ON EVALUATION METRICS FOR GRAPH GENERATIVE MODELS

Anonymous authors
Paper under double-blind review

ABSTRACT

In image generation, generative models can be evaluated naturally by visually inspecting model outputs. However, this is not always the case for graph generative models (GGMs), making their evaluation challenging. The currently popular GGM evaluation metrics are based on calculating the distance between empirical distributions of various graph statistics. However, these metrics suffer from three critical limitations: i) they do not produce a single score which makes model selection challenging, ii) they ignore edge and node features, and iii) they can be prohibitively slow to compute. In this work, we mitigate these issues by finding *a scalar, domain-agnostic, and scalable metric* for evaluating and ranking GGMs. To this end, we study existing GGM metrics and neural-network-based metrics emerging from generative models of images that use embeddings extracted from a task-specific network. Motivated by the power of certain Graph Neural Networks (GNNs) to extract meaningful graph representations *without any training*, we introduce several metrics based on the features extracted by an untrained random GNN. We design experiments to thoroughly test metrics on their ability to measure the diversity and fidelity of generated graphs, as well as their sample and computational efficiency. Our results on several synthetic and real graph datasets demonstrate that embeddings obtained from a random GNN paired with Maximum Mean Discrepancy (MMD) satisfies the aforementioned properties of a strong metric and performs competitively.

1 INTRODUCTION

Graph generation is a key problem in a wide range of domains such as molecule generation (Samanta et al., 2020; Popova et al., 2019; Li et al., 2018; Kong et al., 2021; Jin et al., 2020) and structure generation (Bapst et al., 2019; Thompson et al., 2020). An evaluation metric that is capable of accurately measuring the quality of generated graphs is critical for advancing research on graph generative models (GGMs). The current "gold standard" of evaluation metrics developed by You et al. (2018) use Maximum Mean Discrepancy (MMD) (Gretton et al., 2006) to compare empirical distributions of graph statistics such as orbit counts, degree coefficients, and clustering coefficients. While these metrics are capable of making a meaningful comparison between generated and real graphs (You et al., 2018), this evaluation method yields a metric for each statistic. *The ranking of generative models* may vary between metrics, making it challenging to measure progress. In addition, the metrics from You et al. (2018) are based only on graph structure and do not incorporate edge and node features (or attributes). Therefore, they are less applicable in specific domains such as molecule generation where such features are essential. Finally, the computation of these classical metrics can be prohibitively slow (Liao et al., 2019; O'Bray et al., 2021). A robust and scalable scalar metric is needed to accurately rank GGMs across domains.

While less popular, metrics from image generation literature have been successfully utilized in GGM evaluation. These metrics rely on the use of a task-specific neural network to extract meaningful representations of samples, enabling a more straightforward comparison between generated and reference distributions (Preuer et al., 2018; Liu et al., 2019; Thompson et al., 2020). Although these metrics have been validated empirically in the image domain, they are not universally applicable to GGMs. For example, Fréchet Chemnet Distance (Preuer et al., 2018) uses a language model trained on SMILE strings, rendering it unusable for evaluation of GGMs in other domains. Furthermore, a pretrained GNN cannot be applied to datasets with a different number of edge or node labels.

Pretraining a GNN for every dataset can be prohibitive, making the use of such metrics in GGM evaluation less appealing than in the more established and standardized image domain.

In image generation evaluation, classifiers trained on ImageNet (Deng et al., 2009) are frequently used to extract image embeddings (Bińkowski et al., 2018; Heusel et al., 2017; Kynkäänniemi et al., 2019; Xu et al., 2018; Naeem et al., 2020). While classifiers such as Inception v3 (Szegedy et al., 2016) are consistently used, recent works have investigated the use of randomly-initialized CNNs with no further training (hereafter referred to as *a random network*) in generative model evaluation. Xu et al. (2018); Naeem et al. (2020) found that a random CNN performs similarly to ImageNet classifiers on natural images and is superior outside of the natural image domain. In the graph domain, random GNNs have been shown to extract meaningful features to solve downstream graph tasks without training (Kipf & Welling, 2017; Morris et al., 2019; Xu et al., 2019). However, the applicability of random GNNs for the evaluation of GGMs remains unexplored.

In this work, we explore the use of random GNNs in the evaluation of GGMs using metrics that were developed in the image domain. We tackle the problem of assessing a large number of possible evaluation metrics. We design experiments to thoroughly test each metric on its ability to measure the diversity and fidelity (realism) of generated graphs, as well as their sample and computational efficiency. We study three families of metrics: existing GGM evaluation metrics, which we call *classical metrics* (You et al., 2018); image domain metrics using a random GNN; and image domain metrics using a pretrained GNN. We aim to answer the following questions empirically: ***(Q1)*** *What are the strengths and limitations of each metric?* ***(Q2)*** *Is pretraining a GNN necessary to accurately evaluate GGMs with image domain metrics?* ***(Q3)*** *Is there a strong scalar and domain-agnostic metric for evaluating and ranking GGMs?* Addressing these questions enabled us to reveal several surprising findings that directly impact the future of GGM evaluation. For example, regarding Q1, we identify a failure mode in the classical metrics in that they are poor at measuring the diversity of generated graphs. Consequently, we find several metrics that are more expressive. In terms of Q2, we determine that pretraining is unnecessary to utilize neural-network-based (NN-based) metrics. Regarding Q3, we find two scalar metrics that are ideal for evaluating and ranking GGMs in certain scenarios; they are scalable, powerful, and easily incorporate node and edge features. These findings enable computationally inexpensive and domain-agnostic GGM evaluation. All code will be open-sourced upon publication to provide researchers with off-the-shelf tools to evaluate both GGMs and GGM evaluation metrics.

2 Background & Related Work

Evaluating generative models in any domain is a notoriously difficult task (Theis et al., 2016). However, researchers have found success through the use of sample-based evaluation metrics that estimate the distance ρ between real and generated distributions P_r and P_g by drawing random samples (Heusel et al., 2017; You et al., 2018; Bińkowski et al., 2018). That is, they compute $\hat{\rho}(\mathbb{S}_g, \mathbb{S}_r) \approx \rho(P_g, P_r)$, with $\mathbb{S}_r = \{\mathbf{x}_1^r, \ldots, \mathbf{x}_m^r\} \sim P_r$ and $\mathbb{S}_g = \{\mathbf{x}_1^g, \ldots, \mathbf{x}_n^g\} \sim P_g$, where $\mathbf{x}_i$ is defined as some feature vector extracted from a corresponding graph G_i. These metrics are model agnostic and therefore applicable to all generative models.

2.1 Classical metrics

Metrics based on graph statistics (You et al., 2018) are standard in evaluating GGMs (Liao et al., 2019; Dai et al., 2020). These metrics require extracting the clustering coefficient, node degree, and 4-node orbit count histograms[1] that are then used to compute the empirical MMD between generated and reference sets $\mathbb{S}_g, \mathbb{S}_r$ (Gretton et al., 2006):

$$\mathrm{MMD}(\mathbb{S}_g, \mathbb{S}_r) := \frac{1}{m^2} \sum_{i,j=1}^{m} k(\mathbf{x}_i^r, \mathbf{x}_j^r) + \frac{1}{n^2} \sum_{i,j=1}^{n} k(\mathbf{x}_i^g, \mathbf{x}_j^g) - \frac{2}{nm} \sum_{i=1}^{n} \sum_{j=1}^{m} k(\mathbf{x}_i^g, \mathbf{x}_j^r), \qquad (1)$$

where $k(\cdot, \cdot)$ is the Earth Mover's Distance kernel function. This yields three metrics, one for each graph statistic. The computational cost of these metrics may be decreased by replacing the Earth

[1] While any set of graph statistics can be compared using MMD, these three are the most common and hence are evaluated in this work.

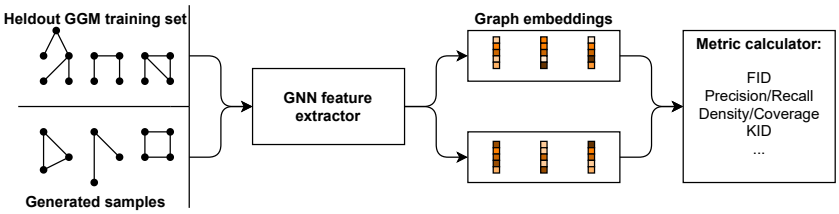

Figure 1: The standard process of evaluating GGMs using NN-based metrics.

Mover's Distance kernel with the total variance kernel (Liao et al., 2019). Nonetheless, this change leads to an indefinite kernel and undefined behaviour when used with MMD (O'Bray et al., 2021). Therefore, we only compute these metrics using EMD (You et al., 2018).

2.2 Graph neural networks

We denote a graph as $G = (\mathbb{V}, E)$ with n vertices $\mathbb{V}$ and edges $E = \{(i,j) \mid i,j \in \{1,\dots,n\}\}$. GNNs allow the extraction of a fixed size representation $\mathbf{x}_i$ of an arbitrary graph G_i. While many GNN formulations exist (Wu et al., 2020), we consider Graph Isomorphism Networks (GINs) (Xu et al., 2019) as a common GNN in practice. GINs consist of L propagation layers followed by a graph readout layer to obtain $\mathbf{x}_i$. For node $v \in \mathbb{V}$, the node embeddings $\mathbf{h}_v^{(l)}$ at layer $l \in [1, L]$ are computed as:

$$\mathbf{h}_v^{(l)} = \text{MLP}^{(l)}\Big(\mathbf{h}_v^{(l-1)} + f^{(l)}\big(\{\mathbf{h}_u^{(l-1)} : u \in \mathcal{N}(v)\}\big)\Big), \tag{2}$$

where $\mathbf{h}_v^{(0)}$ is the input feature of node v, $\mathbf{h}_v^{(l)} \in \mathbb{R}^d \;\; \forall\, l > 0$ denotes a d-dimensional embedding of node v after the l-th graph propagation layer; $\mathcal{N}(v)$ denotes the neighbors of node v; $\text{MLP}^{(l)}$ is a fully-connected neural network; $f^{(l)}$ is some aggregation function over nodes such as mean, max or sum. A graph readout layer with skip connections aggregates features from all nodes at each layer $l \in [1, L]$ and concatenates them into a single $L \cdot d$ dimensional vector $\mathbf{x}_i$ (Xu et al., 2019):

$$\mathbf{x}_i = \text{CONCAT}\Big(\text{READOUT}\big(\{\mathbf{h}_v^{(l)} \mid v \in \mathbb{V}\}\big) \mid l \in 1, 2, ..., L\Big), \tag{3}$$

where READOUT is similar to $f^{(l)}$ and is often chosen as the mean, max, or sum operation.

2.3 Neural-network-based metrics

NN-based metrics utilize a task-specific network to extract meaningful multidimensional embeddings of the input data. In image generation evaluation, the features of a layer in Inception v3 (Szegedy et al., 2016) are frequently used (Heusel et al., 2017). These NN-based metrics can be applied to GGM evaluation by replacing Inception v3 with a GNN (Thompson et al., 2020; Liu et al., 2019). This prompts the use of a wide range of evaluation metrics that have been studied extensively in the image domain. Computation of these metrics follow a common setup and differ only in how the distance between two sets of data are determined (Figure 1).

Fréchet Distance (FD) (Heusel et al., 2017) is one of the most popular generative metrics. FD approximates the graph embeddings as continuous multivariate Gaussians with sample mean and covariance $\mu, \boldsymbol{C}$. The distance between distributions is computed as an approximate measure of sample quality: $\text{FD}(\mathbb{S}_r, \mathbb{S}_g) = \|\mu_r - \mu_g\|_2^2 + \text{Tr}(\boldsymbol{C}_r + \boldsymbol{C}_g - 2(\boldsymbol{C}_r\boldsymbol{C}_g)^{1/2})$.

