph classification and regression tasks. Naturally, it is worth considering *will these more expressive GNNs help in molecular graph generation, and will they perform better than the de-facto network used in GCPN, GraphAF, and GraphEBM?*

In addition, nowadays, there are many metrics used in de-novo molecule design (e.g. molecule’s bioactivity against its corresponding disease targets: DRD2 and JNK3) [10]. However, GCPN, GraphAF, and GraphEBM only consider two molecular generative objectives related to drug design: the Penalised logP and the QED. This brings two major drawbacks. Firstly, there is a lack of consideration of a broader set of generation metrics beyond those related to drug design. Secondly, both Penalised logP and the QED are incapable of effectively distinguishing between various generation models, rendering them disabled for such purpose [10].

In this work, we replace R-GCN in GCPN, GraphAF, and GraphEBM with more expressive GNNs. Then, we evaluate our proposed models: GCPN variants, GraphAF variants, and GraphEBM variants on a wide array of molecular generative objectives (e.g. DRD2, Median1, Median2), which are important metrics for de-novo molecular design, to derive more statistically significant results. The main contributions of our work are as follows:

- Although GNN expressiveness defined by the Weisfeiler-Lehman (1-WL) graph isomorphism test works well on graph classification and regression tasks, it is not a necessary condition for a good GNN-based generative model. We observe empirically that the expressiveness of GNNs does not correlate well with their performance on GNN-based generative models, and GNNs incorporating edge feature extraction can improve GNN-based generative models.
- Although Penalised logP and QED are widely used in evaluating goal-directed graph generative models, they are not effective metrics to differentiate generative models. Other metrics, such as DRD2, Median1 and Median2, can better evaluate the ability of a graph generative model.
- Our findings reveal that while there is no direct correlation between expressiveness and performance in graph generation, substituting the inner GNNs of GCPN, GraphAF, and GraphEBM with advanced GNNs like GearNet yields better performance (e.g. 102.51% better in DRD2 and 48.96% better in Median2). By doing so, these models surpass or reach comparable performance to state-of-the-art non-graph-based generative methods for de-novo molecule generation.

2 Related Work

A variety of deep generative models have been proposed for molecular graph generation recently [3–5, 11–18]. In this paper, we confine our scope to single-objective approaches for molecular generation. Specifically, our focus centres on the generation of organic molecules that possess a desired scalar metric, encompassing key physical, chemical, and biological properties.

2.1 GNN-based Graph Generative Models

Recently, there has been significant progress in molecular graph generation with GNN-based deep generative models, such as the autoregressive generation model (GCPN [3] and GraphAF [4]) and the one-shot generation model (GraphEBM [5]). All of them use R-GCN [6], the state-of-the-art GNN at that time, as their inner graph representation model. Nevertheless, in recent years, the landscape has witnessed the emergence of increasingly expressive GNNs such as GATv2 [7], GSN [8] and GearNet [9]. These advanced GNN models have showcased superior performance in various tasks, including graph classification and regression, surpassing the capabilities of R-GCN. Furthermore, the methods GCPN, GraphAF, and GraphEBM only evaluate their performance in goal-directed molecule generation tasks using two commonly employed metrics in drug design: quantitative estimate of drug-likeness score (QED) and penalised octanol-water partition coefficient (Penalised logP). However, it is worth noting that many advanced graph generative approaches, as discussed in Section 2.2, can achieve upper-limit results on benchmarks for QED and Penalised logP [10, 16]. *QED is likely to have a global maximum of 0.948 and even random sampling could reach that value. Penalised logP is unbounded and the relationship between Penalised logP values and molecular structures is fairly simple: adding carbons monotonically increases the estimated Penalised logP value [3, 10, 19].*

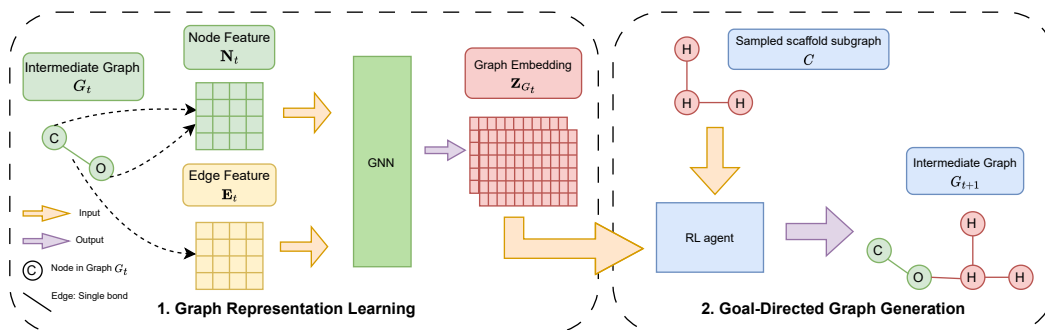


Figure 1: An overview of the GCPN model: this is an example of iterative graph generation from an intermediate graph G_t to an intermediate graph G_{t+1} . Part 1 is the illustration of the graph representation learning process based on a GNN. Part 2 is the illustration of the graph generative procedure based on a reinforcement learning (RL) agent. New nodes or edges are marked in red.

This is presented in Figure 4 in Appendix F. Consequently, one could argue that these scores have reached a saturation point, making them less meaningful as evaluation metrics. Only using these two metrics relevant to drug design on goal-directed graph generation tasks to assess the graph generative models is not convincing enough, and cannot provide insights for distinguishing different algorithms’ de-novo molecule generation ability [10].

2.2 Non-GNN-based Graph Generative Models

There exist also non-GNN-based graph generative models. Genetic Algorithms (GA) are generation approaches by relying on biologically inspired operators such as mutation, crossover and selection. Bayesian Optimisation (BO) [20] is an approach that uses a sequential optimisation technique that leverages probabilistic models to search for the optimal solution. Variational Autoencoders (VAEs) [21] is a type of generative model in machine learning that combines elements of both autoencoders and probabilistic latent variable models to learn and generate data by mapping it to a latent space with continuous distributions. Monte-Carlo Tree Search (MCTS) constructs a search tree by iteratively selecting actions, simulating possible outcomes, and propagating the results to inform future decisions, ultimately aiming to find the optimal solution. Hill Climbing (HC) is an iterative optimisation algorithm that starts with an arbitrary solution to a problem, and then attempts to find a better solution by making an incremental change to the solution. Reinforcement Learning (RL) learns how intelligent agents take actions in an environment to maximise the cumulative reward by transitioning through different states. Details about the non-GNN-based graph generative models we use as baselines in Section 4 can be found in Appendix A.

3 Method

In this section, we provide the theoretical background for graph generative models. A GNN-based graph generative model consists of a GNN model and a graph generative framework. The GNN learns the hidden representations of a graph, such as the node features and the graph features. The goal of the graph generation framework is to generate realistic molecular graph structures based on a given generative objective. The detail of GCPN is presented in Figure 1, and illustrations of the GraphAF and GraphEBM architectures are displayed in Figures 2 and 3 in Appendix E respectively.

3.1 Preliminaries

A graph is defined as $G = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} denotes the set of nodes, and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ denotes the set of edges. A relational graph can be expressed as $G = (\mathcal{V}, \mathcal{E}, \mathcal{R})$ where \mathcal{R} denotes the set of edge relations or edge types. For example, (i, j, r) means the edge from node i to node j with edge type r . A molecular graph can be represented by a tuple of features $(\mathbf{A}, \mathbf{H}, \mathbf{E}, \mathbf{R})$, where

- $\mathbf{A} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ is the adjacency matrix, with each entry a_{ij} representing an edge (if any) between nodes i and j ; note that this is different from the conventional $\{0, 1\}^{|\mathcal{V}| \times |\mathcal{V}|}$ adjacency matrix format, since there are different types of bonds (i.e., single, double, triple, aromatic).

- $\mathbf{H} \in \mathbb{R}^{|\mathcal{V}| \times d}$ is the feature matrix, $\mathbf{h}_i \in \mathbb{R}^d$ is the d -dimensional features of node i .
- $\mathbf{E} \in \mathbb{R}^{|\mathcal{E}| \times d_e}$ is the edge feature matrix, $\mathbf{e}_{ij} \in \mathbb{R}^{d_e}$ is the d_e -dimensional features of edge (i, j) .
- $\mathbf{R} \in \mathbb{N}^{|\mathcal{E}|}$ is a vector containing the edge types of each edge $(i, j) \in \mathcal{E}$ and $r_{(i,j)} \in \mathcal{R}$. This feature vector is explicitly used in R-GCN [6] and GearNet [9] (details in Appendix C).

The degree matrix $\mathbf{D} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ of a graph $G = (\mathcal{V}, \mathcal{E})$ is a diagonal matrix with each diagonal entry $d_{ii} = \text{deg}(v_i)$, which counts the number of edges terminating at a node in an undirected graph.

3.1.1 Subgraphs & Isomorphism

A graph $G' = (\mathcal{V}', \mathcal{E}')$ is a *subgraph* of a graph $G = (\mathcal{V}, \mathcal{E})$ (denoted $G' \subseteq G$) if and only if $\mathcal{V}' \subseteq \mathcal{V}$ and $\mathcal{E}' \subseteq \mathcal{E}$. Two graphs $G = (\mathcal{V}, \mathcal{E})$ and $G' = (\mathcal{V}', \mathcal{E}')$ are *isomorphic* (denoted $G \cong G'$) if and only if there exists an adjacency-preserving bijective mapping $f : \mathcal{V} \rightarrow \mathcal{V}'$, i.e.,

$$\forall i, j \in \mathcal{V}. (i, j) \in \mathcal{E} \iff (f(i), f(j)) \in \mathcal{E}' \quad (1)$$

An *automorphism* of a graph $G = (\mathcal{V}, \mathcal{E})$ is an isomorphism that maps G onto itself.

3.2 Graph Neural Networks

All the GNNs investigated in this paper can be abstracted as Message Passing Neural Networks (MPNNs). A general MPNN operation iteratively updates the node features $\mathbf{h}_i^{(l)} \in \mathbb{R}^d$ from layer l to layer $l + 1$ via propagating messages through neighbouring nodes $j \in \mathcal{N}_i$, which can be formalised as

$$\mathbf{h}_i^{(l+1)} = \text{UPDATE} \left(\mathbf{h}_i^{(l)}, \bigoplus_{j \in \mathcal{N}_i} \text{MESSAGE} \left(\mathbf{h}_i^{(l)}, \mathbf{h}_j^{(l)}, \mathbf{e}_{ij} \right) \right) \quad (2)$$

where MESSAGE and UPDATE are learnable functions, such as Multi-Layer Perceptrons (MLPs), $\mathcal{N}_i = \{j | (i, j) \in \mathcal{E}\}$ is the (1-hop) neighbourhood of node i , and \bigoplus is a permutation-invariant local neighbourhood aggregation function, such as sum, mean or max. After k iterations of aggregation, node i 's representation $\mathbf{h}_i^{(k)}$ can capture the structural information within its k -hop graph neighbourhood. Then, the graph embedding $\mathbf{h}_G \in \mathbb{R}^d$ can be obtained via a READOUT function:

$$\mathbf{h}_G = \text{READOUT}_{i \in \mathcal{V}} \left(\mathbf{h}_i^{(k)} \right) \quad (3)$$

which aggregates the node features to obtain the entire graph's representation \mathbf{h}_G .

