

Graph annotation generative adversarial networks

Yoann Boget

YOANN.BOGET@HESGE.CH

*University of Geneva and Geneva School for Business administration HES-SO
Rue de la Tambourine 17
Carouge, Switzerland*

Magda Gregorova

MAGDA.GREGOROVA@FHWS.DE

*Center for Artificial Intelligence and Robotics (CAIRO), FHWS
Franz-Horn-Strasse 2
Würzburg-Schweinfurt, Germany*

Alexandros Kalousis

ALEXANDR.KALOUSIS@HES-SO.CH

*Geneva School for Business administration HES-SO
Rue de la Tambourine 17
Carouge, Switzerland*

Editors: Emtiyaz Khan and Mehmet Gonen

Abstract

We consider the problem of modelling high-dimensional distributions and generating new examples of data with complex relational feature structure coherent with a graph skeleton. The model we propose tackles the problem of generating the data features constrained by the specific graph structure of each data point by splitting the task into two phases. In the first it models the distribution of features associated with the nodes of the given graph, in the second it complements the edge features conditionally on the node features. We follow the strategy of implicit distribution modelling via generative adversarial network (GAN) combined with permutation equivariant message passing architecture operating over the sets of nodes and edges. This enables generating the feature vectors of all the graph objects in one go (in 2 phases) as opposed to a much slower one-by-one generations of sequential models, prevents the need for expensive graph matching procedures usually needed for likelihood-based generative models, and uses efficiently the network capacity by being insensitive to the particular node ordering in the graph representation. To the best of our knowledge, this is the first method that models the feature distribution along the graph skeleton allowing for generations of annotated graphs with user specified structures. Our experiments demonstrate the ability of our model to learn complex structured distributions through quantitative evaluation over three annotated graph datasets.

Keywords: Graph, Annotation, Generative Model, GAN, adversarial.

1. Introduction

Modern deep learning approaches for learning high-dimensional data distributions and synthesizing new data examples have achieved great successes in a number of domains. We focus here on the particularly challenging problem of generating new data examples $\mathbf{x}_i \in \mathbb{R}^{m_i}$ with different dimensionalities m_i between individual instances and complex relational structures in the feature spaces organized in graphs.

More concretely, we focus on learning the distributions and generating new examples of features associated with nodes and edges of annotated graphs conditioned on the graph structure.

We formulate the problem as that of modeling a conditional distribution, where we learn to generate the node and edge features conditionally on the given graph skeleton (non-annotated graph). In this respect, our paper complements the existing tool-set of models for non-annotated graph generations (e.g. You et al. (2018); Liao et al. (2019)). The non-annotated graph can be generated by one of those or, more interestingly, can be provided by the user based on the needs of a downstream task that requires a specific graph structure (e.g. basis for scaffold-based *de novo* drug discovery or molecular docking, types and interactions of particles in high energy physic, individual features and types of relations in a social network).

The graph annotation generative adversarial networks (GrannGAN) method that we propose follows the strategy of implicit distribution modeling via adversarial training Goodfellow et al. (2014), allowing us to sample new graphs from a distribution p approximating the true p^* without explicitly formulating the p distribution function. In principle, the method can generate node or edge features independently. In practice, we generate the node and edge features in two phases depicted in Figure 1. In the first phase, our model samples the node features conditioning on the graph skeleton. In the second phase, it uses the generated node features as an additional conditional variable to sample the edge features.

The entire model architecture relies on the permutation equivariant transformations of message passing neural networks (MPNN) Gilmer et al.. Thanks to these, we circumvent the difficulties related to the ordering of graph objects in their representation. This critical property of the method allows for efficient use of the model capacity, which needs to learn neither a particular heuristic for unique representation ordering (as in the case of linearized representations) nor the complete set of equivalent permutations.

In the following, we describe the newly proposed GrannGAN method and provide some details of the technical implementation in sections 2. We then position the method within the existing state of research and document its performance competitive with the best of the well-established methods for graph generations on a set of experiments in section 3 and 4. We conclude with a discussion of possible future directions in section 5.