Improved Precision & Recall (P&R) (Kynkäänniemi et al., 2019) decouples the quality of a generator into two separate values to aid in the detection of *mode collapse* and *mode dropping*. Mode dropping refers to the case wherein modes of P_r are underrepresented by P_g, while mode collapse describes a lack of diversity within the modes of P_g. Manifolds are constructed by extending a radius from each embedded sample in a set to its kth nearest neighbour to form a hypersphere, with the union of all hyperspheres representing a manifold. Precision is the percentage of generated samples that fall within the manifold of real samples, while Recall is the percentage of real samples that fall within the manifold of generated samples. The harmonic mean ("F1 PR") of P&R is a scalar metric that can be decomposed into more meaningful values in experiments (Lucic et al., 2018).

Density & Coverage (D&C) (Naeem et al., 2020) have recently been introduced as robust alternatives for Precision and Recall, respectively. Density is calculated as the average number of real examples within whose manifold radius each generated sample falls. Coverage is described as the percentage of real examples that have a generated sample fall within their manifold radius. The manifolds used are found using the kth nearest neighbour as in P&R. As with P&R, the harmonic mean ("F1 DC") of D&C can be used to create a scalar metric (Lucic et al., 2018).

MMD (Gretton et al., 2006) (Equation 1) can also be used to measure the dissimilarity between graph embedding distributions. The original Kernel Inception Distance (KID) (Bińkowski et al., 2018) proposed a polynomial kernel with MMD, $k(\mathbf{x}_i, \mathbf{x}_j) = (\frac{1}{d}\mathbf{x}_i^\top \mathbf{x}_j + 1)^3$, where d is the embedding dimension. The linear kernel $k(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{x}_i \cdot \mathbf{x}_j$ is another parameter-free kernel used with MMD to evaluate generative models (O'Bray et al., 2021). In addition, the RBF kernel is widely used (Xu et al., 2018; Gretton et al., 2006) and is calculated as: $k(\mathbf{x}_i, \mathbf{x}_j) = \exp\left(\frac{-\|\mathbf{x}_i - \mathbf{x}_j\|^2}{2\sigma^2}\right)$. The choice of σ has a significant impact on the output of the kernel, and methods for finding and selecting an optimal value is an important area of research (Bach et al., 2004; Gretton et al., 2012a;b). A common strategy is to select σ as the median pairwise distance between the two comparison sets (Garreau et al., 2018; Gretton et al., 2006; Schölkopf et al., 1998). Another strategy is to evaluate MMD using a small set of σ values and select the value that maximizes MMD: $\sigma = \arg\max_{\sigma \in \Sigma} \mathrm{MMD}(\mathbb{S}_g, \mathbb{S}_r; \sigma)$ (Sriperumbudur et al., 2009). We combine these ideas in our σ selection process, and more details are available in Appendix A.1. We refer to these metrics as "KD," "MMD Linear," and "MMD RBF."

2.4 Benchmarking evaluation metrics

Our work is closely related to other works benchmarking evaluation metrics of generative models (O'Bray et al., 2021; Xu et al., 2018). A consistent method for assessing evaluation metrics is to start with the creation of "reference" and "generated" sets $\mathbb{S}_r$ and $\mathbb{S}_g$ that originate from the same real distribution P_r. Then, one gradually increases distortion to the generated set while recording each metric's response to the distortion. In the image domain, Xu et al. (2018) develop multiple experiments to test metrics for key properties. The metrics are then subjectively evaluated for these properties by manually analyzing each metric's response across experiments. In concurrent work, O'Bray et al. (2021) introduced additional GGM evaluation metrics. Similar to Xu et al. (2018), the metrics are validated through experiments that slowly increase the level of perturbation in each graph in the generated set. However, O'Bray et al. (2021) incorporated an objective evaluation score by computing the Pearson correlation of each metric with the degree of perturbation. Although objective, evaluation using Pearson correlation is biased in that the relationship between the metric and degree of perturbation need not be linear.

3 The effectiveness of random GNNs

While Inception v3 (Szegedy et al., 2016) has found widespread use in evaluation of generated samples (Salimans et al., 2016; Cai et al., 2018; Lunz et al., 2020), there is no analogue in graph-based tasks. This prevents a standardized analysis of GGMs. To tackle this, a single GNN may be pretrained on multiple datasets such that it extracts meaningful embeddings on new datasets (Hu et al., 2020). However, node and edge features often have incompatible dimensions across graph datasets. In addition, the distributions of graphs in the pretraining and target tasks can vary drastically, degrading the performance of a pretrained network. However, random GNNs are capable of extracting meaningful features and solving many graph tasks (Kipf & Welling, 2017; Morris et al., 2019; Xu et al., 2019).[2] Thus, avoiding pretraining and instead utilizing random GNNs may be a viable strategy to bypass these issues.

In particular, Graph Isomorphism Networks (GINs) (Xu et al., 2019) have the potential to extract meaningful graph embeddings without pretraining. These networks are as sensitive to graph isomorphism as the Weisfeiler-Lehman (WL) test, making them one of the most powerful GNNs along with more recent GNNs (Corso et al., 2020; Bouritsas et al., 2020; Bodnar et al., 2021; Beani et al., 2021). To accomplish this sensitivity, GIN utilizes summation as an injective neighborhood and

[2] In the image domain, random CNNs are also beneficial in certain cases, such as large distribution shifts (Naeem et al., 2020).

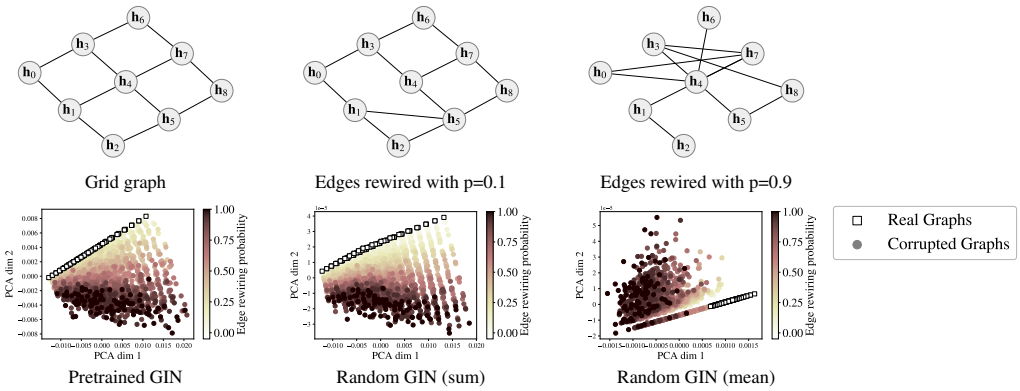

Figure 2: **Top:** Example of a grid graph corrupted by rewiring edges with various probabilities. **Bottom:** PCA of 100 grid graphs with different rewiring probabilities and embeddings obtained using various GINs. For a strong graph representation, we expect the corrupted graphs to diverge from the real graphs in embedding space as the edge rewiring probability grows. The random GIN with sum pooling (bottom middle) extracts as strong a representation as the pretrained GIN with sum pooling (bottom left). The poor representations extracted by a random GIN with mean pooling (bottom right) highlights the need for injective aggregation.

graph aggregation function $f^{(l)}$ and READOUT in Equations 2 and 3. Since this aggregation is injective regardless of model parameters, a random GIN may also be as powerful as the WL graph isomorphism test. To further demonstrate the power of a random GIN in GGM evaluation, we showcase that random weights extract a similar representation to pretrained weights if summation is used (Figure 2). However, one condition for GIN to be as powerful as the WL test is to have a sufficient number of propagation layers L in Equation 2 (Xu et al., 2019). In addition, You et al. (2020) demonstrate that the choice of L is one of the most critical for performance across diverse graph tasks. Furthermore, Equation 3 highlights that L and the node embedding size d are responsible for determining the dimensionality of the graph embedding $\mathbf{x}$, a key factor in many of the statistic-based metrics in Section 2.3. Sample sizes in frequently used GGM datasets tend to be small resulting in a bias towards minimizing the length of $\mathbf{x}$ while retaining discriminability.

4 EXPERIMENTS

In this section, we describe many properties of a strong GGM evaluation metric and thoroughly test each metric for these properties. These properties include a metric's ability to correlate with the fidelity and diversity of generated graphs, its sample efficiency, and its computational efficiency. We argue that these properties capture many desired characteristics of a strong evaluation metric and enables reliable ranking of GGMs. In many cases, we adapt the experiment design of Xu et al. (2018) to the graph domain. We test each metric and architecture combination across 10 random seeds which affects the model weights and perturbations applied. In addition to the classical metrics described in Section 2, we introduce an additional scalar baseline metric. This metric uses the Weisfeiler-Lehman (WL) subtree kernel (Shervashidze et al., 2011) as $k(\mathbf{x}_i, \mathbf{x}_j)$ in the MMD calculation (Equation 1). We refer to this metric as "WL MMD," and additional details regarding this metric are available in Appendix A.1.

Datasets. We experiment using a variety of datasets with varying sizes and characteristics to test each metric's ability to evaluate GGMs across domains. In particular, we include common GGM datasets such as Lobster, Grid, Proteins, Community, and Ego (You et al., 2018; Liao et al., 2019; Dai et al., 2020). In addition, we utilize the molecular dataset ZINC (Irwin et al., 2012) strictly to demonstrate the ability of each metric to detect changes in node and edge feature distributions with no graph structure manipulation. We provide a brief overview of each dataset in Table 1, and thorough descriptions in Appendix A.2.

GNN feature extractor. We consider GIN models (Equations 2 and 3) with $L \in [2, 3, \ldots, 7]$, and $d \in [5, 10, \ldots, 40]$. We randomly select 20 architectures inside these ranges to test in our experiments using both randomly initialized and pretrained GINs. The pretraining process tasks each network with graph classification and the methodology is described in Appendix A.2. Results for met-

rics computed using a pretrained GIN in individual experiments are left to Appendix A.3.1. However, we summarize these results in Table 3 in Section 5 to facilitate discussion. We use node degree features expressed as an integer as an inexpensive way to improve discriminability in both random and pretrained networks. In practice, we utilize orthogonal weight initialization (Saxe et al., 2014) in the random networks as it produces metrics with slightly lower variance across initializations.

Table 1: A summary of the datasets used in our experiments.

DATASET	#SAMPLES	$\lvert\mathbb{V}\rvert$	$\lvert E\rvert$	NODE/EDGE FEATURES
Grid	100	100-400	360-1368	✗
Lobster	100	10-100	10-100	✗
Proteins	918	100-500	186-1575	✗
Ego	757	50-399	57-1071	✗
Community	500	60-160	300-1800	✗
ZINC	1000	10-50	22-82	✓

Evaluating the evaluation metrics. All experiments are designed to begin with $P_g \approx P_r$, and to have a monotonically increasing degree of perturbation $t \in [0, 1]$ that is a measure of the dissimilarity between $\mathbb{S}_g$ and $\mathbb{S}_r$. We evaluate each metric objectively by computing the Spearman rank correlation between the metric scores and the degree of perturbation. Spearman rank correlation is preferable to Pearson correlation as it avoids any bias towards a linear relationship. All metrics are normalized such that $\hat{\rho} = 0$ if $P_r = P_g$ meaning a strong metric will increase with t and a rank correlation of 1.0 is assumed to be ideal. To report the results for a given metric, we first compute the rank correlation for a single random seed (which varies model weights and perturbations applied, if applicable), experiment (e.g. edge rewiring), dataset (e.g. Grid) and GIN configuration (if applicable, e.g. $L = 4, d = 25$). We then aggregate the rank correlation scores across any combination of these factors of variation.

4.1 MEASURING FIDELITY

One of the most important properties of a metric is its ability to measure the fidelity of generated samples. To test metrics for fidelity we construct two experiments. The first experiment tests the metric's ability to detect various amounts of random samples mixed with real samples (Xu et al., 2018), while the second experiment slowly decreases the quality of graphs by randomly rewiring edges (O'Bray et al., 2021). Each of these experiments begin with $\mathbb{S}_g$ as a copy of $\mathbb{S}_r$, which is itself a copy of the dataset.