Ideally, a maximally powerful GNN could distinguish different graph structures by mapping them to different representations in the embedding space. This ability to map any two different graphs to different embeddings, however, implies solving the challenging graph isomorphism problem. That is, we want isomorphic graphs to be mapped to the same representation and non-isomorphic ones to different representations. Thus, the expressiveness of a GNN is defined as the ability to distinguish non-isomorphic graphs, and can be analysed by comparing to the 1-WL graph isomorphism test, which are summarised in Appendix C.

3.3 Graph Generative Frameworks

In this paper, GCPN [3] and GraphAF [4] are used as autoregressive generative frameworks, and GraphEBM [5] is used as the one-shot generative framework for molecular graph generation tasks. GCPN and GraphAF formalise the problem of goal-directed graph generation as a sequential decision process through RL, i.e. the decisions are generated from the generation policy network. GraphEBM uses Langevin dynamics [22] to train the proposed energy function by approximately maximising likelihood and generate molecular graphs with low energies.

In GCPN, the iterative graph generation process is formulated as a general decision process: $M = (\mathcal{S}, \mathcal{A}, P, R, \gamma)$, where $\mathcal{S} = \{s_i\}$ is the set of states that consists of all possible intermediate and final graphs; $\mathcal{A} = \{a_i\}$ is the set of actions that describe the modification made to the current graph at each time step determined by a group of MLPs predicting a distribution of actions; P is the Markov transition dynamics that specifies the possible outcomes of carrying out an action, $p(s_{t+1} | s_t, \dots, s_0, a_t) = p(s_{t+1} | s_t, a_t)$; $R(s_t)$ is a reward function that specifies the reward after

reaching state s_t ; and γ is the discount factor. In each iteration, GCPN takes the intermediate graph G_t at time t and the collection of proposed scaffold subgraphs C as inputs, and outputs the action a_t , which predicts a new link to be added to G_t .

GraphAF defines an invertible transformation from a base distribution (e.g. multivariate Gaussian) to a molecular graph structure $G = (\mathcal{V}, \mathcal{E})$ as the generation policy network. Starting from an empty graph G_0 , in each step a new node v_i is generated based on the current sub-graph structure G_i , i.e., $p(v_i|G_i)$ by the policy network. Next, the edges between this new node and existing nodes are sequentially generated according to the current graph structure, i.e., $p(\mathcal{E}_{ij}|G_i, v_i, \mathcal{E}_{i,1:j-1}, \mathcal{E}_{1:j-1,i})$. This process is repeated until all the nodes and edges are generated.

In GraphEBM, in order to generate molecules with a specific desirable property at one time, it parameterises an energy function $E_\theta(G)$ by a GNN, assigning lower energies to data points that correspond to desirable molecular graphs and higher energies to other data points. After training the energy function network $E_\theta(G)$, it initialises a random sample G' and applies K steps of Langevin dynamics [22] to obtain a desirable data point G'' that has low energy.

4 Evaluation

4.1 Experimental Setup

Dataset. We use the ZINC-250k [23] dataset for both pre-training and fine-tuning proposed models for its pharmaceutical relevance, moderate size, and popularity. All molecules are presented in the kekulised form with hydrogen removed. ZINC-250k contains around 250,000 drug-like molecules with 9 atom types and 3 edge types, and with a maximum graph size of 38 sampled from the ZINC database. The molecules in ZINC are readily synthesisable molecules: it contains over 120 million purchasable “drug-like” compounds—effectively all organic molecules that are for sale—a quarter of which are available for immediate delivery. Other datasets, such as QM9 [24] (a subset of GDB9 [25]), contain, at least to some degree, virtual molecules that are likely to be synthesisable but have not been made yet, including many molecules with complex annulated ring systems [16]. In addition, the original works of GCPN, GraphAF, GraphEBM and the benchmark [10] we used to compare graph generative models have also been trained on ZINC-250k. Thus, ZINC-250k is well-suited for models to learn representations of drug-like and synthesisable molecules in de-novo molecule generation.

Implementation details. We use the open-source platform TorchDrug [26] and DIG [27] in dataset preparation and graph generative models training. Advanced GNNs are implemented in PyTorch [28] in the MPNN framework aligned with TorchDrug [26]. The basic architectures of GCPN and GraphAF are implemented in TorchDrug. Note that the original GraphAF cannot allow GNN module to aggregate edge features in the message-passing mechanism, since the intermediate molecular graph generated by the autoregressive flow doesn’t contain edge features. Therefore, we propose an improved version of GraphAF considering edge features, named GraphAF+e. To conduct a fairly analogous evaluation on GCPN, we considered GCPN without considering edge features, named GCPN-e. A detailed description of how to incorporate edge features is in Appendix D. The basic architectures of GraphEBM are implemented in DIG [27].

In the experiments, we replace R-GCN in GCPN (both with and without edge features), GraphAF (both with and without edge features), and GraphEBM with six more expressive GNNs: GIN [29], GAT [30], GATv2 [7], PNA [31], GSN [8] and GearNet [9]. We set all GNNs in the experiments to have 3 hidden layers with batch normalisation [32] and ReLU [33] activation applied after each layer. We find little improvement when further adding GCN layers. Each GNN uses a linear layer to transform the edge features in the graph to a hidden embedding, and concatenates with the node features for message passing. In addition, they use sum as the READOUT function for graph representations. For PNA and GSN, the MESSAGE and UPDATE functions are parameterised by linear layers. For GSN, we set the graph substructure set to contain cyclic graphs of sizes between 3 and 8 (both inclusive), which are some of the most important substructures in molecules [8]. For GAT and GATv2, we use multi-head attention with $k = 3$ after a manual grid search for attention heads.

The autoregressive generative models (GCPN and GraphAF), with all different GNNs, are pre-trained on the ZINC-250k dataset for 1 epoch, as we do not observe too much performance gain from increasing pre-training epochs. They are then fine-tuned towards the target properties with RL, using the proximal policy optimisation (PPO) algorithm [34]. The models are fine-tuned for 5 epochs for

all goal-directed generation tasks with early stopping if all generation results in one batch collapse to singleton molecules. We set the agent update interval for GCPN and GraphAF models to update their RL agents every 3 and 5 batches respectively, as considering the cost of computing resources. During generating process, we allow our RL agents in all models to make 20 resamplings on the intermediate generated graphs if they cannot generate chemically valid molecular graphs. The max graph size is set as 48 empirically. For GCPN, the collection of proposed scaffold subgraphs C for GCPN are sampled from non-overlapping scaffolds in the ZINC-250k dataset. For GraphAF, we define multivariate Gaussian as our base distribution and use node MLPs and edge MLPs which have two fully-connected layers equipped with tanh non-linearity to generate the nodes and edges respectively. We notice limited improvements in performance when increasing the number of MLP layers. Adam [35] is used as the optimiser for both pre-training and fine-tuning tasks. Reward temperature for each metric, the number of neurons in each hidden layer and the learning rate for fine-tuning) for each model on each task through grid search by Optuna [36] independently.

In the one-shot generative models (GraphEBM), with all different GNNs, we adopt an energy network of $L = 3$ layers with hidden dimension $d = 64$ for each GNN. For training, we tune the following hyperparameters: the sample step K of Langevin dynamics, the standard deviation of the added noise in Langevin dynamics σ and the step size λ . All models are trained for up to 20 epochs. In addition, we clip the gradient used in Langevin dynamics so that its value magnitude can be less than 0.01.

The experiments were run with a mix of NVIDIA A100 GPUs with 40GB memory and NVIDIA V100 GPUs with 16GB memory. The total amount of training time for all GCPN, GraphAF and GraphEBM variants under all metrics is around 1650 GPU hours. All details are provided in Appendix D.

Baselines. We compare our proposed models based on the original GCPN (both with and without edge features), GraphAF (both with and without edge features), and GraphEBM on six goal-directed molecule generation tasks: Penalised logP [37], QED [38], synthetic accessibility (SA) [39], DRD2 [1], Median1 and Median2. All generation metrics are taken from the Therapeutic Data Commons (TDC) [40]. Due to the practicality of de-novo molecule design, we only consider the single generation objective for all generation tasks. Specifically, SA stands for how hard or how easy it is to synthesise a given molecule. Penalised logP is a logP score that also accounts for ring size and SA, while QED is an indicator of drug-likeness. DRD2 is derived from a support vector machine (SVM) classifier with a Gaussian kernel fitting experimental data to predict the bioactivities against their corresponding disease targets. Median1 [40] measures the average score of the molecule’s Tanimoto similarity [41] to Camphor and Menthol. Median2 [40] measures the average score of the molecule’s Tanimoto similarity to Tadalafil and Sildenafil. Penalised logP has an unbounded range, while QED, DRD2, Median1, Median2 and SA have a range of $[0, 1]$ by definition. Higher scores in Penalised logP, QED, DRD2, Median1 and Median2 and a lower score in SA are desired. Note that all scores are calculated from empirical prediction models.

We choose to use DRD2, Median1 and Median2 to evaluate generative models, since they are mature and representative generative metrics [10, 16]. In addition, some other metrics are data-missing or inappropriate and thus cannot reflect the ability of generative models properly. For example, GSK3 cannot evaluate all the generated molecules; multi-property objectives (MPO) measure the geometric means of several scores, which will be 0 if one of the scores is 0; Valsartan Smarts is implemented incorrectly in TDC: it computes the geometric means of several scores instead of arithmetic means of several scores, which lead to incorrect results in the benchmark [10].

We compare the best GCPN variant, GraphAF variant, and GraphEBM variant with eight state-of-the-art approaches for molecule generation [10] on 6 generation metrics: Penalised logP, SA, DRD2, Median1, Median2 and QED. All results of baselines are taken from original papers unless stated.

4.2 Results

4.2.1 Improving GNN-based Graph Generative Methods

De-novo molecule design with more expressive GNNs. As we re-evaluate the Penalised logP and QED scores of the top-3 molecules found by GCPN, we note that our results turn out to be higher than the results reported in the original GCPN paper. We hypothesise this to be due to our more extensive hyperparameter searching. In Table 6 in Appendix E, we explore a set of simpler generation metrics based on the GCPN framework, such as Penalised logP, QED and SA, which have been widely used as objectives in previous work on GNN-based graph generative models.

Table 1: Comparison of the top-3 DRD2, Median1 and Median2 scores of the generated molecules by GCPN variants, with the top-3 property scores of molecules in the ZINC dataset for reference.

Model	DRD2			Median1			Median2		
	1st	2nd	3rd	1st	2nd	3rd	1st	2nd	3rd
ZINC	0.9872	0.9815	0.9773	0.3243	0.3096	0.3096	0.2913	0.2765	0.2749
R-GCN	0.4790	0.4790	0.4790	0.3367	0.3367	0.3242	0.1921	0.1891	0.1891
GIN	0.3460	0.3094	0.3094	0.3243	0.3243	0.3235	0.1770	0.1766	0.1730
GAT	0.4946	0.4946	0.4946	0.3367	0.3367	0.3328	0.1648	0.1640	0.1637
GATv2	0.5101	0.4946	0.4946	0.3367	0.3367	0.3331	0.1759	0.1720	0.1697
PNA	0.5828	0.4448	0.4448	0.3472	0.3254	0.3254	0.1629	0.1619	0.1605
GSN	0.5363	0.4946	0.4946	0.3243	0.3243	0.3235	0.1770	0.1766	0.1730
GearNet	0.9696	0.9684	0.9404	0.3367	0.3367	0.3367	0.2862	0.2794	0.2794

Table 2: Comparison of the top-1 DRD2, Median1, Median2 and QED scores with the selected non-GNN-based generative models. The full table can be found in Table 11 in Appendix E.

Model	DRD2	Median1	Median2	QED
GCPN (R-GCN)	0.479	0.337	0.192	0.948
GCPN (GearNet)	0.970 (+ 102.51%)	0.337 (+ 0.00%)	0.286 (+ 48.96%)	0.948 (+ 0.00%)
GraphAF (R-GCN)	0.928	0.281	0.143	0.946
GraphAF (GearNet)	0.987 (+ 6.36%)	0.290 (+ 3.20%)	0.183 (+ 27.97%)	0.947 (+ 0.11%)
GraphEBM (R-GCN)	0.691	0.281	0.148	0.948
GraphEBM (GearNet)	0.944 (+ 36.61%)	0.400 (+ 42.35%)	0.207 (+ 39.86%)	0.948 (+ 0.00%)
LSTM HC (SMILES) [16]	0.999	0.388	0.339	0.948
DoG-Gen [17]	0.999	0.322	0.297	0.948
GP BO [13]	0.999	0.345	0.337	0.947
SynNet [11]	0.999	0.244	0.259	0.948
GA+D [12]	0.836	0.219	0.161	0.945
VAE BO (SMILES) [14]	0.940	0.231	0.206	0.947
Graph MCTS [15]	0.586	0.242	0.148	0.928
MolDQN [18]	0.049	0.188	0.108	0.871

As summarised in Table 6, after replacing the inner R-GCN in GCPN with more expressive GNNs, we can observe a significant improvement of GCPN in Penalised logP: GCPN with GIN, GearNet and GSN can achieve the saturated 11.19 in Penalised logP. Moreover, with GearNet, the performance of the GCPN variants can outperform the original GCPN on all three metrics. However, we find that *these benchmarks are not suited to differentiate between different models*, since we see many GCPN variants with different GNNs achieve similar and saturated results on those metrics. In addition, these metrics are not representative enough to obtain meaningful conclusions, as discussed in Section 2.1. Therefore, *there is a need of better graph generation objectives*.

De-novo molecule design with better graph generation objectives. From Table 6, we notice that both Penalised logP and SA are saturating on advanced GNNs, making them ineffective metrics for distinguishing the capability of different GNN models, and we need better de-novo molecule generation metrics. Therefore, we introduce three more representative metrics: DRD2, Median1 and Median2, as described in Section 4.1, and report the top-3 property scores of molecules generated by each model trained on those three metrics in Table 1 (GCPN), Table 8 (GraphAF), and Table 10 (GraphEBM) in Appendix E. The results show that GCPN, GraphAF, and GraphEBM with more expressive GNNs, such as GearNet, can outperform the original GCPN, GraphAF and GraphEBM with R-GCN on all generation tasks by a significant margin. Specifically, on metrics such as DRD2 and Median2, GCPN, GraphAF, and GraphEBM with more expressive GNNs can vastly improve the original performance. This observation further indicates that by combining with more expressive GNNs, GCPN, GraphAF, and GraphEBM can successfully capture the distribution of desired molecules. Therefore, we suggest that DRD2, Median1, Median2 and QED are better graph generation metrics for differentiating different GNNs.

Table 3: Graph classification metrics among GCPN and GraphAF variants. NLL_{all} means the average negative log likelihood loss for all node and edge classification tasks. Acc means the average accuracy for all node and edge classification tasks. NLL_e and NLL_n mean the average negative log likelihood evaluated on classification tasks for edge and node respectively. By considering Acc , we derive a ranking of all GNNs in our practical setting in decreasing order of GNN expressiveness: GSN, PNA, GIN, GearNet, R-GCN, GATv2 and GAT. By considering NLL_e , we derive another ranking of the GNNs in decreasing order of one-hot encoding edge feature extraction ability: GearNet, R-GCN, PNA, GSN, GIN, GATv2, and GAT (More details in Appendix C).

Model	GCPN		GraphAF	
	NLL_{all}	Acc	NLL_e	NLL_n
GearNet	2.1265 ± 0.0017	0.8677 ± 0.0003	0.9275 ± 0.0048	2.8981 ± 0.0398
R-GCN	2.1427 ± 0.1326	0.8656 ± 0.0088	1.0674 ± 0.0204	3.3312 ± 0.1557
PNA	2.1172 ± 0.0036	0.8716 ± 0.0003	1.1043 ± 0.1148	3.2840 ± 0.2154
GSN	1.4219 ± 0.0026	0.9308 ± 0.0002	1.1319 ± 0.0678	3.1280 ± 0.0566
GIN	2.1500 ± 0.3158	0.8680 ± 0.0220	1.1691 ± 0.1167	3.6067 ± 0.2589
GATv2	2.8285 ± 0.2878	0.8013 ± 0.0244	1.2426 ± 0.0370	3.1325 ± 0.0681
GAT	2.8705 ± 0.2354	0.7957 ± 0.0211	1.2737 ± 0.0340	3.1541 ± 0.0844

Comparison with non-GNN-based graph generative methods. We report the top-1 DRD2, Median1, Median2 and QED scores found by all the GNN-based and non-GNN-based graph generative models in Table 2. As displayed in Table 2, original GCPN, GraphAF, and GraphEBM are not competitive among graph generative models on all the goal-directed molecule generation tasks. However, after modifying their inner GNNs with more advanced GNNs, such as GearNet, they can outperform or match state-of-the-art results across other generative approaches on the de-novo molecule generation task. Specifically, replacing R-GCN with GearNet can achieve an average of 50.49% improvement on GCPN, 12.51% on GraphAF, and 39.61% on GraphEBM in de-novo molecule design, with the proposed generation metrics.

Visualisations of the generated molecules with desired generative metrics by GCPN and GraphAF variants are presented in Figure 5 in Appendix F. In addition, we only report eight selected graph generative methods in the benchmark [10] in Table 2, and GCPN with GearNet can achieve comparable results across 17 other non-GNN-based graph generative methods on the proposed metrics, which are fully reported in Table 11 in Appendix E.

4.2.2 Correlation Between GNN Expressiveness and Graph Generation

In the pre-training phase, the GNNs are used to predict all node types and edge types in all masked graphs in the training data. We report the graph classification results for all GNNs in Table 3. In the GCPN framework, we can see GSN perform the best with 93.08% accuracy on the graph classification task. In the GraphAF framework, we can see GearNet perform the best with the lowest edge and node average negative log likelihood: 0.9275 and 2.8981 respectively. Both GSN and GearNet are more expressive GNNs than R-GCN demonstrated in Appendix C. It is not surprising that expressive GNNs can outperform other GNNs on the graph classification task, since the expressiveness of a GNN is defined as the ability to distinguish non-isomorphic graphs.

However, by comparing Table 3 with Tables 7, 1, 8, 9, it is worth noting that *more expressive GNNs cannot ensure better performance of GNN-based graph generative models in molecular generation tasks*. For example, PNA and GSN perform better than R-GCN on graph classification tasks (Table 3), but GCPN with PNA or GSN cannot surpass the original GCPN with R-GCN on all generation metrics (Table 1). As also demonstrated in Appendix C, we find graph generation models with strong one-hot encoding edge feature extraction ability GNNs, such as GearNet and R-GCN, can perform better than models with other GNNs, such as GATv2 and GAT. Details about GNN expressiveness and edge feature extraction ability are described in Appendix C. We conclude that the graph generation tasks requires other abilities of GNNs than graph prediction tasks, such as edge feature extraction.

De-novo molecule design with GNNs incorporating edge features. We investigate the performance of GCPN and GraphAF using GNNs with and without edge features. It is worth mentioning that the original GCPN considers edge features but GraphAF does not. The results of the GCPN without edge features (GCPN-e) are summarised in Table 7 in Appendix E, and those of the original GCPN are in Table 1. Both sets of results demonstrate that including edge features can significantly improve the top-3 scores on all three metrics for most GNNs. For instance, with the help of edge feature extraction, GCPN with GIN, GAT, GATv2, PNA and GSN can improve GCPN-e by 24.0%, 149.8%, 70.5%, 31.0% and 12.0% respectively on the top-1 score on the metric DRD2.

The original GraphAF comes without edge feature extraction and we improved it by incorporating the edge features. Table 8 (the original GraphAF) and Table 9 (GraphAF+e) in Appendix E summarised the results of both implementations accordingly. The results align with the GCPN case, where GraphAF+e improves the top-3 scores on all three metrics for most GNNs than the original GraphAF. With the help of edge feature extraction, GraphAF+e with GAT, GATv2, PNA and GSN can increase by 60.0%, 20.0%, 230.4% and 312.4% respectively on the top-1 score on the metric DRD2, compared with the original GraphAF with those GNNs.

The results above indicate the importance of *aggregating edge features for GNN-based generative models*. By harnessing the power to extract knowledge from edge information, the potential for generating more refined and accurate graphs is enhanced. We also notice that, when the GNNs used are either R-GCN or GearNet, we find GCPN and GraphAF perform well with and without considering edge information on generation metrics except DRD2, so we hypothesise the reason is that they absorb the number of edge relations as prior information in the model.

In summary, we conclude the following results:

1. More expressive GNN can lead to better results in the graph classification task. However, GNN expressiveness is not a necessary condition for a good GNN-based generative model. Generation tasks require other abilities of GNNs, such as edge feature extraction.
2. Although Penalised logP and QED are widely used for generative metrics in evaluating goal-directed graph generative models, they are not effective metrics to differentiate different generative models. Other metrics, such as DRD2, Median1 and Median2, can better evaluate the ability of a graph generative model. Under those metrics, we can see the performance of GCPN, GraphAF and GraphEBM can be enhanced by using more robust GNNs.
3. After applying advanced GNN to current GNN-based graph generative methods, such as GCPN, GraphAF and GraphEBM, they can outperform or match state-of-the-art results across 17 other generative approaches in the de-novo molecule generation task.

5 Limitation and Conclusion

Due to computation cost, we acknowledge several limitations of the current study: we cannot exhaustively explore every method, such as other Relational GNNs, and thoroughly tune every hyperparameter; we cannot evaluate all generative models on other complicated datasets besides ZINC-250k, such as ChEMBL [42], and other generation metrics [10]. However, our efforts have still provided valuable insights into investigating the expressiveness of GNN on the graph generation task, because of our focus on many different generative models and diverse generation objectives.