In the first experiment, we utilize random graphs to impact the quality of $\mathbb{S}_g$. To decrease the similarity between $\mathbb{S}_r$ and $\mathbb{S}_g$, we slowly *mix random graphs* by increasing the ratio t of random graphs to real graphs in $\mathbb{S}_g(t)$. The random graphs are Erdős-Rényi (E-R) graphs (Erdős & Rényi, 1960) with sizes and p values chosen to resemble $\mathbb{S}_r$: for every $G_r \in \mathbb{S}_r$, there is a corresponding E-R graph G_g with $|\mathbb{V}|_g = |\mathbb{V}|_r$ and $p = \frac{|E|_r}{|\mathbb{V}|_r^2}$.

The second experiment increases distance between P_r and P_g by randomly *rewiring edges* in $\mathbb{S}_g$. Here, the degree of perturbation t is the probability of rewiring each edge $(i, j) \in E$. For each $G \in \mathbb{S}_g$ and each $(i, j) \in E$, a sample $x_{i,j} \sim \text{Bernoulli}(t)$ is drawn. Edges with $x_{i,j} = 1$ are rewired, another sample $y_{i,j} \sim \text{Bernoulli}(0.5)$ is drawn to decide which node $\{i, j\}$ of the edge is kept, and a new connection for this edge is chosen uniformly from $\mathbb{V}$.

With the exception of Recall and Orbits MMD, the majority of tested metrics excel in these experiments as indicated by a rank correlation close to 1.0 (Figure 3, top). However, Recall is specifically designed to measure diversity of $\mathbb{S}_g$ rather than fidelity, and its low sensitivity to fidelity here is expected. Surprisingly, Coverage demonstrates strong sensitivity to the fidelity of $\mathbb{S}_g$ although it is also designed to measure the diversity of $\mathbb{S}_g$.

4.2 MEASURING DIVERSITY

The next property we investigate is a metric's ability to measure the diversity of generated samples in $\mathbb{S}_g$. We focus on two common pitfalls of generative models that a strong metric must be sensitive to: mode dropping and mode collapse. We test each metric for both sensitivities independently by adapting two experiments from Xu et al. (2018) to the graph domain, both of which begin by clustering the dataset to identify the modes of P_r. We employ Affinity Propagation (Frey & Dueck, 2007) to automatically determine the number of clusters, and set the similarity between graphs to be the value obtained from the WL-subtree kernel (Shervashidze et al., 2011) using the GraKel Python library (Siglidis et al., 2020). This clustering method avoids any dependency on the GNN model

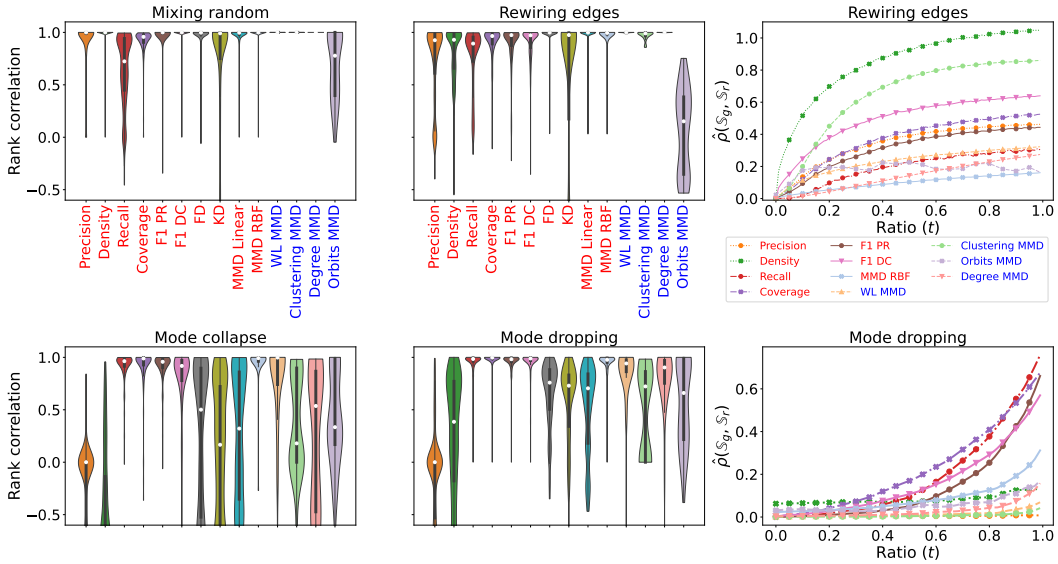

Figure 3: Results from the fidelity (top) and diversity experiments (bottom). NN-based metrics using a randomly initialized GIN are highlighted in red, while baseline metrics are in blue. Results are aggregated across all datasets and all GIN configurations (if applicable). For the violin plots, white dots are the median, thick black lines are the IQR, and thin black lines are the whiskers. **Fidelity:** Several NN-based and baseline metrics perform nearly optimally across both experiments with median rank correlations close to 1.0. **Diversity:** Classical metrics (You et al., 2018) are below average in the mode collapse experiment, and suboptimal in the mode dropping experiment. Scalar metrics such as MMD RBF and WL MMD have median rank correlations close to 1.0 and perform extremely well across both experiments.

architecture and ensures the clusters found are consistent across all experiments. Both of these experiments begin with $\mathbb{S}_r$ and $\mathbb{S}_g$ as disjoint halves of the dataset.

To simulate *mode collapse*, we progressively replace each datapoint with its cluster centre. The degree of perturbation t represents the ratio of clusters that have been collapsed in this manner. To simulate *mode dropping*, we progressively remove clusters from $\mathbb{S}_g$. To keep $|\mathbb{S}_g|$ constant, we randomly select samples from the remaining clusters to duplicate. In this experiment, the degree of perturbation t is the ratio of clusters that have been deleted from $\mathbb{S}_g$.

In the mode collapse experiment, all classical metrics (You et al., 2018) perform poorly with a rank correlation less than 0.5 (Figure 3, bottom). Classical metrics obtain slightly better, though still suboptimal results in the mode dropping experiment. As expected, Recall and Coverage exhibit strong positive correlation with the diversity of $\mathbb{S}_g$, while Precision and Density are negatively correlated. In addition, several scalar metrics such as MMD RBF and WL MMD exhibit strong correlation with the diversity of $\mathbb{S}_g$ across both experiments.

4.3 Sensitivity to node and edge features

In this experiment, we measure the sensitivity of each metric to changes in node or edge feature distributions while the underlying graph structure is static. Similar to the rewiring edges experiment, this is performed by randomizing features with a monotonically increasing probability t, and we provide more thorough descriptions of these experiments in Appendix A.2. We exclude classical metrics (You et al., 2018) and WL MMD from this experiment as they are unable to incorporate both edge and node features in evaluation. We find all NN-based metrics are sensitive to these perturbations, and a summary is shown in Table 3. More detailed results are available in Appendix A.3.1.

4.4 Sample efficiency

Frequently used datasets in the GGM literature are relatively small in size, meaning the sample efficiency of each metric is extremely important. In this experiment, we determine the sample efficiency of each metric by finding the minimum number of samples to discriminate a set of random

graphs $\mathbb{S}_g$ from real samples $\mathbb{S}_r$. The random graphs are E-R graphs generated using the same process described in Section 4.1. We begin this experiment by sampling two disjoint sets $\mathbb{S}'_r$ and $\mathbb{S}''_r$ from $\mathbb{S}_r$, and a set of random graphs $\mathbb{S}'_g$ from $\mathbb{S}_g$ with $|\mathbb{S}'_r| = |\mathbb{S}''_r| = |\mathbb{S}'_g| = n$ and small n. A metric with high sample efficiency should measure $\hat{\rho}(\mathbb{S}'_r, \mathbb{S}''_r) < \hat{\rho}(\mathbb{S}'_r, \mathbb{S}'_g)$ with a small n. Rather than using rank correlation, in this experiment we record the sample size of each metric as the smallest n where $\hat{\rho}(\mathbb{S}'_r, \mathbb{S}''_r) < \hat{\rho}(\mathbb{S}'_r, \mathbb{S}'_g) \;\; \forall i \geq n$. All of the metrics based on K-nearest neighbours and many of the classical metrics exhibit high sample efficiency and require a minimal number of samples to correctly score $\mathbb{S}''_r$ and $\mathbb{S}'_g$ (Table 3).

4.5 Computational efficiency

Computational efficiency is the final property of an evaluation metric that we examine. Metrics that are efficient to compute are ideal as they can be easily used throughout training to measure progress and perform model selection. Graph datasets can scale in several dimensions: the number of samples, average number of nodes, and average number of edges. We make use of E-R graphs to generate graphs with an arbitrary number of nodes and edges enabling us to independently scale graphs in each dimension. The results highlighting each metric's computational efficiency as dataset size increases to 10,000 samples is shown in Table 3, while the results for all dimensions along with each metric's RAM usage throughout these experiments are presented in Figures 10 and 12 in Appendix A.3. The classical metrics (You et al., 2018) quickly become prohibitive to compute as the number of samples increase, while the NN-based metrics are faster by several orders of magnitude and are inexpensive to compute at any scale.

4.6 GGM selection

While several NN-based metrics have performed consistently in our experiments, this does not necessarily mean they are suited for evaluating GGMs. With the lack of an Inception v3 (Szegedy et al., 2016) analogue for GGM evaluation, these metrics must be consistent across possibly random model parameterizations. Thus, in this section we evaluate both GRAN (Liao et al., 2019) and GraphRNN (You et al., 2018) models trained on the Grid and Lobster datasets with 80% of the graphs randomly selected for training. We use the implementations in the official GitHub repositories and train using the recommended hyperparameters. To simulate the model selection process, we select models at various stages of training for evaluation. We then generate n graphs from each model where n is the size of the dataset, and use the remaining 20% of graphs as $\mathbb{S}_r$. We compute $\hat{\rho}(\mathbb{S}_g, \mathbb{S}_r)$ across 10 different random GINs using the strongest GIN configuration identified in Section 5. Similarly, we present one of the strongest NN-based metrics identified in Section 5, MMD RBF, and compare it to a model selection process using classical metrics (You et al., 2018). MMD RBF is extremely low variance across random GINs making it a possible candidate for evaluating GGMs (Table 2). Additional strong NN-based metrics such as F1 PR and F1 DC are tested in Table 7 in Appendix A.3.4 and are also found to have low variance across random GINs. Furthermore, classical metrics are unable to unanimously rank GGMs, making model selection difficult. We perform this experiment using pretrained GINs of the same architecture in Table 7 and in some cases find that pretraining induces slightly more variance within metrics.

Table 2: Evaluation of different GGMs at various percentages of total epochs trained on the Grid and Lobster datasets. Models are evaluated using the strongest NN-based metric (MMD RBF) and classical metrics (You et al., 2018). MMD RBF is averaged across 10 different random GINs with the strongest configuration. Cells are colored according to their rank in a given column. 50/50 split represents the metric computed using a random 50/50 split of the dataset and represents an ideal score for each metric.

	Grid				Lobster			
GGM	MMD RBF	Clus.	Deg.	Orbit	MMD RBF	Clus.	Deg.	Orbit
50/50 split	0.042	0.0	$6.51e^{-5}$	0.018	0.049	0.0	0.003	0.037
Erdős-Rényi	0.203 ± 0.006	0.023	1.001	0.560	0.538 ± 0.029	0.002	0.736	0.729
GraphRNN-100%	$0.184 \pm 5e^{-5}$	$4.59e^{-8}$	0.032	0.252	0.235 ± 0.005	0.095	0.038	0.104
GraphRNN-66%	$0.154 \pm 7e^{-5}$	$1.69e^{-6}$	0.015	0.169	0.244 ± 0.006	0.099	0.050	0.104
GraphRNN-33%	$0.248 \pm 6e^{-5}$	$1.81e^{-6}$	0.022	0.134	0.137 ± 0.010	0.137	0.047	0.104
GRAN-100%	0.063 ± 0.001	$1.24e^{-6}$	$1.68e^{-4}$	0.037	0.146 ± 0.009	0.101	0.048	0.098
GRAN-66%	0.061 ± 0.001	$3.15e^{-6}$	$3.20e^{-5}$	0.047	0.161 ± 0.010	0.103	0.045	0.094
GRAN-33%	$0.060 \pm 2e^{-4}$	$1.46e^{-6}$	$4.85e^{-5}$	0.054	0.096 ± 0.013	0.125	0.080	0.099

5 DISCUSSION

We aggregate the results across experiments for a high-level overview of each metric's strengths and weaknesses in Table 3. Metrics computed using a pretrained GIN are nearly indistinguishable from those using a random GIN across rank correlation experiments (columns 3–4). Although metrics using a pretrained GIN benefit from slightly higher sample efficiency, they can have higher variance across model parameters than a simple random initialization (Table 7 in Appendix A.3.4). Similar to Table 3, we aggregate the results across all experiments and metrics for a specific GIN configuration in Table 4b in Appendix A.3. We find that the mean rank correlation taken across NN-based metrics and all experiments is approximately 10× higher variance than the mean taken across GIN configurations (0.078 and 0.007, respectively). In general, our findings indicate that the key to strong GIN-based GGM evaluation metrics is *the choice of metric rather than the feature space used to extract graph embeddings*.