After exploring (1) unexplored underlying GNNs and (2) non-trivial generative objectives, we would like to conclude that expressiveness is not a necessary condition for a good GNN-based generative model. By evaluating GCPN variants, GraphAF variants, and GraphEBM variants on effective metrics, we demonstrate that GNN-based generative methods (GCPN, GraphAF and GraphEBM) can be improved by using more robust GNNs (e.g. strong edge features extraction and edge relation detection). With more robust GNNs, GNN-based graph generative models can outperform or match state-of-the-art results across 17 other generative approaches on de-novo molecule design tasks [10]. In the future, we plan to explore the necessary conditions for GNN to enhance the performance of GNN-based graph generative models.

Acknowledgement

This work was performed using the Sulis Tier 2 HPC platform hosted by the Scientific Computing Research Technology Platform at the University of Warwick, and the Cirrus UK National Tier-2 HPC at EPCC. Sulis is funded by EPSRC Grant EP/T022108/1 and the HPC Midlands+ consortium. Cirrus is funded by the University of Edinburgh and EPSRC Grant EP/P020267/1. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

References

- [1] Marcus Olivecrona, Thomas Blaschke, Ola Engkvist, and Hongming Chen. Molecular de-novo design through deep reinforcement learning. In *Journal of Cheminformatics*, 2017. 1, 6, 14, 24
- [2] Pavel G Polishchuk, Timur I Madzhidov, and Alexandre Varnek. Estimation of the size of drug-like chemical space based on GDB-17 data. In *Journal of Computer-Aided Molecular Design*, 2013. 1
- [3] Jiakuan You, Bowen Liu, Rex Ying, Vijay Pande, and Jure Leskovec. Graph Convolutional Policy Network for Goal-Directed Molecular Graph Generation. In *NeurIPS*, 2018. 1, 2, 4
- [4] Chence Shi, Minkai Xu, Zhaocheng Zhu, Weinan Zhang, Ming Zhang, and Jian Tang. GraphAF: a Flow-based Autoregressive Model for Molecular Graph Generation. In *ICLR*, 2020. 1, 2, 4, 20
- [5] Meng Liu, Keqiang Yan, Bora Oztekin, and Shuiwang Ji. GraphEBM: Molecular Graph Generation with Energy-Based Models. In *ICLR*, 2021. 1, 2, 4
- [6] Michael Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. Modeling Relational Data with Graph Convolutional Networks. In *The Semantic Web: 15th International Conference*, 2018. 2, 4, 15, 16, 18, 19
- [7] Shaked Brody, Uri Alon, and Eran Yahav. How Attentive are Graph Attention Networks? In *ICLR*, 2022. 2, 5, 17, 18
- [8] Giorgos Bouritsas, Fabrizio Frasca, Stefanos Zafeiriou, and Michael M. Bronstein. Improving Graph Neural Network Expressivity via Subgraph Isomorphism Counting. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 2, 5, 16, 19
- [9] Zuobai Zhang, Minghao Xu, Arian Jamasb, Vijil Chenthamarakshan, Aurelie Lozano, Payel Das, and Jian Tang. Protein Representation Learning by Geometric Structure Pretraining. In *ICLR*, 2023. 2, 4, 5, 16, 18
- [10] Wenhao Gao, Tianfan Fu, Jimeng Sun, and Connor W. Coley. Sample Efficiency Matters: A Benchmark for Practical Molecular Optimization. In *NeurIPS*, 2022. 2, 3, 5, 6, 8, 9, 14, 24
- [11] Wenhao Gao, Rocío Mercado, and Connor W Coley. Amortized tree generation for bottom-up synthesis planning and synthesizable molecular design. In *ICLR*, 2022. 2, 7, 13, 24
- [12] AkshatKumar Nigam, Pascal Friederich, Mario Krenn, and Alán Aspuru-Guzik. Augmenting genetic algorithms with deep neural networks for exploring the chemical space. In *ICLR*, 2020. 7, 13, 24
- [13] Austin Tripp, Gregor NC Simm, and José Miguel Hernández-Lobato. A fresh look at de novo molecular design benchmarks. In *NeurIPS 2021 AI for Science Workshop*, 2021. 7, 13, 24
- [14] Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Jorge Aguilera-Iparraguirre Dennis Sheberla, Timothy D Hirzel, Ryan P Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. In *ACS central science*, 2018. 7, 13, 24
- [15] Jan H Jensen. A graph-based genetic algorithm and generative model/monte carlo tree search for the exploration of chemical space. In *Chemical science*, 2019. 7, 13, 24
- [16] Nathan Brown, Marco Fiscato, Marwin H.S. Segler, and Alain C. Vaucher. GuacaMol: Benchmarking Models for de Novo Molecular Design. In *Journal of Chemical Information and Modeling*, 2019. 2, 5, 6, 7, 13, 14, 24

- [17] John Bradshaw, Brooks Paige, Matt J Kusner, Marwin Segler, and José Miguel Hernández-Lobato. Barking up the right tree: an approach to search over molecule synthesis dags. In *Advances in Neural Information Processing Systems*, 2020. 7, 13, 24
- [18] Zhenpeng Zhou, Steven Kearnes, Li Li, Richard N Zare, and Patrick Riley. Optimization of molecules via deep reinforcement learning. In *Scientific reports*, 2019. 2, 7, 13, 24
- [19] Tianfan Fu, Wenhao Gao, Cao Xiao, Jacob Yasonik, Connor W. Coley, and Jimeng Sun. Differentiable Scaffolding Tree for Molecular Optimization. In *ICLR, 2022*. 2, 24
- [20] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. A review of Bayesian optimization. In *Proceedings of the IEEE*, 2015. 3, 13
- [21] Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. In *ICLR, 2014*. 3, 13
- [22] Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *ICML*, 2011. 4, 5
- [23] Teague Sterling and John J Irwin. ZINC 15–ligand discovery for everyone. In *Journal of chemical information and modeling*, 2015. 5
- [24] Raghunathan Ramakrishnan, Pavlo O Dral, Matthias Rupp, and O Anatole von Lilienfeld. Quantum chemistry structures and properties of 134 kilo molecules. In *Scientific data*, 2014. 5
- [25] Lars Ruddigkeit, Ruud van Deursen, Lorenz C. Blum, and Jean-Louis Reymond. Enumeration of 166 billion organic small molecules in the chemical universe database GDB-17. In *Journal of Chemical Information and Modeling*, 2012. 5
- [26] Zhaocheng Zhu, Chence Shi, Zuobai Zhang, Shengchao Liu, Minghao Xu, Xinyu Yuan, Yangtian Zhang, Junkun Chen, Huiyu Cai, Jiarui Lu, Chang Ma, Runcheng Liu, Louis-Pascal Xhonneux, Meng Qu, and Jian Tang. TorchDrug: A Powerful and Flexible Machine Learning Platform for Drug Discovery. *arXiv preprint arXiv:2202.08320*, 2022. 5
- [27] Meng Liu, Youzhi Luo, Limei Wang, Yaochen Xie, Hao Yuan, Shurui Gui, Haiyang Yu, Zhao Xu, Jingtun Zhang, Yi Liu, Keqiang Yan, Haoran Liu, Cong Fu, Bora M Oztekin, Xuan Zhang, and Shuiwang Ji. DIG: A Turnkey Library for Diving into Graph Deep Learning Research. In *Journal of Machine Learning Research*, 2021. 5
- [28] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NIPS 2017 Workshop Autodiff*, 2017. 5
- [29] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How Powerful are Graph Neural Networks? In *ICLR*, 2019. 5, 15
- [30] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph Attention Networks. In *ICLR*, 2018. 5, 17, 19
- [31] Gabriele Corso, Luca Cavalleri, Dominique Beaini, Pietro Liò, and Petar Veličković. Principal Neighbourhood Aggregation for Graph Nets. In *NeurIPS*, 2020. 5, 15, 18, 19
- [32] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on Machine Learning*, 2015. 5
- [33] Abien Fred Agarap. Deep Learning using Rectified Linear Units (ReLU). *arXiv preprint arXiv:1803.08375*, 2018. 5
- [34] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal Policy Optimization Algorithms. *arXiv preprint arXiv:1707.06347*, 2017. 5
- [35] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference for Learning Representations*, 2015. 6
- [36] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A Next-generation Hyperparameter Optimization Framework. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery and data mining*, 2019. 6
- [37] Matt J. Kusner, Brooks Paige, and José Miguel Hernández-Lobato. Grammar Variational Autoencoder. In *Proceedings of the 34th International Conference on Machine Learning*, 2017. 6, 14

- [38] G. Richard Bickerton, Gaia V. Paolini, Jérémy Besnard, Sorel Muresan, and Andrew L. Hopkins. Quantifying the chemical beauty of drugs. In *Nature Chemistry*, 2012. 6, 14
- [39] Peter Ertl and Ansgar Schuffenhauer. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. In *Journal of Cheminformatics*, 2009. 6, 14
- [40] Kexin Huang, Tianfan Fu, Wenhao Gao, Yue Zhao, Yusuf Roohani, Jure Leskovec, Connor W Coley, Cao Xiao, Jimeng Sun, and Marinka Zitnik. Therapeutics data commons: Machine learning datasets and tasks for therapeutics. In *NeurIPS Track Datasets and Benchmarks*, 2021. 6, 14
- [41] Dávid Bajusz, Anita Rácz, and Károly Héberger. Why is Tanimoto index an appropriate choice for fingerprint-based similarity calculations? In *Journal of Cheminformatics*, 2015. 6
- [42] Anna Gaulton, Louisa J. Bellis, A. Patricia Bento, Jon Chambers, Mark Davies, Anne Hersey, Yvonne Light, Shaun McGlinchey, David Michalovich, Bissan Al-Lazikani, and John P. Overington. ChEMBL: a large-scale bioactivity database for drug discovery. In *Nucleic Acids Research*, 2012. 9
- [43] Thomas N. Kipf and Max Welling. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*, 2017. 15
- [44] Pan Li and Jure Leskovec. The Expressive Power of Graph Neural Networks. In *Graph Neural Networks: Foundations, Frontiers, and Applications*, 2022. 18, 19
- [45] George Cybenko. Approximation by superpositions of a sigmoidal function. In *Mathematics of control, signals and systems*, 1989. 18
- [46] Allan Pinkus. Approximation Theory of the MLP Model in Neural Networks. In *Acta Numerica*, 1999. 18
- [47] AkshatKumar Nigam, Robert Pollice, Mario Krenn, Gabriel dos Passos Gomes, and Alan Aspuru-Guzik. Beyond generative models: superfast traversal, optimization, novelty, exploration and discovery (STONED) algorithm for molecules using SELFIES. In *Chemical science*, 2021. 24
- [48] Yutong Xie, Chence Shi, Hao Zhou, Yuwei Yang, Weinan Zhang, Yong Yu, and Lei Li. MARS: Markov molecular sampling for multi-objective drug discovery. In *ICLR*, 2021. 24
- [49] David E Graff, Eugene I Shakhnovich, and Connor W Coley. Accelerating high-throughput virtual screening through molecular pool-based active learning. In *Chemical science*, 2021. 24
- [50] Tianfan Fu, Cao Xiao, Xinhao Li, Lucas M Glass, and Jimeng Sun. MIMOSA: Multi-constraint molecule sampling for molecule optimization. In *AAAI*, 2021. 24
- [51] Fredrik Svensson, Ulf Norinder, and Andreas Bender. Improving screening efficiency through iterative screening using docking and conformal prediction. In *Journal of chemical information and modeling*, 2017. 24
- [52] Yoshua Bengio, Tristan Deleu, Edward J. Hu, Salem Lahlou, Mo Tiwari, and Emmanuel Bengio. GFlowNet foundations. In *CoRR*, 2021. 24
- [53] Cynthia Shen, Mario Krenn, Sagi Eppel, and Alan Aspuru-Guzik. Deep molecular dreaming: Inverse machine learning for de-novo molecular design and interpretability with surjective representations. In *Machine Learning: Science and Technology*, 2021. 24
- [54] Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction tree variational autoencoder for molecular graph generation. In *ICML*, 2018. 24

A Appendix: Non-GNN-based Graph Generative Models

Genetic Algorithm (GA). The Genetic Algorithm (GA) is a widely used heuristic technique that draws inspiration from natural evolutionary mechanisms. It combines mutation and/or crossover perturbing a mating pool to enable exploration in the design space. SynNet [11] utilises a genetic algorithm on binary fingerprints and subsequently decodes them into synthetic pathways. GA+D [12] constitutes a genetic algorithm improved through the incorporation of a neural network (DNN) based discriminator model.

Bayesian Optimisation (BO) [20]. Bayesian Optimisation (BO) represents a broad category of techniques that constructs a surrogate model for the objective function through the application of Bayesian machine learning methods like Gaussian process (GP) regression. It then employs an acquisition function that integrates information from the surrogate model and its associated uncertainty to determine optimal sampling points. GPBO [13] optimises the GP acquisition function with Graph GA methods in an inner loop.

Variational Autoencoders (VAEs) [21]. Variational Autoencoders (VAEs) belong to a category of generative techniques that focus on maximising a lower bound of the likelihood, known as the evidence lower bound (ELBO), as opposed to directly estimating the likelihood itself. A VAE typically learns to map molecules to and from real space to enable the indirect optimisation of molecules by numerically optimising latent vectors. SMILES-VAE [14] uses a VAE to model molecules represented as SMILES strings.

Monte-Carlo Tree Search (MCTS). Monte-Carlo Tree Search (MCTS) conducts a localised and stochastic exploration of each branch originating from the present state, which could be a molecule or an incomplete molecule. It then identifies the most promising branches, which usually exhibit the highest property scores, to be considered for the subsequent iteration. Graph-MCTS [15] is an MCTS algorithm based on atom-level searching over molecular graphs.

Hill Climbing (HC). Hill Climbing (HC) is an iterative learning method that incorporates the generated high-scored molecules into the training data and fine-tunes the generative model for each iteration. SMILES-LSTM [16] leverages a LSTM to learn the molecular distribution represented in SMILES strings, and modifies it to a SELFIES version. DoG-Gen [17] instead learn the distribution of synthetic pathways as Directed Acyclic Graph (DAGs) with an RNN generator.

Reinforcement Learning (RL). In molecular design, a state is usually a partially generated molecule; actions are manipulations at the level of graph or string representations; rewards are determined by the desired properties of the molecules generated. MolDQN [18] uses a deep Q-network to generate molecular graphs in an atom-wise manner.

B Appendix: Details of Generation Metrics

The details about the de-novo molecular generation metrics we use [10, 40] are described as follows:

1. **Penalised LogP (PlogP):** It is a logP score that also accounts for ring size and synthetic accessibility (SA). The penalised logP score measures the solubility and synthetic accessibility of a compound [37]. It can range between zero (all properties unfavourable) and an unbounded limit (all properties favourable).
2. **Quantitative Estimate of Drug-likeness (QED):** The empirical rationale of the QED measure reflects the underlying distribution of molecular properties including molecular weight, logP, topological polar surface area, number of hydrogen bond donors and acceptors, the number of aromatic rings and rotatable bonds, and the presence of unwanted chemical functionalities [38]. It can range between zero (all properties unfavourable) and one (all properties favourable).
3. **Synthetic Accessibility (SA):** Synthetic Accessibility score stands for how hard or how easy it is to synthesise a given molecule, based on a combination of the molecule’s fragments contributions. The oracle is calculated via RDKit, using a set of chemical rules [39]. The method is based on the combination of molecule complexity and fragment contributions obtained by analyzing structures of a million already synthesised chemicals, and in this way captures also historical synthetic knowledge. It can range between zero (all properties favourable) and one (all properties unfavourable).
4. **DRD2:** it uses dopamine type 2 receptor as the biological target. The oracle is constructed by using a support vector machine classifier with a Gaussian kernel and ECFP6 fingerprints on the ExCAPE-DB dataset to measure a molecule’s biological activity against dopamine type 2 receptor [1]. It ranges between zero (all properties unfavourable) and one (all properties favourable).
5. **Median1:** it measures the average score of the molecule’s Tanimoto similarity to Camphor and Menthol [40]. In the median molecules benchmarks, the similarity to several molecules has to be maximised simultaneously. Besides measuring the obtained top score, it is instructive to study if the models also explore the chemical space between the target structures [16]. It can range between zero (all properties unfavourable) and one (all properties favourable).
6. **Median2:** it measures the average score of the molecule’s Tanimoto similarity to Tadalafil and sildenafil [40]. In the median molecules benchmarks, the similarity to several molecules has to be maximised simultaneously. Besides measuring the obtained top score, it is instructive to study if the models also explore the chemical space between the target structures [16]. It can range between zero (all properties unfavourable) and one (all properties favourable).

C Appendix: Graph Neural Networks

C.1 GNN Architectures

Relational Graph Convolutional Network (R-GCN). The graph convolution operation of the original GCN [43] can be defined as follows:

$$\mathbf{h}_i^{(l+1)} = \sigma \left(\sum_{j \in \mathcal{N}_i} c_{ij} \mathbf{W}^{(l)} \mathbf{h}_j^{(l)} \right) \quad (4)$$

where c_{ij} is a normalisation constant for each edge \mathcal{E}_{ij} which originates from using the symmetrically normalised adjacency matrix $\mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}$, $\mathbf{W}^{(l)}$ is a learnable weight matrix, and σ is a non-linear activation function. R-GCN [6] makes use of the relational data of the graphs, and extends the graph convolution operation to the following: let \mathcal{R} be the edge relation type (for molecular graphs, this can be the bond type), then

$$\mathbf{h}_i^{(l+1)} = \sigma \left(\sum_{r \in \mathcal{R}} \sum_{j \in \mathcal{N}_i^r} c_{i,r} \mathbf{W}_r^{(l)} \mathbf{h}_j^{(l)} + \mathbf{W}_0^{(l)} \mathbf{h}_i^{(l)} \right) \quad (5)$$

where \mathcal{N}_i^r denotes the set of neighbouring nodes of node i under relation $r \in \mathcal{R}$, $c_{i,r}$ is a problem-specific normalisation constant that can either be learnt or chosen in advance and $\mathbf{W}_r^{(l)}$ denotes the learnable matrix for edge type r . It has been shown that *neither GCN nor R-GCN is as expressive as the 1-WL test* [29].

Graph Isomorphism Network (GIN). Each GIN [29] layer updates the node features as follows:

$$\mathbf{h}_i^{(l+1)} = \phi^{(l)} \left(\left(1 + \epsilon^{(l)} \right) \mathbf{h}_i^{(l)} + \sum_{j \in \mathcal{N}_i} \mathbf{h}_j^{(l)} \right) \quad (6)$$

where $\phi^{(l)}$ is an MLP, and $\epsilon^{(l)}$ is a learnable scalar. *GIN is provably as expressive as the 1-WL test*, which makes it one of the maximally-expressive GNNs (proof in [29]).

Principal Neighbourhood Aggregation (PNA). The PNA [31] operator defines its aggregation function \oplus as a combination of neighbourhood-aggregators and degree-scalers, as defined by the following equation, with \otimes being the tensor product:

$$\oplus = \underbrace{\begin{bmatrix} \text{identity} \\ \text{amplification} \\ \text{attenuation} \end{bmatrix}}_{\text{scalers}} \otimes \underbrace{\begin{bmatrix} \text{mean} \\ \text{max} \\ \text{min} \\ \text{std} \end{bmatrix}}_{\text{aggregators}} \quad (7)$$

The PNA operator can then be inserted into the standard MPNN framework, obtaining the PNA layer:

$$\mathbf{h}_i^{(l+1)} = \phi^{(l)} \left(\mathbf{h}_i^{(l)}, \bigoplus_{j \in \mathcal{N}_i} \psi^{(l)} \left(\mathbf{h}_i^{(l)}, \mathbf{h}_j^{(l)}, \mathbf{e}_{ij} \right) \right) \quad (8)$$

where $\phi^{(l)}$ and $\psi^{(l)}$ are MLPs. According to the theorem that *in order to discriminate between multisets of size n whose underlying set is R , at least n aggregators are needed* (proof in [31]), PNA pushes its expressivity closer towards the 1-WL limit than GIN, by including more aggregators, thereby increasing the probability that at least one of the aggregators can distinguish different graphs.

Graph Substructure Network (GSN). GSN [8] adopts a feature-augmented message passing style by counting the appearance of certain graph substructures and encoding them into the features. The feature augmentation of GSN then works as follows: let $\mathcal{G} = \{G_1, \dots, G_K\}$ be a set of pre-computed small (connected) graphs. For each $G_k = (\mathcal{V}_k, \mathcal{E}_k)$ in \mathcal{G} , we first find its isomorphic subgraphs $G'_k = (\mathcal{V}'_k, \mathcal{E}'_k)$ in $G = (\mathcal{V}, \mathcal{E})$. Then, for each node $i \in V$ and $1 \leq k \leq K$, we count the number of subgraphs G'_k node i belongs to, as defined by the equation below:

$$x_{G_k}^{\mathcal{V}}(i) = |\{G'_k \cong G_k | i \in \mathcal{V}'_k\}| \quad (9)$$

We then obtain the node structural features for each node $i \in \mathcal{V}$: $\mathbf{x}_i^{\mathcal{V}} = [x_{G_1}^{\mathcal{V}}(i), \dots, x_{G_K}^{\mathcal{V}}(i)] \in \mathbb{N}^K$. Similarly, we can derive the edge structural features for each edge $(i, j) \in \mathcal{E}$: $\mathbf{x}_{ij}^{\mathcal{E}} = [x_{G_1}^{\mathcal{E}}(i, j), \dots, x_{G_K}^{\mathcal{E}}(i, j)] \in \mathbb{N}^K$ by counting the numbers of subgraphs it belongs to:

$$x_{G_k}^{\mathcal{E}}(i, j) = |\{G'_k \cong G_k | i \in \mathcal{E}'_k\}| \quad (10)$$

The augmented features can then be inserted into the messages and follow the standard MPNN, obtaining two variants of GSN layer, GSN-v (vertex-count) and GSN-e (edge-count):

$$\begin{aligned} \mathbf{h}_i^{(l+1)} &= \phi^{(l)} \left(\mathbf{h}_i^{(l)}, \bigoplus_{j \in \mathcal{N}_i} \psi^{(l)} \left(\mathbf{h}_i^{(l)}, \mathbf{h}_j^{(l)}, \mathbf{x}_i^{\mathcal{V}}, \mathbf{x}_j^{\mathcal{V}}, \mathbf{e}_{ij} \right) \right) \text{ (GSN-v)} \\ \mathbf{h}_i^{(l+1)} &= \phi^{(l)} \left(\mathbf{h}_i^{(l)}, \bigoplus_{j \in \mathcal{N}_i} \psi^{(l)} \left(\mathbf{h}_i^{(l)}, \mathbf{h}_j^{(l)}, \mathbf{x}_{ij}^{\mathcal{E}}, \mathbf{e}_{ij} \right) \right) \text{ (GSN-e)} \end{aligned} \quad (11)$$

where $\phi^{(l)}$ and $\psi^{(l)}$ are MLPs and \bigoplus is a permutation-invariant local neighbourhood aggregation function, such as sum, mean or max. It can be proved that *GSN is strictly more expressive than the 1-WL test, when G_k is any graph except for the star graphs (i.e., one center nodes connected to one of multiple outer nodes) of any size, and structural features are inferred by subgraph matching* (proof in [8]). This essentially suggests that GSN is more expressive than R-GCN which is at most as expressive as the 1-WL test in general.