Table 3: Summary of each metric's performance across experiments. Column headings indicate the experiment across which results are aggregated. NN-based metrics are aggregated across all GIN configurations using random networks unless otherwise stated. Computational efficiency is taken to be the maximum recorded time reported in Figure 10. Values reported are the mean ± std. error, and the average in the final row is taken strictly across the NN-based metrics. Cells are colored if the results can be interpreted objectively for the given experiment.

METRIC	FIDELITY	DIVERSITY	FIDELITY & DIVERSITY (RANDOM/PRETRAINED)		NODE/EDGE FEATS.	SAMPLE EFF. (RANDOM/PRETRAINED)		COMP. EFF. (S)
Orbits MMD	0.37 ± 0.048	0.49 ± 0.046	0.43 ± 0.034		N/A	122 ± 22		$1.4e^4$
Degree MMD	1.00 ± 0.000	0.51 ± 0.061	0.76 ± 0.035		N/A	9 ± 1		$7.5e^3$
Clustering MMD	0.99 ± 0.003	0.43 ± 0.047	0.72 ± 0.030		N/A	7 ± 0		$1.1e^5$
WL MMD	1.00 ± 0.000	0.82 ± 0.037	0.91 ± 0.019		N/A	10 ± 2		110
FD	0.98 ± 0.002	0.44 ± 0.013	0.71 ± 0.008	0.74 ± 0.007	0.96 ± 0.010	58 ± 3	55 ± 3	4.5
KD	0.62 ± 0.015	0.32 ± 0.014	0.47 ± 0.010	0.58 ± 0.009	0.94 ± 0.011	89 ± 4	63 ± 3	5.1
Precision	0.82 ± 0.007	-0.25 ± 0.010	0.29 ± 0.011	0.29 ± 0.011	0.99 ± 0.001	7 ± 0	7 ± 0	18
Recall	0.70 ± 0.007	0.93 ± 0.003	0.82 ± 0.004	0.81 ± 0.005	0.80 ± 0.018	7 ± 0	7 ± 0	18
Density	0.90 ± 0.004	-0.10 ± 0.015	0.40 ± 0.011	0.36 ± 0.012	0.99 ± 0.001	7 ± 0	7 ± 0	12
Coverage	0.91 ± 0.003	0.95 ± 0.003	0.93 ± 0.002	0.93 ± 0.003	0.99 ± 0.000	7 ± 0	7 ± 0	12
F1 PR	0.92 ± 0.004	0.93 ± 0.003	0.93 ± 0.003	0.93 ± 0.002	0.99 ± 0.000	7 ± 0	7 ± 0	18
F1 DC	0.95 ± 0.002	0.86 ± 0.007	0.91 ± 0.004	0.88 ± 0.004	0.99 ± 0.000	7 ± 0	7 ± 0	12
MMD Linear	0.98 ± 0.002	0.37 ± 0.012	0.68 ± 0.008	0.75 ± 0.007	0.99 ± 0.005	57 ± 3	31 ± 2	4.5
MMD RBF	0.97 ± 0.002	0.95 ± 0.003	0.96 ± 0.002	0.97 ± 0.002	1.00 ± 0.001	42 ± 2	12 ± 1	120
Average (NN-based)	0.88 ± 0.039	0.54 ± 0.144	0.71 ± 0.078	0.72 ± 0.076	0.96 ± 0.019	29 ± 10	20 ± 7	—

In practice, we recommend the use of either the MMD RBF or F1 PR metrics. Although we found the choice of GIN architecture to be unimportant, we suggest the use of a random GIN with 3 rounds of graph propagation and a node embedding size of 35 for consistency in future work. This is the strongest GIN configuration across all experiments (Table 4b in Appendix A.3.2). Both metrics have been shown to be sensitive to changes in both fidelity and diversity while having minimal variance across random initializations. In addition, they are capable of detecting changes in node and edge feature distributions (Table 3), making them useful for a wide array of datasets. While MMD RBF has slightly stronger correlation with the fidelity and diversity of generated graphs, F1 PR has superior sample and computational efficiency. Thus, we suggest the use of F1 PR in cases where there are less samples than the measured sample size of MMD RBF (i.e., 42) or MMD RBF is extremely prohibitive to compute, otherwise, we suggest MMD RBF.

6 CONCLUSION

We introduced the use of an untrained random GIN in GGM evaluation, inspired by metrics popular in the image domain. We discovered that pre-existing GGM metrics fail to capture the diversity of generated graphs, and find several random GIN-based metrics that are more expressive while being domain-agnostic at a significantly reduced computational cost. Conversely, we find that pretrained GIN-based metrics have superior sample efficiency to metrics utilizing a random GIN. An interesting direction for future work is to explore more sophisticated GNNs to improve sample efficiency.

REFERENCES

Francis R. Bach, Gert R. G. Lanckriet, and Michael I. Jordan. Multiple kernel learning, conic duality, and the smo algorithm. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2004.

Victor Bapst, Alvaro Sanchez-Gonzalez, Carl Doersch, Kimberly Stachenfeld, Pushmeet Kohli, Peter Battaglia, and Jessica Hamrick. Structured agents for physical construction. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019.

Dominique Beani, Saro Passaro, Vincent Létourneau, Will Hamilton, Gabriele Corso, and Pietro Liò. Directional graph networks. In *International Conference on Machine Learning*, pp. 748–758. PMLR, 2021.

Mikołaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying MMD GANs. *arXiv preprint arXiv:1801.01401*, 2018.

Cristian Bodnar, Fabrizio Frasca, Nina Otter, Yu Guang Wang, Pietro Liò, Guido Montúfar, and Michael Bronstein. Weisfeiler and lehman go cellular: Cw networks. *arXiv preprint arXiv:2106.12575*, 2021.

Giorgos Bouritsas, Fabrizio Frasca, Stefanos Zafeiriou, and Michael M Bronstein. Improving graph neural network expressivity via subgraph isomorphism counting. *arXiv preprint arXiv:2006.09252*, 2020.

Haoye Cai, Chunyan Bai, Yu-Wing Tai, and Chi-Keung Tang. Deep video generation, prediction and completion of human action sequences. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.

Gabriele Corso, Luca Cavalleri, Dominique Beaini, Pietro Liò, and Petar Velickovic. Principal neighbourhood aggregation for graph nets. *arXiv preprint arXiv:2004.05718*, 2020.

Hanjun Dai, Azade Nazi, Yujia Li, Bo Dai, and Dale Schuurmans. Scalable deep generative modeling for sparse graphs. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2020.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.

Paul D. Dobson and Andrew J. Doig. Distinguishing enzyme structures from non-enzymes without alignments. *Journal of molecular biology*, 330(4):771–783, 2003.

Paul Erdős and Alfréd Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5(1):17–60, 1960.

Brendan J. Frey and Delbert Dueck. Clustering by passing messages between data points. *Science*, 315(5814):972–976, 2007.

Damien Garreau, Wittawat Jitkrittum, and Motonobu Kanagawa. Large sample analysis of the median heuristic. *arXiv preprint arXiv:1707.07269*, 2018.

Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander J. Smola. A kernel method for the two-sample problem. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 19, 2006.

Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander J. Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(1):723–773, 2012a.

Arthur Gretton, Dino Sejdinovic, Heiko Strathmann, Sivaraman Balakrishnan, Massimiliano Pontil, Kenji Fukumizu, and Bharath K Sriperumbudur. Optimal kernel choice for large-scale two-sample tests. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 22, 2012b.

Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, Günter Klambauer, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a nash equilibrium. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, 2017.

Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. Strategies for pre-training graph neural networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2020.

John J. Irwin, Teague Sterling, Michael M. Mysinger, Erin S. Bolstad, and Ryan G. Coleman. ZINC: A free tool to discover chemistry for biology. *Journal of Chemical Information and Modeling*, 52 (7):1757–1768, 2012.

Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Multi-objective molecule generation using interpretable substructures. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2020.

Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.

Xiangzhe Kong, Zhixing Tan, and Yang Liu. GraphPiece: Efficiently generating high-quality molecular graph with substructures. *arXiv preprint arXiv:2106.15098*, 2021.

Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019.

Yujia Li, Oriol Vinyals, Chris Dyer, Razvan Pascanu, and Peter Battaglia. Learning deep generative models of graphs. *arXiv preprint arXiv:1803.03324*, 2018.

Renjie Liao, Yujia Li, Yang Song, Shenlong Wang, William L Hamilton, David Duvenaud, Raquel Urtasun, and Richard Zemel. Efficient graph generation with graph recurrent attention networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32, 2019.

Chia-Cheng Liu, Harris Chan, and Kevin Luk. Auto-regressive Graph Generation Modeling with Improved Evaluation Methods. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

Mario Lucic, Karol Kurach, Marcin Michalski, Sylvain Gelly, and Olivier Bousquet. Are GANs created equal? a large-scale study. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 31, 2018.

Sebastian Lunz, Yingzhen Li, Andrew Fitzgibbon, and Nate Kushman. Inverse graphics GAN: Learning to generate 3D shapes from unstructured 2D data. *arXiv preprint arXiv:2002.12674*, 2020.

Christopher Morris, Martin Ritzert, Matthias Fey, William L Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. Weisfeiler and Leman go neural: Higher-order graph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2019.

Muhammad Ferjad Naeem, Seong Joon Oh, Youngjung Uh, Yunjey Choi, and Jaejun Yoo. Reliable fidelity and diversity metrics for generative models. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2020.

Leslie O'Bray, Max Horn, Bastian Rieck, and Karsten M. Borgwardt. Evaluation metrics for graph generative models: Problems, pitfalls, and practical solutions. *arXiv preprint arXiv:2106.01098*, 2021.

Mariya Popova, Mykhailo Shvets, Junier Oliva, and Olexandr Isayev. MolecularRNN: Generating realistic molecular graphs with optimized properties. *arXiv preprint arXiv:1905.13372*, 2019.

Kristina Preuer, Philipp Renz, Thomas Unterthiner, Sepp Hochreiter, and Günter Klambauer. Fréchet chemnet distance: a metric for generative models for molecules in drug discovery. *Journal of chemical information and modeling*, 58(9):1736–1741, 2018.

Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training GANs. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 29, 2016.

Bidisha Samanta, Abir De, Gourhari Jana, Vicenç Gómez, Pratim Kumar Chattaraj, Niloy Ganguly, and Manuel Gomez-Rodriguez. NeVAE: A deep generative model for molecular graphs. *Journal of Machine Learning Research*, 21(114):1–33, 2020.

Andrew M. Saxe, James L. McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014.

Bernhard Schölkopf, Alexander J. Smola, and Klaus-Robert Müller. Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation*, 10(5):1299–1319, 1998.

Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective classification in network data. *AI Magazine*, 29(3):93, Sep. 2008.

Nino Shervashidze, Pascal Schweitzer, Erik Jan van Leeuwen, Kurt Mehlhorn, and Karsten M. Borgwardt. Weisfeiler-Lehman graph kernels. *Journal of Machine Learning Research*, 12(77): 2539–2561, 2011.