Geometry Aware Relational Graph Neural Network (GearNet). GearNet [9] uses an R-GCN [6] as a fundamental framework to develop a node features and edge features message passing mechanism. Each GearNet layer updates the node features as follows:

$$\mathbf{h}_i^{(l+1)} = \sigma \left(\text{BN} \left(\sum_{r \in \mathcal{R}} \mathbf{W}_r^{(l)} \sum_{j \in \mathcal{N}_i^r} \mathbf{h}_j^{(l)} \right) \right) + \mathbf{h}_i^{(l)} \text{ (GearNet-v)} \quad (12)$$

where \mathcal{R} is the edge relation type, \mathcal{N}_i^r denotes the set of neighbouring nodes of node i under relation $r \in \mathcal{R}$, BN denotes a batch normalisation layer, $\mathbf{W}_r^{(l)}$ is the learnable convolutional kernel matrix for edge type r , and σ is a non-linear activation function.

To model the interactions between edges, we first construct a relational graph $G' = (\mathcal{V}', \mathcal{E}', \mathcal{R}')$ among edges. Each node in the graph G' corresponds to an edge in the original graph. G' links edge (i, j, r_1) in the original graph to edge (w, k, r_2) if and only if $j = w$ and $i \neq k$. The type of this edge is determined by the angle between (i, j, r_1) and (w, k, r_2) . *The angular information reflects the relative position between two edges that determines the strength of their interaction* [9]. Similar to R-GCN, the GearNet edge message passing layer is defined as:

$$\mathbf{e}_{i,j,r_1}^{(l+1)} = \sigma \left(\text{BN} \left(\sum_{r \in \mathcal{R}'} \mathbf{W}_r^{(l)} \sum_{(w,k,r_2) \in \mathcal{N}'_{rr}(i,j,r_1)} \mathbf{e}_{(i,j,r_1)}^{(l)} \right) \right) \text{ (GearNet-e)} \quad (13)$$

Similar as Eq. (12), the message function for edge (i, j, r_1) will be updated by aggregating features from its neighbours $\mathcal{N}''_{(i,j,r_1)}$, where $\mathcal{N}''_{(i,j,r_1)} = \{(w, k, r_2) \in \mathcal{V}' | ((w, k, r_2), (i, j, r_1), r) \in \mathcal{E}'\}$.

Finally, the entire GearNet message passing layer can be expressed as:

$$\mathbf{h}_i^{(l+1)} = \sigma \left(\text{BN} \left(\sum_{r \in \mathcal{R}} \mathbf{W}_r^{(l)} \sum_{j \in \mathcal{N}_i^r} \left(\mathbf{h}_j^{(l)} + \text{FC} \left(\mathbf{e}_{j,i,r_1}^{(l)} \right) \right) \right) \right) + \mathbf{h}_i^{(l)} \quad (14)$$

where FC denotes a linear transformation on the message function. GearNet is more expressive than R-GCN for its sparse edge message passing mechanism which encodes spatial information of a graph.

Graph Attention Network (GAT). In order to generalise the standard averaging or max-pooling aggregators in GNNs, GAT [30] applies attention-based neighbourhood aggregation as its aggregation function to obtain sufficient expressive power to transform the input features into higher-level features. The normalised masked attention coefficient for node i is defined as:

$$\forall j \in \mathcal{N}_i, \alpha_{ij} = \frac{\exp(\text{LeakyReLU}(\mathbf{a}[\mathbf{W}\mathbf{h}_i \parallel \mathbf{W}\mathbf{h}_j]))}{\sum_{k \in \mathcal{N}_i} \exp(\text{LeakyReLU}(\mathbf{a}[\mathbf{W}\mathbf{h}_i \parallel \mathbf{W}\mathbf{h}_k]))} \quad (15)$$

where \mathbf{a} is a learnable weight vector, representing the attention mechanism a : a single-layer feedforward neural network, \mathbf{W} is a learnable input linear transformation's weight matrix and \parallel represents concatenation operation.

Each K multi-head attention GAT layer updates the node features as follows:

$$\mathbf{h}_i^{(l+1)} = \parallel_{k=1}^K \sigma \left(\sum_{j \in \mathcal{N}_i} \alpha_{ij}^k \mathbf{W}^k \mathbf{h}_j^{(l)} \right) \quad (16)$$

where \parallel represents concatenation, α_{ij}^k are normalised attention coefficients computed by the k -th attention mechanism (a^k), and \mathbf{W}^k is the corresponding input linear transformation's weight matrix. GAT computes a representation for every node as a weighted average of its neighbours through the attention mechanism, which is more flexible than the neighbourhood aggregation in R-GCN.

Graph Attention Network v2 (GATv2). GATv2 [7] adopts a strictly more expressive *dynamic graph attention mechanism* [7] in its aggregation function to learn the graph representation. The graph attention variant that has a universal approximator attention function.

The normalised masked dynamic attention coefficient for node i is defined as:

$$\forall j \in \mathcal{N}_i, \alpha_{ij} = \frac{\exp(\mathbf{a}(\text{LeakyReLU}(\mathbf{W} \cdot [\mathbf{h}_i \parallel \mathbf{h}_j])))}{\sum_{k \in \mathcal{N}_i} \exp(\mathbf{a}(\text{LeakyReLU}(\mathbf{W} \cdot [\mathbf{h}_i \parallel \mathbf{h}_k])))} \quad (17)$$

where \mathbf{a} is a learnable weight vector, representing the attention mechanism a : a single-layer feedforward neural network, \mathbf{W} is a learnable input linear transformation's weight matrix and \parallel represents concatenation operation.

Each K multi-head attention GATv2 layer updates the node features as follows:

$$\mathbf{h}_i^{(l+1)} = \parallel_{k=1}^K \sigma \left(\sum_{j \in \mathcal{N}_i} \alpha_{ij}^k \mathbf{W}^k \mathbf{h}_j^{(l)} \right) \quad (18)$$

where \parallel represents concatenation, α_{ij}^k are normalised dynamic attention coefficients computed by the k -th attention mechanism (a^k), and \mathbf{W}^k is the corresponding input linear transformation's weight matrix. It has been proved that GATv2 – a graph attention variant that has a universal approximator attention function, and is thus strictly more expressive than GAT (proof in [7]).

C.2 GNN Expressiveness

Weisfeiler-Lehman (1-WL) Graph Isomorphism Test. Similar to GNNs, the 1-WL test iteratively updates the node embeddings of a graph by neighbourhood aggregation: for each node $i \in \mathcal{V}$ in a graph, an initial colour $c_i^{(0)}$ is assigned, and is iteratively updated using random hashes of sums:

$$c_i^{(t+1)} = \text{HASH} \left(\sum_{j \in \mathcal{N}_i} c_j^{(t)} \right) \quad (19)$$

The 1-WL test terminates when stable node colouring of the graph is reached, and outputs a histogram of colours. Two graphs with different colour histograms are non-isomorphic, and two graphs with the same colour histograms are possibly, but not necessarily, isomorphic. The neighbourhood aggregation in the 1-WL test can also be seen as a form of message passing, with GNNs being the learnable analogue. It has been proved that *message-passing GNNs are at most as expressive as the 1-WL test over discrete features* [44].

GNN Expressiveness in Graph Isomorphism Tests. We first illustrate a general graph representation learning framework [44] and then formulate GNN expressiveness defined by the 1-WL test in a graph representation learning framework. Finally, we explore the theoretical GNN expressiveness for all GNN models we proposed above and analyse the GNN expressiveness in our practical setting.

Definition 1 (Graph Representation Learning). The feature space is defined as $\mathcal{X} := \mathcal{A} \times \mathcal{B}$, where \mathcal{A} is the space of graph-structured data and \mathcal{B} includes all the node subsets of interest, given a graph, given a graph $G \in \mathcal{A}$. Then, a point in \mathcal{X} can be denoted as (G, S) , where S is a subset of nodes that are in G . Later, we call (G, S) as a graph representation learning (GRL) example. Each GRL example $(G, S) \in \mathcal{X}$ is associated with a target y in the target space \mathcal{Y} . Suppose the ground-truth association function between features and targets is denoted by $f^* : \mathcal{X} \rightarrow \mathcal{Y}$, $f^*(G, S) = y$. Given a set of training examples $\phi = \{(G^{(i)}, S^{(i)}, y^{(i)})\}_{i=1}^k$ and a set of testing examples $\psi = \{(\hat{G}^{(i)}, \hat{S}^{(i)}, \hat{y}^{(i)})\}_{i=1}^k$, a graph representation learning problem is to learn a GNN model f based on ϕ such that f is close to f^* on ψ .

Note that many frequently investigated learning problems can be formulated as graph representation learning problems by properly defining \mathcal{X} and \mathcal{Y} , such as graph classification and node classification problems [44]. Based on the graph representation learning framework and the 1-WL test, we can define GNN expressiveness [44] as below.

Definition 2 (GNN Expressiveness). Consider a feature space \mathcal{X} of a graph representation learning problem and a GNN model f defined on \mathcal{X} . Define another space $\mathcal{X}(f)$ as a subspace of the quotient space \mathcal{X}/\cong such that for two GRL examples $(G^{(1)}, S^{(1)}), (G^{(2)}, S^{(2)}) \in \mathcal{X}(f)$, $f(G^{(1)}, S^{(1)}) \neq f(G^{(2)}, S^{(2)})$. Then, the size of $\mathcal{X}(f)$ characterises the expressiveness of GNN model f . For two GNN models $f^{(1)}$ and $f^{(2)}$, if $\mathcal{X}(f^{(1)}) \supset \mathcal{X}(f^{(2)})$, then $f^{(1)}$ is more expressive than $f^{(2)}$.

Note that the expressive power of GNNs in Definition 2, characterised by how a model can distinguish non-isomorphic GRL examples, does not exactly match the traditional expressive power used for feed-forward neural networks in the sense of functional approximation [44]. Since the universal approximation theorem [45], current studies have proved that feed-forward neural networks can approximate any function of interest [46]. *However, these results have not been applied to GNNs due to the inductive bias imposed by additional constraints on the GNN parameter space* [44].