Giannis Siglidis, Giannis Nikolentzos, Stratis Limnios, Christos Giatsidis, Konstantinos Skianis, and Michalis Vazirgiannis. GraKeL: A graph kernel library in Python. *Journal of Machine Learning Research*, 21(54):1–5, 2020.

Bharath K. Sriperumbudur, Kenji Fukumizu, Arthur Gretton, Gert R. G. Lanckriet, and Bernhard Schölkopf. Kernel choice and classifiability for RKHS embeddings of probability distributions. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 22, 2009.

Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

Lucas Theis, Aäron van den Oord, and Matthias Bethge. A note on the evaluation of generative models, 2016.

Rylee Thompson, Elahe Ghalebi, Terrance DeVries, and Graham W. Taylor. Building LEGO using deep generative models of graphs. In *NeurIPS Workshop on Machine Learning for Engineering Modeling, Simulation, and Design (ML4Eng)*, 2020.

Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1):4–24, 2020.

Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.

Qiantong Xu, Gao Huang, Yang Yuan, Chuan Guo, Yu Sun, Felix Wu, and Kilian Weinberger. An empirical study on evaluation metrics of generative adversarial networks. *arXiv preprint arXiv:1806.07755*, 2018.

Jiaxuan You, Rex Ying, Xiang Ren, William L Hamilton, and Jure Leskovec. GraphRNN: Generating realistic graphs with deep auto-regressive models. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2018.

Jiaxuan You, Zhitao Ying, and Jure Leskovec. Design space for graph neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, 2020.

A APPENDIX

In this material, we cover the details that are missing in the main article.

A.1 ADDITIONAL METRIC DETAILS

For Precision, Recall, Density, and Coverage, we set $k = 5$ for all experiments (Naeem et al., 2020). For the MMD RBF metric, we compute MMD as $\mathrm{MMD}(\mathbb{S}_g, \mathbb{S}_r) = \max\{\mathrm{MMD}(\mathbb{S}_g, \mathbb{S}_r; \sigma) \mid \sigma \in \Sigma\}$, where the values of Σ are multiplied by the mean pairwise distance of $\mathbb{S}_g$ and $\mathbb{S}_r$. While the median pairwise distance is the heuristic (Garreau et al., 2018), we find the mean to be more robust in our experiments. Before this scaling factor is applied, we use $\Sigma = \{0.01, 0.1, 0.25, 0.5, 0.75, 1.0, 2.5, 5.0, 7.5, 10.0\}$.

For the WL MMD metric, we perform all experiments with four rounds of graph propagation. As with the classical metrics from (You et al., 2018), it does not rely on a GNN and is unable to include continuous node and edge features in evaluation. In addition, the WL-subtree kernel is only positive-semidefinite (Shervashidze et al., 2011), while MMD requires a positive definite kernel (Gretton et al., 2006). It is for these reasons that it was left as an additional baseline for NN-based metrics.

A.2 ADDITIONAL EXPERIMENTAL DETAILS

Following Section 2.2, we denote a graph as $G = (\mathbb{V}, E)$ with n vertices $\mathbb{V}$ and edges $E = \{(i, j) \mid i, j \in \{1, \ldots, n\}\}$. Node features for node $i \in \mathbb{V}$ are denoted as $\mathbf{h}_i \in \mathbb{R}^b$ ($\mathbf{h}_i^{(0)}$ in Section 2.2). Similarly, $\boldsymbol{A} \in \mathbb{R}^{n \times n \times a}$ stores edge information, with the $(i, j)^{th}$ entry corresponding to the edge features for edge (i, j). Unless otherwise specified in Appendix A.2.1, $\boldsymbol{A}$ is not utilized and we set $\mathbf{h}_i$ to the degree of node i expressed as an integer as an inexpensive way to improve discriminability. We handle edge features by concatenating $\boldsymbol{A}_{i,j,:}$ to messages from node i to node j at each round of graph propagation (Hu et al., 2020).

A.2.1 DATASETS

We perform experiments using a variety of datasets with varying sizes and characteristics to thoroughly test each metric's ability to evaluate GGMs across domains. We place a slight emphasis on datasets that are frequently used in the GGM literature.

Lobster. A set of 100 stochastic graphs where each node is at most 2 hops away from a backbone path. We generate lobster graphs with $10 \leq |\mathbb{V}| \leq 100$ (Dai et al., 2020).

Grid. 100 2D grid graphs generated with $100 \leq |\mathbb{V}| \leq 400$ (Dai et al., 2020; You et al., 2018; Liao et al., 2019).

Proteins. 918 protein graphs where nodes are amino acids and two nodes are connected by an edge if they are less than 6 Angstroms away (Dobson & Doig, 2003). We select graphs with $100 \leq |\mathbb{V}| \leq 500$ (You et al., 2018; Liao et al., 2019; Dai et al., 2020).

Ego. 757 3-hop ego networks with $50 \leq |\mathbb{V}| \leq 399$ (You et al., 2018). Graphs are extracted from the CiteSeer network where nodes represent documents and edges represent citation relationships (Sen et al., 2008).

Community. 500 two-community graphs with $60 \leq |\mathbb{V}| \leq 160$. Each community is generated by the Erdős-Rényi model (E-R) (Erdős & Rényi, 1960) with $n = |\mathbb{V}|/2$ and $p = 0.3$. We then add $0.05|\mathbb{V}|$ inter-community edges with uniform probability (You et al., 2018).

ZINC. 250k real-world molecular graphs (Irwin et al., 2012) of which we randomly select 1000 samples for efficiency with $10 \leq |\mathbb{V}| \leq 50$. We set $\mathbf{h}_i$ to be a one-hot encoding of the element of node i, and $\boldsymbol{A}_{i,j,:}$ to be a one-hot encoding $\boldsymbol{a} \in \mathbb{R}^a$ of the bond type between nodes i and j. This dataset is unique in that it is the only one to naturally contain node and edge features. Thus, we can use this dataset to determine a metric's sensitivity to changes in the edge and node feature distributions without generating artificial labels.

A.2.2 RANDOMIZING EDGE FEATURES

This experiment is performed exclusively on the ZINC dataset as it naturally contains edge features. Here, the degree of perturbation t is the probability of randomizing each edge feature $\boldsymbol{A}_{i,j,:}, \forall (i, j) \in E$. For each $G \in \mathbb{S}_g$ and each edge $(i, j) \in E$, a sample $x_{i,j} \sim \mathrm{Bernoulli}(t)$ is drawn. Edges with $x_{i,j} = 1$ have their edge feature randomized. As these features are a one-hot encoding, we sample

the new edge feature uniformly from the set of valid edge features. The results of this experiment are shown in Figure 9.

A.2.3 RANDOMIZING NODE FEATURES

This experiment is performed exclusively on the ZINC dataset as it naturally contains node features. Here, the degree of perturbation t is the probability of randomizing each node feature $\mathbf{h}_i, \forall i \in \mathbb{V}$. For each $G \in \mathbb{S}_g$ and each node $i \in \mathbb{V}$, a sample $x_i \sim \text{Bernoulli}(t)$ is drawn. Nodes with $x_i = 1$ have their feature randomized. As these features are a one-hot encoding, we sample the new node feature uniformly from the set of valid node features. The results of this experiment are shown in Figure 9.

A.2.4 COMPUTATIONAL EFFICIENCY

All experiments in this section were conducted on an Intel Platinum 8160F Skylake @ 2.1Ghz with 4 CPU cores. We benchmark NN-based metrics using the most expensive GIN configuration tested in our experiments. This was a configuration with 7 propagation layers and a node embedding size of 20. The results for this experiment on CPU are shown in Figure 10 and on GPU in Figure 11. In addition, the approximate RAM usage for each metric throughout these experiments are shown in Figure 12 and Table 6.

Scaling number of samples. We generate various sets of E-R graphs (Erdős & Rényi, 1960) to scale the number of samples in a dataset independently from the number of nodes and number of edges per graph. We generate sets of graphs with the number of samples in $[100, 1000, 2000, \ldots, 10\text{k}]$ with $|\mathbb{V}| = 50$ for all graphs. We choose $p = \frac{|E|_{\text{avg}}}{|\mathbb{V}|^2_{\text{avg}}}$ for all graphs, where the averages are taken across the Proteins dataset.

Scaling number of edges. To scale the number of edges independently, we generate 50 E-R graphs per set with $|\mathbb{V}| = 1000$. The graphs are generated with p in $[0.01, 0.1, 0.2, \ldots, 1.0]$.

Scaling number of nodes. We generate 50 graphs of a constant size per set with $|\mathbb{V}|$ in $[1000, 10\text{k}, 20\text{k}, \ldots, 100\text{k}]$. For each set, we use $p = \frac{10000}{|\mathbb{V}|^2}$ to generate graphs with approximately 10000 edges per graph.

A.2.5 GNN PRETRAINING

We isolate the comparison between randomly initialized and pretrained weights to datasets without node or edge features (i.e. excluding ZINC). The GNN is trained under the multiclass classification setting to classify a given graph into the remaining datasets. To increase the difficulty of the problem, for each graph G_r in the remaining datasets, we generate an E-R graph (Erdős & Rényi, 1960) G_g with $|\mathbb{V}|_g = |\mathbb{V}|_r$ and $p = \frac{|E|_r}{|\mathbb{V}|^2_r}$. These E-R graphs have separate labels, thereby doubling the number of classes present. We use the same optimizer across all experiments and select the model with the lowest validation loss. We train each configuration across 10 random seeds to perform a fair comparison with the 10 different random initializations used in our experiments.

A.3 ADDITIONAL RESULTS

A.3.1 INDIVIDUAL EXPERIMENTS

We provide violin plots for each individual experiment in Figures 4 through 8.

A.3.2 ALL EXPERIMENTS

We aggregate results by GIN configuration across all experiments in Tables 4a and 5b. In addition, we aggregate results by GIN configuration and metric combination and highlight the top 20 results in Tables 5a and 5b.

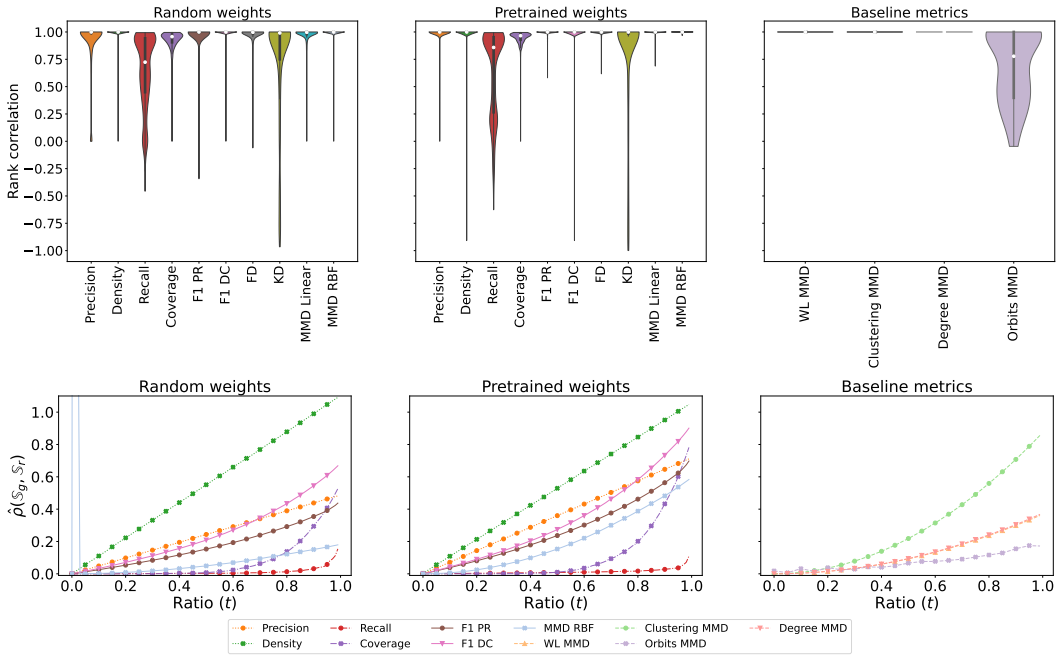

Figure 4: Results from the **mixing random graphs** experiment across all datasets and all GIN configurations (if applicable).