The expressiveness ability of each GNN model is discussed in a theoretical setting (with a sufficient and optimal number of GNN layers) above [6–9, 31, 44]. We can make the following conclusions:

1. R-GCN is theoretically the least expressive GNN among seven GNN models, since it is not as expressive as the 1-WL test [6]. Other GNN models in the message-passing mechanism are at most as powerful as the 1-WL test in distinguishing different graph-structured features [44].
2. GearNet is more expressive than R-GCN due to its sparse edge message-passing mechanism which encodes spatial information in a graph [9].
3. It has been proved that GATv2 – a graph attention variant that has a universal approximation attention function, and is thus strictly more expressive than GAT (proof in [7]).

4. GIN is provably as expressive as the 1-WL test [6]. PNA pushes its expressivity closer towards the 1-WL limit than GIN (proof in [31]).
5. It can be proved that GSN is strictly more expressive than the 1-WL test under certain constraints on the input graph space (proof in [8]).

Although there is a gap in comparing the GNN expressiveness among different GNN models based on current theoretical research, we can rank the GNN expressiveness based on our graph classification task. As shown in Table 3, we use the average accuracy for all node and edge classification tasks NLL_{all} to implicitly derive a ranking for all GNN models in our practical setting: GSN, PNA, GIN, GearNet, R-GCN, GATv2 and GAT (in decreasing order of GNN expressiveness), which corresponds to our theoretical analysis above. Thus, we can conclude that GSN, PNA and GearNet generally have stronger expressiveness and GIN, R-GCN, GAT and GATv2 have weaker expressiveness.

GNN Expressiveness in Edge Feature Extraction Ability. GNN expressiveness defined by the 1-WL test in Definition 2 is not a necessary condition for a good GNN-based generative model, since the GNN expressiveness is defined as the ability to distinguish non-isomorphic graphs without valuing the ability of GNN to distinguish the edge type in graphs [6, 44]. Actually, GNN expressiveness in Definition 2 is weak because distinguishing any non-isomorphic GRL examples does not necessarily indicate that we can approximate any function f defined over \mathcal{X} . Thus, in order to better compare different GNNs based on other abilities, we propose the general edge feature extraction ability of GNN models below.

Definition 3 (GNN Edge Feature Extraction Ability). Consider a graph representation learning problem formulated in Definition 1. In a GRL example (G, S) , S corresponds to a pair of nodes of interest. G for each example can be an induced subgraph around S or the entire graph. The target space \mathcal{Y} can be defined by low-level or high-level features related to edges in G . For example, the target space \mathcal{Y} can be defined as the edge type in S or the attention [30] between the edge type and the nodes in S . Each GRL example $(G, S) \in \mathcal{X}$ is associated with a target y in the target space \mathcal{Y} . Suppose the ground-truth association function between features and targets is denoted by $f^* : \mathcal{X} \rightarrow \mathcal{Y}$, $f^*(G, S) = y$. Given a set of training examples $\phi = \{(G^{(i)}, S^{(i)}, y^{(i)})\}_{i=1}^k$ and a set of testing examples $\psi = \{(\hat{G}^{(i)}, \hat{S}^{(i)}, \hat{y}^{(i)})\}_{i=1}^k$, a graph representation learning problem is to learn a GNN model f based on ϕ such that f is close to f^* on ψ .

Based on Definition 3, we consider a specific target space \mathcal{Y} and define the GNN one-hot encoding edge feature extraction ability below.

Definition 4 (GNN One-hot Encoding Edge Feature Extraction Ability). Consider the setting proposed on Definition 3 and define the target space \mathcal{Y} contains the one-hot encoding of the edge relation between two nodes. Given a set of training examples $\phi = \{(G^{(i)}, S^{(i)}, y^{(i)})\}_{i=1}^k$ and a set of testing examples $\psi = \{(\hat{G}^{(i)}, \hat{S}^{(i)}, \hat{y}^{(i)})\}_{i=1}^k$, consider a GNN model f learned based on ϕ , we define the space $\mathcal{C}(f)$ as a set $\{(\hat{G}^{(i)}, \hat{S}^{(i)}) | f(\hat{G}^{(i)}, \hat{S}^{(i)}) = \hat{y}^{(i)}\}_{i=1}^j$ consisted of correctly classified test examples. Then, the size of $\mathcal{C}(f)$ characterises the one-hot encoding edge feature extraction ability of the GNN model f . For two GNN models $f^{(1)}$ and $f^{(2)}$, if $\mathcal{C}(f^{(1)}) \supset \mathcal{C}(f^{(2)})$, we say that $f^{(1)}$ has stronger ability in edge feature extraction than $f^{(2)}$.

We can implicitly rank the one-hot encoding edge relation detection ability of GNN models based on the average negative log likelihood evaluated on classification tasks for the edge in Table 3 our paper: GearNet, R-GCN, PNA, GSN, GIN, GATv2, GAT (in decreasing order of one-hot encoding edge relation detection ability), which almost corresponds to the performance of their generative models shown in Table 1 and Table 9. Thus, we conclude that more expressive GNN can lead to better results in the graph classification task. However, GNN expressiveness defined by the 1-WL test is not a necessary condition for a good GNN-based generative model. Generation tasks require other abilities of GNNs, such as edge feature extraction and edge relation detection.

D Appendix: Experiment Details

Implementation of GNN incorporating edge features. As illustrated in Equation 2, a general MPNN can aggregate both node and edge embeddings to update the node embedding of the targeted node. In order to incorporate edge embeddings in GNNs, we use a single-layer MLP $f : \mathbb{R}^{d_e} \rightarrow \mathbb{R}^d$ to project relevant edge features of the targeted node to the same dimension of node features and then concatenate them to node features during the neighbourhood aggregation step in GNNs.

In GCPN-e, we use GNNs without incorporating edge features as the representational module in GCPN, i.e. there is no edge feature aggregated during the neighbourhood aggregation step in the message-passing mechanism. The general GNN without considering edge features iteratively updates the node features $\mathbf{h}_i^{(l)} \in \mathbb{R}^d$ from layer l to layer $l+1$ via propagating messages through neighbouring nodes $j \in \mathcal{N}_i$, which can be formalised by the following equation:

$$\mathbf{h}_i^{(l+1)} = \text{UPDATE} \left(\mathbf{h}_i^{(l)}, \bigoplus_{j \in \mathcal{N}_i} \text{MESSAGE} \left(\mathbf{h}_i^{(l)}, \mathbf{h}_j^{(l)} \right) \right) \quad (20)$$

where MESSAGE and UPDATE are learnable functions, such as Multi-Layer Perceptrons (MLPs), $\mathcal{N}_i = \{j | (i, j) \in \mathcal{E}\}$ is the (1-hop) neighbourhood of node i , and \bigoplus is a permutation-invariant local neighbourhood aggregation function, such as sum, mean or max.

In the original GraphAF, the intermediate molecular graph generated by the autoregressive flow doesn't contain edge features. Thus, in GraphAF+e, during the graph generation step, we apply a one-hot encoding to each edge according to the edge type to obtain edge features in the intermediate generated molecular graph.

Reward design implementation. For the property optimisation task, we use the same reward design in GraphAF, which incorporates both intermediate and final rewards for training the policy network. A small penalisation will be introduced as the intermediate reward if the edge predictions violate the valency check. The final rewards include both the score of targeted-properties of generated molecules and the chemical validity reward. The final reward is distributed to all intermediate steps with a discounting factor to stabilise the training [4]. The property-targeted reward r for a molecule m on a metric d is defined as follows:

$$r(m) = \exp \left(\frac{d(m)}{t} \right) \quad (21)$$

where t is the temperature for reward design decided by the grid search.

Hyper-parameter tuning. All the pre-training works are trained with an Adam optimiser with a learning rate of 0.001. For GCPN, we fixed batch sizes for pre-training and fine-tuning as 128 and 32. For GraphAF, we fixed batch sizes for pre-training and fine-tuning as 32 and 32. Each reward scale factor for Penalised logP, QED and SA is set as 1 after a manual grid search in the space $\{1, 5, 10\}$ based on evaluating each generative metric based on the performance of the original GCPN and GraphAF. The reward scale factors for DRD2, Median1 and Median2 are set as 0.5, 0.05 and 0.05 respectively after a manual grid search in the space $\{0.0001, 0.001, 0.05, 0.1, 0.5\}$ based on evaluating each generative metric based on the performance of original GCPN and GraphAF. We use Optuna to conduct a parallelised hyper-parameter grid search to determine the optimal hyper-parameters: the number of neurons in each hidden layer in the search space $\{64, 128, 256\}$ and the learning rate for fine-tuning in the search space $\{0.001, 0.0001, 0.00001, 0.000001\}$. The number of neurons in each hidden layer and the learning rate for fine-tuning for each generative task are summarised below in Table 4 and Table 5. For GraphEBM, we set the sample step of Langevin dynamics $K = 150$, the standard deviation of the added noise in Langevin dynamics $\sigma = 0.005$ and the step size $\lambda = 30$ based on hyperparameters of the original energy-based model.

Table 4: Detailed (batch size, learning rate) setup for GCPN and GraphAF variants on Penalised logP, QED and SA.

Model	Penalised logP	QED	SA
GCPN (R-GCN)	256, 0.00001	256, 0.00001	128, 0.001
GCPN (GIN)	256, 0.0001	256, 0.00001	256, 0.0001
GCPN (GAT)	128, 0.0001	256, 0.00001	64, 0.001
GCPN (GATv2)	256, 0.0001	256, 0.00001	256, 0.0001
GCPN (PNA)	256, 0.00001	64, 0.001	256, 0.0001
GCPN (GSN)	256, 0.0001	64, 0.001	256, 0.0001
GCPN (GearNet)	256, 0.0001	256, 0.00001	128, 0.001
GraphAF (R-GCN)	64, 0.0001	256, 0.00001	128, 0.0001
GraphAF (GearNet)	64, 0.0001	256, 0.000001	128, 0.0001

Table 5: Detailed (batch size, learning rate) setup for GCPN, GraphAF and GraphEBM variants on DRD2, Median1 and Median2.

Model	DRD2	Median1	Median2
GCPN (R-GCN)	256, 0.0001	128, 0.0001	256, 0.00001
GCPN (GIN)	256, 0.00001	256, 0.00001	256, 0.00001
GCPN (GAT)	256, 0.00001	64, 0.001	256, 0.00001
GCPN (GATv2)	128, 0.0001	128, 0.001	256, 0.00001
GCPN (PNA)	64, 0.001	256, 0.0001	256, 0.00001
GCPN (GSN)	256, 0.00001	64, 0.001	256, 0.00001
GCPN (GearNet)	256, 0.00001	256, 0.0001	256, 0.00001
GraphAF (R-GCN)	256, 0.000001	256, 0.00001	256, 0.000001
GraphAF (GIN)	256, 0.000001	256, 0.000001	256, 0.000001
GraphAF (GAT)	256, 0.000001	256, 0.00001	256, 0.00001
GraphAF (GATv2)	256, 0.000001	256, 0.000001	256, 0.000001
GraphAF (PNA)	256, 0.000001	256, 0.000001	256, 0.000001
GraphAF (GSN)	256, 0.0001	256, 0.000001	256, 0.000001
GraphAF (GearNet)	256, 0.00001	256, 0.00001	256, 0.000001
GraphEBM (All)	128, 0.0001	128, 0.0001	128, 0.0001

E Appendix: Additional Results




Figure 2: An overview of the GraphAF model, demonstrating an example of iterative graph generation from an intermediate graph G_t to an intermediate graph G_{t+1} . Part 1 is the illustration of the graph representation learning process based on a GNN. Part 2 is the illustration of the graph generative procedure based on a flow-based model. New nodes or edges are marked in red.