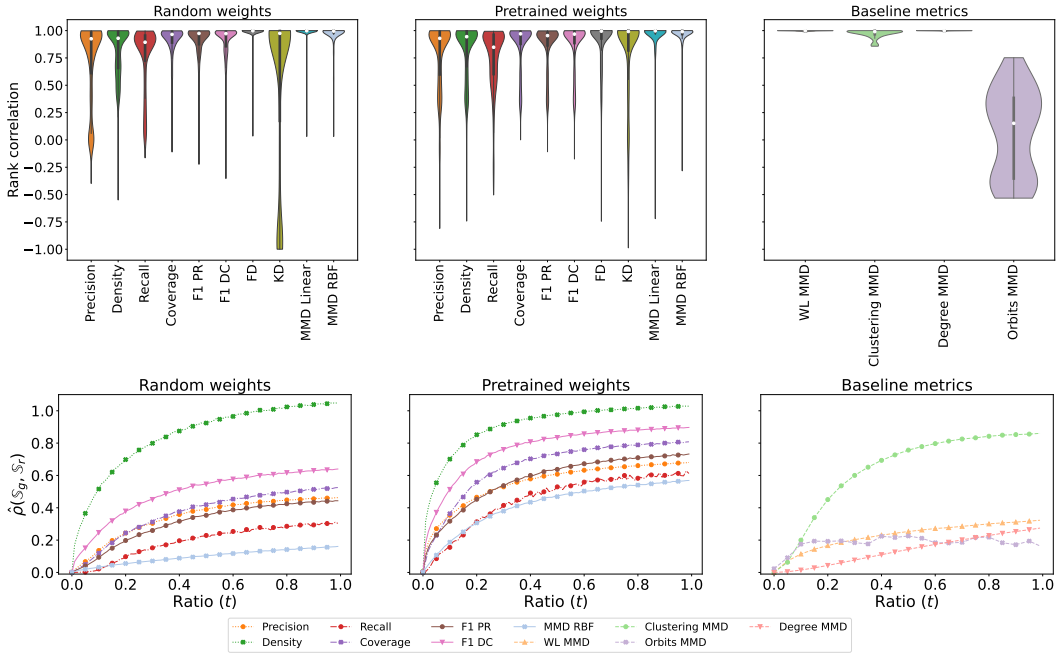

Figure 5: Results from the **rewiring edges** experiment across all datasets and all GIN configurations (if applicable).

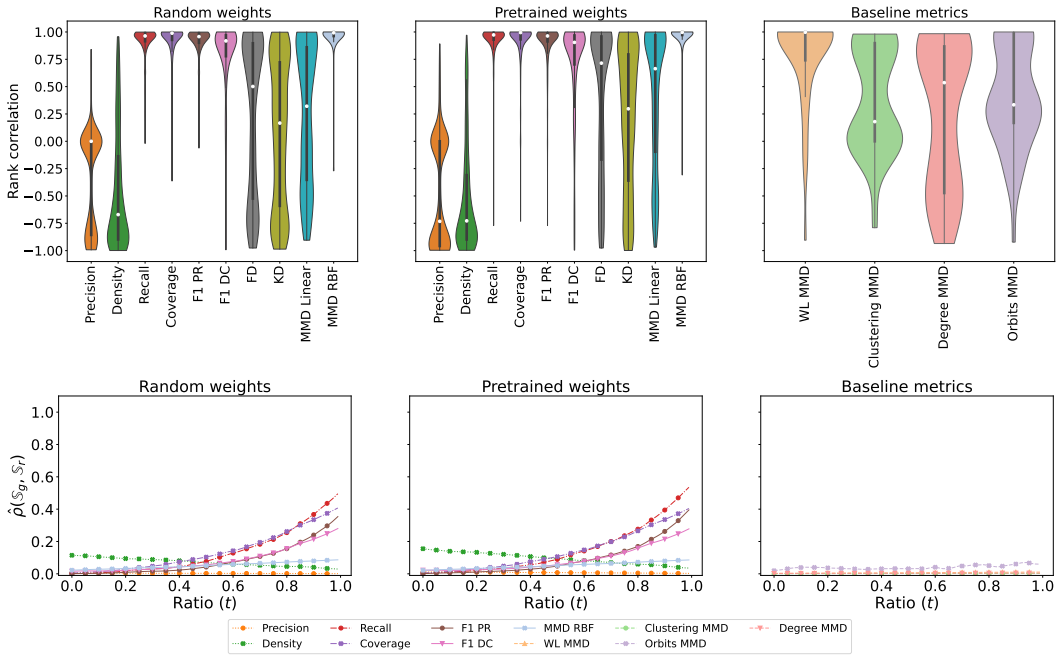

Figure 6: Results from the **mode collapse** experiment across all datasets and all GIN configurations (if applicable).

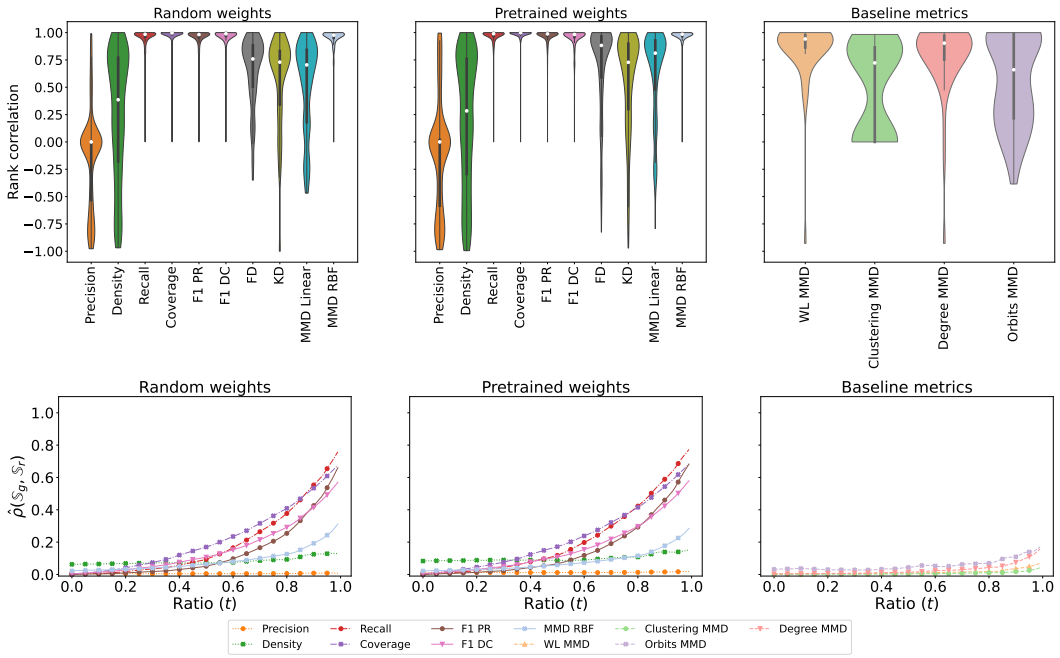

Figure 7: Results from the **mode dropping** experiment across all datasets and all GIN configurations (if applicable).

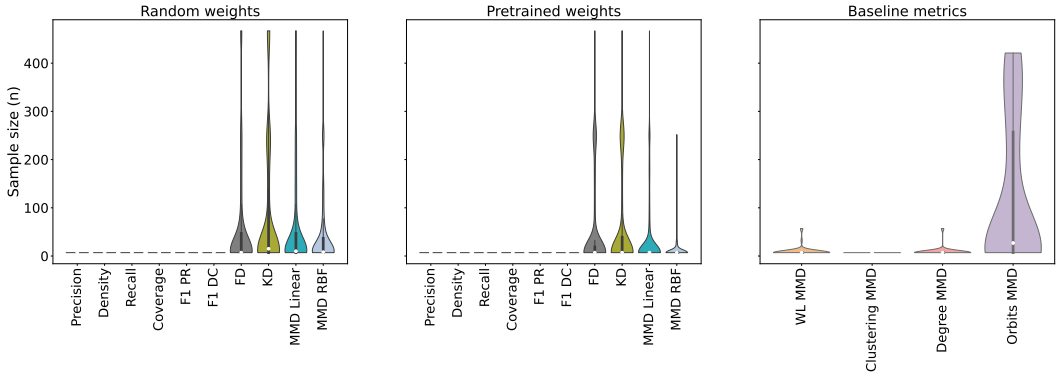

Figure 8: Results from the **sample efficiency** experiment across all datasets and all GIN configurations (if applicable).

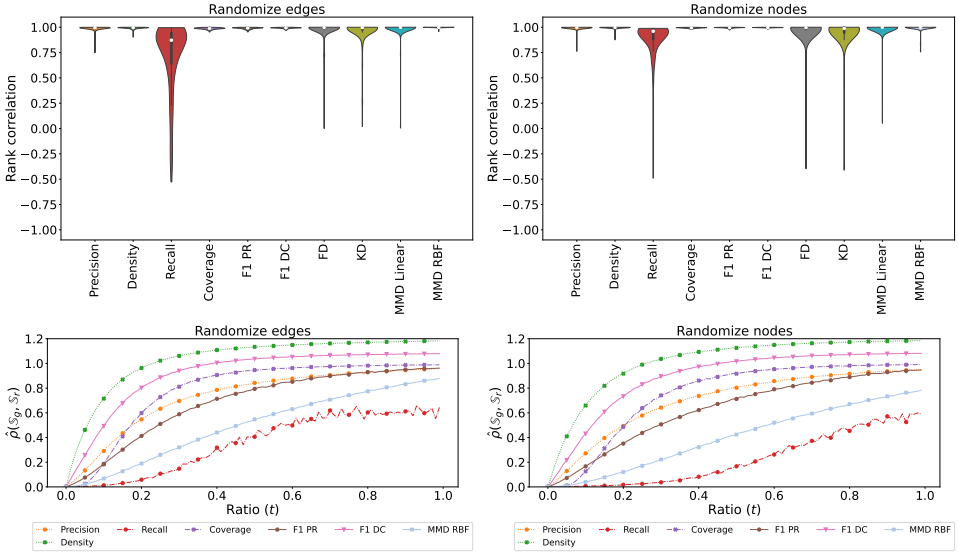

Figure 9: Results from the randomizing edge features (left) and randomizing node features (right) experiments on the ZINC dataset.

Table 4: Results across fidelity, diversity, and sample efficiency experiments and all datasets grouped by GIN configuration. N is the node embedding size, P is the number of graph propagation rounds, and G is the dimensionality of the obtained graph embedding. Configurations are sorted by their mean rank correlation across fidelity and diversity experiments.

(a) Pretrained GIN

GIN CONFIG.	FIDELITY	DIVERSITY	FIDELITY & DIVERSITY	SAMPLE EFF.
N: 25, P: 3, G: 50	0.90 ± 0.007	0.59 ± 0.019	0.75 ± 0.011	17 ± 2
N: 35, P: 5, G: 140	0.89 ± 0.008	0.59 ± 0.019	0.74 ± 0.011	19 ± 2
N: 35, P: 3, G: 70	0.91 ± 0.007	0.56 ± 0.020	0.74 ± 0.011	19 ± 2
N: 30, P: 5, G: 120	0.90 ± 0.008	0.57 ± 0.019	0.73 ± 0.011	20 ± 2
N: 40, P: 6, G: 200	0.89 ± 0.008	0.57 ± 0.019	0.73 ± 0.011	20 ± 2
N: 40, P: 2, G: 40	0.90 ± 0.007	0.56 ± 0.020	0.73 ± 0.011	22 ± 3
N: 40, P: 5, G: 160	0.87 ± 0.009	0.59 ± 0.019	0.73 ± 0.011	21 ± 2
N: 20, P: 5, G: 80	0.90 ± 0.008	0.56 ± 0.020	0.73 ± 0.011	19 ± 2
N: 15, P: 6, G: 75	0.89 ± 0.008	0.56 ± 0.020	0.73 ± 0.011	20 ± 2
N: 15, P: 2, G: 15	0.89 ± 0.008	0.56 ± 0.019	0.73 ± 0.011	23 ± 3
N: 20, P: 7, G: 120	0.90 ± 0.007	0.55 ± 0.020	0.73 ± 0.011	22 ± 3
N: 20, P: 3, G: 40	0.91 ± 0.007	0.54 ± 0.020	0.72 ± 0.011	18 ± 2
N: 10, P: 7, G: 60	0.89 ± 0.009	0.56 ± 0.020	0.72 ± 0.011	22 ± 3
N: 5, P: 5, G: 20	0.90 ± 0.008	0.54 ± 0.020	0.72 ± 0.012	18 ± 2
N: 10, P: 3, G: 20	0.89 ± 0.008	0.55 ± 0.020	0.72 ± 0.011	16 ± 2
N: 15, P: 7, G: 90	0.88 ± 0.009	0.55 ± 0.020	0.72 ± 0.011	20 ± 2
N: 25, P: 4, G: 75	0.88 ± 0.009	0.56 ± 0.020	0.72 ± 0.011	21 ± 2
N: 5, P: 7, G: 30	0.89 ± 0.008	0.54 ± 0.020	0.72 ± 0.012	22 ± 3
N: 10, P: 2, G: 10	0.86 ± 0.010	0.56 ± 0.020	0.71 ± 0.011	26 ± 3
N: 5, P: 2, G: 5	0.86 ± 0.009	0.51 ± 0.019	0.69 ± 0.011	24 ± 3
Average	0.89 ± 0.003	0.56 ± 0.004	0.73 ± 0.003	20 ± 1

(b) Random GIN.

GIN CONFIG.	FIDELITY	DIVERSITY	FIDELITY & DIVERSITY	SAMPLE EFF.
N: 35, P: 3, G: 70	0.92 ± 0.008	0.55 ± 0.019	0.73 ± 0.011	31 ± 3
N: 10, P: 3, G: 20	0.92 ± 0.008	0.55 ± 0.019	0.73 ± 0.011	34 ± 4
N: 25, P: 3, G: 50	0.92 ± 0.008	0.55 ± 0.019	0.73 ± 0.011	31 ± 3
N: 35, P: 5, G: 140	0.91 ± 0.008	0.55 ± 0.019	0.73 ± 0.011	18 ± 2
N: 25, P: 4, G: 75	0.91 ± 0.009	0.55 ± 0.019	0.73 ± 0.011	26 ± 3
N: 20, P: 3, G: 40	0.91 ± 0.008	0.54 ± 0.020	0.72 ± 0.011	34 ± 4
N: 40, P: 5, G: 160	0.90 ± 0.008	0.54 ± 0.019	0.72 ± 0.011	20 ± 2
N: 30, P: 5, G: 120	0.90 ± 0.009	0.54 ± 0.019	0.72 ± 0.011	22 ± 2
N: 15, P: 7, G: 90	0.90 ± 0.008	0.54 ± 0.019	0.72 ± 0.011	26 ± 3
N: 40, P: 6, G: 200	0.90 ± 0.008	0.53 ± 0.020	0.72 ± 0.011	20 ± 3
N: 20, P: 5, G: 80	0.90 ± 0.009	0.54 ± 0.020	0.72 ± 0.012	30 ± 4
N: 15, P: 6, G: 75	0.89 ± 0.009	0.54 ± 0.019	0.72 ± 0.011	26 ± 3
N: 20, P: 7, G: 120	0.90 ± 0.007	0.53 ± 0.020	0.72 ± 0.011	22 ± 3
N: 10, P: 7, G: 60	0.90 ± 0.008	0.53 ± 0.019	0.71 ± 0.011	25 ± 3
N: 5, P: 7, G: 30	0.90 ± 0.009	0.53 ± 0.020	0.71 ± 0.012	25 ± 3
N: 5, P: 5, G: 20	0.87 ± 0.011	0.55 ± 0.019	0.71 ± 0.012	36 ± 4
N: 40, P: 2, G: 40	0.80 ± 0.011	0.55 ± 0.019	0.68 ± 0.011	32 ± 3
N: 15, P: 2, G: 15	0.78 ± 0.012	0.56 ± 0.019	0.67 ± 0.011	46 ± 5
N: 10, P: 2, G: 10	0.75 ± 0.014	0.52 ± 0.021	0.64 ± 0.013	29 ± 3
N: 5, P: 2, G: 5	0.74 ± 0.012	0.50 ± 0.019	0.62 ± 0.012	44 ± 4
Average	0.88 ± 0.013	0.54 ± 0.003	0.71 ± 0.007	29 ± 2

Table 5: Results across fidelity, diversity, and sample efficiency experiments and all datasets grouped by GIN configuration and metric used. N is the node embedding size, P is the number of graph propagation rounds, and G is the dimensionality of the obtained graph embedding. Configurations are sorted by their mean rank correlation across fidelity and diversity experiments and the top 20 results are shown.

(a) Pretrained GIN.

GIN CONFIG., METRIC	FIDELITY	DIVERSITY	FIDELITY & DIVERSITY	SAMPLE EFF.
N: 35, P: 3, G: 70, MMD RBF	0.99 ± 0.005	0.97 ± 0.005	0.98 ± 0.004	9 ± 1
N: 25, P: 4, G: 75, MMD RBF	0.98 ± 0.005	0.97 ± 0.007	0.98 ± 0.004	10 ± 2
N: 20, P: 3, G: 40, MMD RBF	0.99 ± 0.006	0.97 ± 0.013	0.98 ± 0.007	9 ± 1
N: 30, P: 5, G: 120, MMD RBF	0.98 ± 0.007	0.97 ± 0.007	0.97 ± 0.005	11 ± 3
N: 20, P: 5, G: 80, MMD RBF	0.98 ± 0.007	0.97 ± 0.010	0.97 ± 0.006	9 ± 1
N: 40, P: 2, G: 40, MMD RBF	0.99 ± 0.004	0.96 ± 0.014	0.97 ± 0.007	12 ± 2
N: 10, P: 7, G: 60, MMD RBF	0.98 ± 0.004	0.96 ± 0.014	0.97 ± 0.007	12 ± 2
N: 10, P: 2, G: 10, MMD RBF	0.99 ± 0.004	0.96 ± 0.014	0.97 ± 0.007	19 ± 5
N: 10, P: 3, G: 20, MMD RBF	0.98 ± 0.006	0.96 ± 0.013	0.97 ± 0.007	11 ± 3
N: 25, P: 3, G: 50, MMD RBF	0.98 ± 0.008	0.96 ± 0.013	0.97 ± 0.007	10 ± 2
N: 15, P: 7, G: 90, MMD RBF	0.98 ± 0.006	0.96 ± 0.013	0.97 ± 0.007	9 ± 1
N: 20, P: 7, G: 120, MMD RBF	0.98 ± 0.006	0.96 ± 0.014	0.97 ± 0.008	10 ± 1
N: 40, P: 6, G: 200, MMD RBF	0.98 ± 0.010	0.97 ± 0.010	0.97 ± 0.007	9 ± 1
N: 5, P: 5, G: 20, MMD RBF	0.98 ± 0.011	0.97 ± 0.012	0.97 ± 0.008	11 ± 2
N: 15, P: 6, G: 75, MMD RBF	0.98 ± 0.010	0.96 ± 0.014	0.97 ± 0.009	9 ± 1
N: 5, P: 7, G: 30, MMD RBF	0.98 ± 0.011	0.96 ± 0.013	0.97 ± 0.008	14 ± 5
N: 15, P: 2, G: 15, MMD RBF	0.98 ± 0.007	0.96 ± 0.014	0.97 ± 0.008	17 ± 5
N: 35, P: 5, G: 140, MMD RBF	0.96 ± 0.015	0.97 ± 0.008	0.97 ± 0.008	9 ± 1
N: 40, P: 5, G: 160, MMD RBF	0.97 ± 0.010	0.96 ± 0.011	0.96 ± 0.007	9 ± 1
N: 15, P: 2, G: 15, F1 PR	0.94 ± 0.011	0.95 ± 0.010	0.95 ± 0.007	7 ± 0

(b) Random GIN.

GIN CONFIG., METRIC	FIDELITY	DIVERSITY	FIDELITY & DIVERSITY	SAMPLE EFF.
N: 10, P: 7, G: 60, MMD RBF	0.99 ± 0.002	0.95 ± 0.013	0.97 ± 0.007	25 ± 6
N: 25, P: 4, G: 75, MMD RBF	0.99 ± 0.002	0.95 ± 0.013	0.97 ± 0.007	30 ± 7
N: 25, P: 3, G: 50, MMD RBF	0.99 ± 0.002	0.95 ± 0.014	0.97 ± 0.007	51 ± 10
N: 20, P: 5, G: 80, MMD RBF	0.99 ± 0.005	0.95 ± 0.014	0.97 ± 0.007	31 ± 7
N: 10, P: 3, G: 20, MMD RBF	0.99 ± 0.003	0.95 ± 0.014	0.97 ± 0.007	56 ± 11
N: 35, P: 3, G: 70, MMD RBF	0.99 ± 0.003	0.95 ± 0.014	0.97 ± 0.007	52 ± 10
N: 20, P: 3, G: 40, MMD RBF	0.99 ± 0.005	0.95 ± 0.014	0.97 ± 0.007	52 ± 11
N: 5, P: 5, G: 20, MMD RBF	0.99 ± 0.006	0.95 ± 0.013	0.97 ± 0.007	41 ± 9
N: 5, P: 7, G: 30, MMD RBF	0.99 ± 0.007	0.95 ± 0.013	0.97 ± 0.008	25 ± 7
N: 30, P: 5, G: 120, MMD RBF	0.99 ± 0.008	0.95 ± 0.014	0.97 ± 0.008	29 ± 6
N: 15, P: 7, G: 90, MMD RBF	0.98 ± 0.008	0.96 ± 0.013	0.97 ± 0.008	20 ± 4
N: 15, P: 6, G: 75, MMD RBF	0.98 ± 0.008	0.96 ± 0.013	0.97 ± 0.008	21 ± 4
N: 20, P: 7, G: 120, MMD RBF	0.98 ± 0.007	0.96 ± 0.013	0.97 ± 0.008	23 ± 5
N: 35, P: 5, G: 140, MMD RBF	0.98 ± 0.008	0.95 ± 0.013	0.97 ± 0.008	17 ± 3
N: 40, P: 6, G: 200, MMD RBF	0.97 ± 0.009	0.95 ± 0.013	0.96 ± 0.008	17 ± 4
N: 40, P: 5, G: 160, MMD RBF	0.97 ± 0.010	0.95 ± 0.014	0.96 ± 0.008	20 ± 4
N: 35, P: 3, G: 70, F1 PR	0.98 ± 0.005	0.94 ± 0.011	0.96 ± 0.006	7 ± 0
N: 20, P: 3, G: 40, F1 PR	0.98 ± 0.005	0.94 ± 0.010	0.96 ± 0.006	7 ± 0
N: 10, P: 3, G: 20, F1 PR	0.99 ± 0.002	0.93 ± 0.011	0.96 ± 0.006	7 ± 0
N: 25, P: 4, G: 75, Coverage	0.95 ± 0.005	0.96 ± 0.012	0.96 ± 0.007	7 ± 0

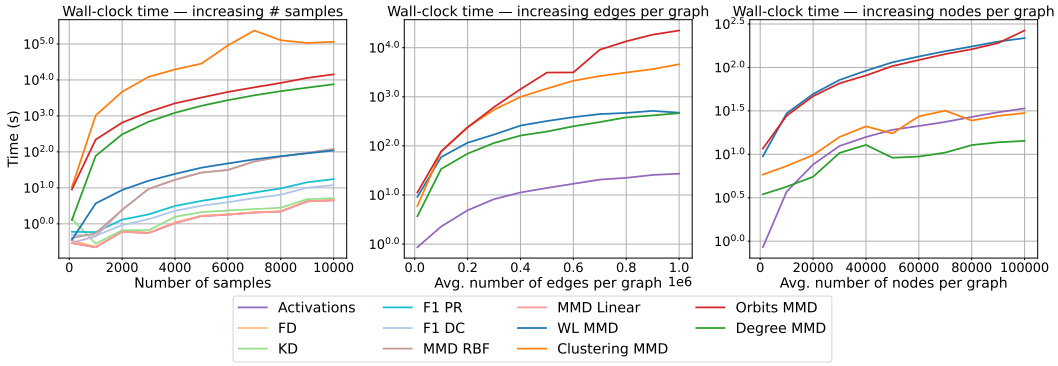

Figure 10: The wall-clock time of each metric as datasets scale in a single dimension. Activations is the time to extract graph embeddings from GIN on a **CPU**.