Figure 3: An overview of the GraphEBM model, demonstrating an example of the one-shot graph generation from an initialised node feature matrix \mathbf{N} and adjacency matrix \mathbf{A} to a generated molecular graph G^* with validity correction. It is the illustration of the graph generative procedure based on the Langevin dynamics with a trained energy function $E_{\theta^*}(\mathbf{N}, \mathbf{A})$ (parameterised by a GNN).

Table 6: Comparison of the top-3 Penalised logP, QED and SA scores of the generated molecules by GCPN variants, with the top-3 property scores of molecules in the ZINC dataset for reference.

Model	Penalised logP			QED			SA		
	1st	2nd	3rd	1st	2nd	3rd	1st	2nd	3rd
ZINC	4.52	4.30	4.23	0.948	0.948	0.948	1.0	1.0	1.0
R-GCN (baseline)	7.98	7.85	7.80	0.948	0.947	0.946	—	—	—
R-GCN (ours)	8.67	8.67	8.67	0.948	0.948	0.948	1.0	1.0	1.0
GIN	11.19	11.19	11.19	0.942	0.926	0.923	1.2	1.2	1.2
GAT	7.70	7.47	7.44	0.926	0.911	0.911	1.0	1.0	1.0
GATv2	8.08	7.48	7.34	0.945	0.908	0.907	1.0	1.0	1.0
PNA	8.66	8.61	8.22	0.833	0.825	0.754	1.0	1.0	1.0
GSN	11.19	11.19	11.19	0.804	0.783	0.783	1.0	1.0	1.0
GearNet	11.19	11.19	11.19	0.948	0.948	0.948	1.0	1.0	1.0

Table 7: Comparison of the top-3 DRD2, Median1 and Median2 scores of the generated molecules by GCPN-e variants, with the top-3 property scores of molecules in the ZINC dataset for reference.

Model	DRD2			Median1			Median2		
	1st	2nd	3rd	1st	2nd	3rd	1st	2nd	3rd
ZINC	0.9872	0.9815	0.9773	0.3243	0.3096	0.3096	0.2913	0.2765	0.2749
R-GCN	0.8315	0.7576	0.7551	0.3152	0.3152	0.3001	0.1932	0.1613	0.1592
GIN	0.2791	0.1980	0.1752	0.3152	0.3152	0.3152	0.1140	0.1113	0.1069
GAT	0.1980	0.1580	0.1580	0.3243	0.3243	0.3175	0.1196	0.1100	0.1098
GATv2	0.2992	0.2992	0.2992	0.3281	0.3281	0.3281	0.1042	0.1009	0.0911
PNA	0.4448	0.4448	0.4448	0.3243	0.3202	0.3175	0.0911	0.0764	0.0716
GSN	0.4790	0.4790	0.4448	0.3175	0.3175	0.3015	0.0982	0.0978	0.0897
GearNet	0.9990	0.9705	0.9574	0.3482	0.3482	0.3482	0.2084	0.2043	0.2037

Table 8: Comparison of the top-3 DRD2, Median1 and Median2 scores of the generated molecules by GraphAF variants, with the top-3 property scores of molecules in the ZINC dataset for reference.

Model	DRD2			Median1			Median2		
	1st	2nd	3rd	1st	2nd	3rd	1st	2nd	3rd
ZINC	0.9872	0.9815	0.9773	0.3243	0.3096	0.3096	0.2913	0.2765	0.2749
R-GCN	0.9277	0.9133	0.9080	0.2810	0.2641	0.2449	0.1426	0.1426	0.1417
GIN	0.5847	0.1835	0.1495	0.2651	0.2649	0.2393	0.1094	0.1042	0.1037
GAT	0.2992	0.2992	0.2992	0.2453	0.2212	0.2208	0.1031	0.1025	0.1023
GATv2	0.5909	0.4790	0.4790	0.2437	0.2335	0.2331	0.1239	0.1239	0.1232
PNA	0.1495	0.1411	0.1411	0.2897	0.2651	0.2651	0.1023	0.0764	0.0761
GSN	0.1238	0.0614	0.0601	0.2896	0.2773	0.2449	0.1025	0.1017	0.0977
GearNet	0.9872	0.9725	0.9714	0.2897	0.2896	0.2651	0.1826	0.1666	0.1616

Table 9: Comparison of the top-3 DRD2, Median1 and Median2 scores of the generated molecules by GraphAF+e variants, with the top-3 property scores of molecules in the ZINC dataset for reference.

Model	DRD2			Median1			Median2		
	1st	2nd	3rd	1st	2nd	3rd	1st	2nd	3rd
ZINC	0.9872	0.9815	0.9773	0.3243	0.3096	0.3096	0.2913	0.2765	0.2749
R-GCN	0.7833	0.7702	0.5714	0.2651	0.2667	0.2632	0.1526	0.1507	0.1481
GIN	0.8391	0.8016	0.7145	0.2735	0.2343	0.2304	0.1505	0.1496	0.1459
GAT	0.4790	0.4790	0.4790	0.2651	0.2509	0.2471	0.1358	0.1328	0.1152
GATv2	0.7087	0.5452	0.3986	0.2887	0.2342	0.2331	0.1362	0.1343	0.1316
PNA	0.4940	0.4940	0.4940	0.2633	0.2633	0.2449	0.1423	0.1313	0.1110
GSN	0.5106	0.4940	0.4940	0.2572	0.2530	0.2518	0.1448	0.1420	0.1090
GearNet	0.9699	0.9688	0.9623	0.2810	0.2582	0.2453	0.1829	0.1795	0.1672

Table 10: Comparison of the top-3 DRD2, Median1 and Median2 scores of the generated molecules by GraphEBM variants, with the top-3 property scores of molecules in the ZINC dataset for reference.


Model	DRD2			Median1			Median2		
	1st	2nd	3rd	1st	2nd	3rd	1st	2nd	3rd
ZINC	0.9872	0.9815	0.9773	0.3243	0.3096	0.3096	0.2913	0.2765	0.2749
R-GCN	0.6907	0.6907	0.6907	0.2810	0.2810	0.2810	0.1481	0.1481	0.1481
GIN	0.6907	0.6907	0.6907	0.2357	0.2357	0.2357	0.1420	0.1420	0.1420
GAT	0.6907	0.6907	0.6907	0.2897	0.2897	0.2897	0.1313	0.1313	0.1313
GATv2	0.8305	0.8305	0.8305	0.3243	0.3243	0.3243	0.1640	0.1640	0.1640
PNA	0.7652	0.7652	0.7652	0.2053	0.2053	0.2053	0.1343	0.1343	0.1343
GSN	0.7188	0.7188	0.7188	0.2810	0.2810	0.2810	0.1672	0.1672	0.1672
GearNet	0.9443	0.9443	0.9443	0.3999	0.3999	0.3999	0.2067	0.2067	0.2067

Table 11: Comparison of the top-1 DRD2, Median1, Median2 and QED scores between the experimented GNN-based generative model variants (GCPN, GraphAF, and GraphEBM) and all non-GNN-based graph generative models in the benchmark [10].


Model	DRD2	Median1	Median2	QED
GCPN (R-GCN)	0.479	0.337	0.192	0.948
GCPN (GearNet)	0.970 (+ 102.51%)	0.337 (+ 0.00%)	0.286 (+ 48.96%)	0.948 (+ 0.00%)
GraphAF (R-GCN)	0.928	0.281	0.143	0.946
GraphAF (GearNet)	0.987 (+ 6.36%)	0.290 (+ 3.20%)	0.183 (+ 27.97%)	0.947 (+ 0.11%)
GraphEBM (R-GCN)	0.691	0.281	0.148	0.948
GraphEBM (GearNet)	0.944 (+ 36.61%)	0.400 (+ 42.35%)	0.207 (+ 39.86%)	0.948 (+ 0.00%)
REINVENT (SMILES) [1]	0.999	0.399	0.332	0.948
LSTM HC (SMILES) [16]	0.999	0.388	0.339	0.948
Graph GA [15]	0.999	0.350	0.324	0.948
REINVENT (SELFIES) [1]	0.999	0.399	0.313	0.948
DoG-Gen [17]	0.999	0.322	0.297	0.948
GP BO [13]	0.999	0.345	0.337	0.947
STONED [47]	0.997	0.295	0.265	0.947
LSTM HC (SELFIES) [16]	0.999	0.362	0.274	0.948
DST [19]	0.999	0.281	0.201	0.947
SMILES GA [16]	0.986	0.207	0.210	0.948
SynNet [11]	0.999	0.244	0.259	0.948
MARS [48]	0.994	0.233	0.203	0.946
MolPal [49]	0.964	0.309	0.273	0.948
MIMOSA [50]	0.993	0.296	0.238	0.947
GA+D [12]	0.836	0.219	0.161	0.945
VAE BO (SELFIES) [14]	0.940	0.231	0.206	0.947
DoG-AE [17]	0.999	0.203	0.201	0.944
Screening [51]	0.949	0.271	0.244	0.947
GFlowNet [52]	0.951	0.237	0.198	0.945
VAE BO (SMILES) [14]	0.940	0.231	0.206	0.947
Pasithea [53]	0.592	0.216	0.194	0.943
JT-VAE BO [54]	0.778	0.212	0.192	0.946
GFlowNet-AL [52]	0.863	0.229	0.191	0.944
Graph MCTS [15]	0.586	0.242	0.148	0.928
MolDQN [18]	0.049	0.188	0.108	0.871

F Appendix: Visualisation of Generated Molecule Graphs

We present visualisations of generated molecules with the highest score on metric Penalised logP and QED. The visualisations illustrate that generated molecules with the highest Penalised logP score only contain a long chain of carbons.



(a) Generated molecules with the highest Penalised logP scores



(b) Generated molecules with the highest QED scores

Figure 4: Molecules with highest generation metrics: Penalised logP and QED generated by GNN-based graph generative models on de-novo molecule design tasks.

We present visualisations of generated molecules by GCPN, GraphAF and their optimal variants with GearNet in Figure 5(a), Figure 5(b), Figure 5(c) and Figure 5(d) respectively. The visualisations demonstrate that GCPN and GraphAF with advanced GNN have strong abilities to model different graph structures on the de-novo molecule design task.




Figure 5: Molecules with highest generation metrics: DRD2, Median1, Median2 and QED generated by proposed GNN-based graph generative models: (a) GCPN with R-GCN (b) GraphAF with R-GCN (c) GCPN with GearNet (d) GraphAF with GearNet on de-novo molecule design tasks.