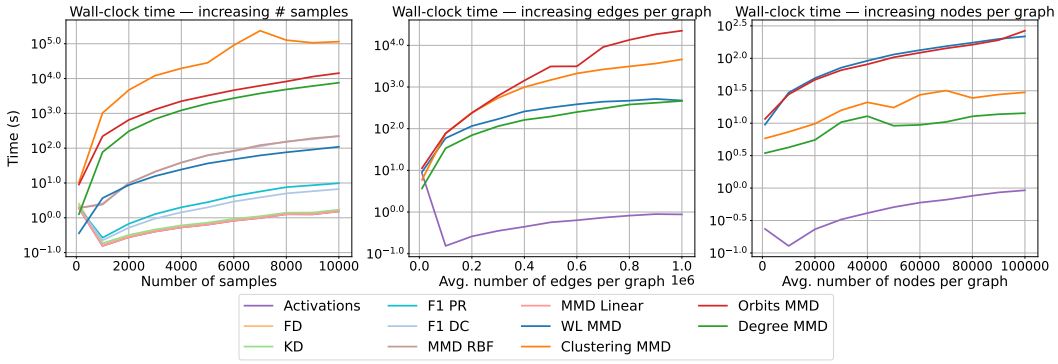

Figure 11: The wall-clock time of each metric as datasets scale in a single dimension. Activations is the time to extract graph embeddings from GIN on a **GPU**.

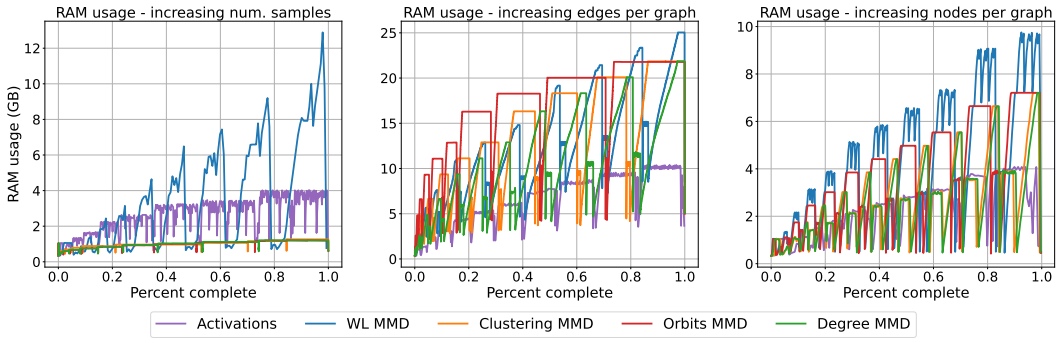

Figure 12: The estimated memory usage of each metric as the dataset scales in a single dimension. Activations is the memory required for all classifier-based metrics combined.

A.3.3 Computational efficiency

In Figure 12 we plot each metric's RAM usage throughout the computational efficiency experiments in Section 4.5 by tracking the memory usage of the main Python process. While this is likely imperfect and slightly noisy, it provides an idea of the memory requirements of each metric. Unfortunately all classifier-based metrics are grouped into the same experiment and more in-depth breakdowns are unavailable for this experiment. However, it is likely that the bulk of the memory requirements specific to these metrics come from storing the GIN parameters and extracted graph embeddings. The maximum recorded value for each metric in each experiment is recorded in Table 6.

Table 6: The maximum recorded memory usage in GB of each metric during each of the computational efficiency experiments.

METRIC	SCALING EXPERIMENT		
	Num. samples	Nodes	Edges
MMD Degree	1.20	7.20	21.88
MMD Clustering	1.26	7.20	21.85
MMD Orbits	1.19	7.20	21.80
MMD WL	12.89	9.75	25.05
Classifier metrics	4.01	4.08	10.40

A.3.4 GGM MODEL SELECTION

Here we evaluate GGMs at various stages of training on Grid, Lobster and Proteins datasets using classical metrics (You et al., 2018), NN-based metrics using a random GIN, and NN-based metrics using a pretrained GIN. The results are shown in Table 7.

Table 7: Evaluation of different GGMs at various percentages of total epochs trained on the Grid, Lobster, and Proteins datasets. Models are evaluated using three of the strongest NN-based metrics and classical metrics (You et al., 2018). All NN-based metrics are averaged across 10 different GINs with the strongest configuration. Cells are colored according to their rank in a given column. 50/50 split represents the metric computed using a random 50/50 split of the dataset and represents the theoretical ideal score for each metric.

(a) Random GIN, Grid dataset.

	MMD RBF	F1 PR	F1 DC	Clus.	Deg.	Orbit
50/50 split	0.042	0.998	0.995	0.0	$6.51e^{-5}$	0.018
Erdős-Rényi	0.203 ± 0.006	0.669 ± 0.041	0.89 ± 0.026	0.023	1.001	0.56
GraphRNN-100%	$0.184 \pm 5e^{-5}$	$0.919 \pm 1e-16$	0.896 ± 0.002	$4.59e^{-8}$	0.032	0.252
GraphRNN-66%	$0.154 \pm 7e^{-5}$	0.912 ± 0.007	$0.942 \pm 7e^{-4}$	$1.69e^{-6}$	0.015	0.169
GraphRNN-33%	$0.248 \pm 6e^{-5}$	$0.953 \pm 1e-16$	0.878 ± 0.003	$1.81e^{-6}$	0.022	0.134
GRAN-100%	0.063 ± 0.001	1.0 ± 0.0	1.073 ± 0.001	$1.24e^{-6}$	$1.68e^{-4}$	0.037
GRAN-66%	0.061 ± 0.001	1.0 ± 0.0	1.065 ± 0.002	$3.15e^{-6}$	$3.20e^{-5}$	0.047
GRAN-33%	$0.06 \pm 2e^{-4}$	0.982 ± 0.012	1.067 ± 0.005	$1.46e^{-6}$	$4.85e^{-5}$	0.054

(b) Pretrained GIN, Grid dataset.

	MMD RBF	F1 PR	F1 DC	Clus.	Deg.	Orbit
50/50 split	0.05	0.997	0.969	0.0	$6.51e^{-5}$	0.018
Erdős-Rényi	1.278 ± 0.05	0.0 ± 0.0	0.0 ± 0.0	0.023	1.001	0.56
GraphRNN-100%	0.21 ± 0.017	0.928 ± 0.021	0.82 ± 0.03	$4.59e^{-8}$	0.032	0.252
GraphRNN-66%	0.156 ± 0.012	0.907 ± 0.018	0.871 ± 0.033	$1.69e^{-6}$	0.015	0.169
GraphRNN-33%	0.176 ± 0.009	0.951 ± 0.015	0.848 ± 0.035	$1.81e^{-6}$	0.022	0.134
GRAN-100%	0.078 ± 0.016	0.98 ± 0.017	1.004 ± 0.05	$1.24e^{-6}$	$1.68e^{-4}$	0.037
GRAN-66%	0.066 ± 0.009	0.963 ± 0.032	0.994 ± 0.052	$3.15e^{-6}$	$3.20e^{-5}$	0.047
GRAN-33%	0.073 ± 0.015	0.953 ± 0.029	0.99 ± 0.052	$1.46e^{-6}$	$4.85e^{-5}$	0.054

(c) Random GIN, Lobster dataset

	MMD RBF	F1 PR	F1 DC	Clus.	Deg.	Orbit
50/50 split	0.049	0.989	0.987	0.0	0.003	0.037
Erdős-Rényi	0.538 ± 0.029	0.446 ± 0.126	0.554 ± 0.135	0.002	0.736	0.729
GraphRNN-100%	0.235 ± 0.005	0.979 ± 0.011	0.921 ± 0.026	0.095	0.038	0.104
GraphRNN-66%	0.244 ± 0.006	0.983 ± 0.008	0.914 ± 0.023	0.099	0.05	0.104
GraphRNN-33%	0.137 ± 0.01	0.939 ± 0.056	0.893 ± 0.031	0.137	0.047	0.104
GRAN-100%	0.146 ± 0.009	0.932 ± 0.041	0.905 ± 0.022	0.101	0.048	0.098
GRAN-66%	0.161 ± 0.01	0.956 ± 0.007	0.942 ± 0.022	0.103	0.045	0.094
GRAN-33%	0.096 ± 0.013	0.921 ± 0.054	0.907 ± 0.032	0.125	0.08	0.099

(d) Pretrained GIN, Lobster dataset

	MMD RBF	F1 PR	F1 DC	Clus.	Deg.	Orbit
50/50 split	0.043	0.97	0.977	0.0	0.003	0.037
Erdős-Rényi	0.676 ± 0.23	0.088 ± 0.085	0.077 ± 0.077	0.002	0.736	0.729
GraphRNN-100%	0.193 ± 0.026	0.7 ± 0.058	0.658 ± 0.037	0.095	0.038	0.104
GraphRNN-66%	0.222 ± 0.039	0.671 ± 0.06	0.582 ± 0.039	0.099	0.05	0.104
GraphRNN-33%	0.196 ± 0.033	0.652 ± 0.041	0.615 ± 0.043	0.137	0.047	0.104
GRAN-100%	0.153 ± 0.02	0.698 ± 0.04	0.661 ± 0.042	0.101	0.048	0.098
GRAN-66%	0.186 ± 0.025	0.68 ± 0.051	0.602 ± 0.032	0.103	0.045	0.094
GRAN-33%	0.194 ± 0.016	0.595 ± 0.035	0.565 ± 0.051	0.125	0.08	0.099

(e) Random GIN, Proteins dataset.

	MMD RBF	F1 PR	F1 DC	Clus.	Deg.	Orbit
50/50 split	0.005	0.988	0.988	$7.18e^{-4}$	$8.64e^{-4}$	0.004
Erdős-Rényi	0.733 ± 0.099	0.023 ± 0.011	0.01 ± 0.004	1.796	1.523	0.124
GRAN-100%	0.186 ± 0.065	0.756 ± 0.1	0.568 ± 0.114	0.28	0.542	0.099
GRAN-75%	0.09 ± 0.034	0.853 ± 0.055	0.819 ± 0.1	0.323	0.304	0.044
GRAN-50%	0.226 ± 0.071	0.66 ± 0.109	0.474 ± 0.111	0.382	0.646	0.103
GRAN-25%	0.019 ± 0.009	0.922 ± 0.004	0.995 ± 0.045	0.291	0.077	0.028

(f) Pretrained GIN, Proteins dataset

	MMD RBF	F1 PR	F1 DC	Clus.	Deg.	Orbit
50/50 split	0.005	0.973	0.987	$7.18e^{-4}$	$8.64e^{-4}$	0.004
Erdős-Rényi	1.097 ± 0.102	$4.34e^{-4} \pm 9e^{-4}$	$1.68e^{-4} \pm 3e^{-4}$	1.796	1.523	0.124
GRAN-100%	0.514 ± 0.083	0.561 ± 0.141	0.306 ± 0.085	0.28	0.542	0.099
GRAN-75%	0.304 ± 0.061	0.691 ± 0.074	0.488 ± 0.072	0.323	0.304	0.044
GRAN-50%	0.549 ± 0.086	0.463 ± 0.13	0.228 ± 0.076	0.382	0.646	0.103
GRAN-25%	0.066 ± 0.014	0.885 ± 0.022	0.83 ± 0.033	0.291	0.077	0.028