2. GrannGAN

Let $\mathcal{G} = \{\mathcal{V}, \mathcal{E}, V, E\}$ be an undirected graph with a set of nodes (vertices) \mathcal{V} and a set of edges \mathcal{E} between pairs of nodes in the graph, V and E are the corresponding node and edge features. Let ν_i denote a node i and $\mathbf{v}_i \in \mathbb{R}^d$ the features associated with that node. An edge ϵ_{ij} is connecting the pair of nodes ν_i and ν_j and $\mathbf{e}_{ij} \in \mathbb{R}^c$ are features corresponding to that edge. We consider only the case where the edges are undirected so that $\epsilon_{ij} = \epsilon_{ji}$ and $\mathbf{e}_{ij} = \mathbf{e}_{ji}$, but the model can be easily extended to the case with directed edges. We further use the term *skeleton* and the letter \mathcal{S} to refer to the non-annotated graph $\mathcal{S} = \{\mathcal{V}, \mathcal{E}\}$ corresponding to \mathcal{G} .

2.1. Model factorization

In our approach, we model the underlying graph distribution in the following factorization

$$p(\mathcal{G}) = p(\mathcal{S}, V, E) = p(\mathcal{S})p(V|\mathcal{S})p(E|\mathcal{S}, V) . \quad (1)$$

Our model consists of modeling $p(V|\mathcal{S})$ and $p(E|\mathcal{S}, V)$ with conditional GAN. Note that we could equivalently reverse the conditioning order and model $p(E|\mathcal{S})$ and $p(V|\mathcal{S}, E)$. In our experiments, we found the first option yielding better results. So, we keep this ordering for the following.

During training, we sample the graph skeleton from the data distribution $p^*(\mathcal{S})$. At inference, we can either sample the graph skeleton or keep it fixed depending on the task.

2.2. Implicit data generation

We use the Wasserstein-GAN (WGAN) [Arjovsky et al. \(2017\)](#) formulation of adversarial training with spectral normalization [Miyato et al. \(2018\)](#) in all linear layers of the critic to enforce 1-lipschitzness. To sample new examples of data \mathbf{x} from a model distribution $p_\theta(\mathbf{x})$ WGAN uses a generator g_θ mapping from a random latent variable \mathbf{z} to the output $\mathbf{x} = g_\theta(\mathbf{z})$. The generator is learned to minimize the Wasserstein-1 (or Earth-Mover) distance $W(p_\theta(\mathbf{x}), p^*(\mathbf{x}))$ between the implicit model distribution and the true generative distribution through a min-max optimization.

$$\min_{\theta} \max_{\varphi} \mathbb{E}_{(\mathbf{y}) \sim p^*} f_{\varphi}(\mathbf{y}) - \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} f_{\varphi}(g_{\theta}(\mathbf{z})) , \quad (2)$$

where f_{φ} is the K-Lipschitz critic function mapping from the data \mathbf{y} (real or generated) to a real-valued score $f_{\varphi}(\mathbf{y}) \in \mathbb{R}$.

We use the conditional WGAN formulation twice in the GrannGAN pipeline sketched out in figure 1. Following the factorization in equation (1) we first model the conditional distribution $p(V|\mathcal{S})$ in a *node-annotation* phase. Here, the method generates the node features $\hat{V} = g_{V,\theta}(\mathcal{Z}, \mathcal{S})$ from the set of latent noise variables \mathcal{Z} and the skeleton \mathcal{S} . So, equation 2 becomes

$$\min_{\theta} \max_{\varphi} \mathbb{E}_{(V, \mathcal{S}) \sim p^*} f_{V, \varphi}(V, \mathcal{S}) - \mathbb{E}_{\mathcal{z} \sim p(\mathcal{z}), \mathcal{S} \sim p^*} f_{V, \varphi}(g_{V, \theta}(\mathcal{Z}, \mathcal{S}), \mathcal{S}) . \quad (3)$$

Similarly in the following *edge-annotation* phase, we model the conditional $p(E|\mathcal{S}, V)$. The edge features $\hat{E} = g_{E,\theta}(\mathcal{Z}, V, \mathcal{S})$ are generated conditionally on the previously generated node features by including these as an additional input variable to the generator and critic. In this phase, equation 2 becomes

$$\min_{\theta} \max_{\varphi} \mathbb{E}_{(E, V, \mathcal{S}) \sim p^*} f_{E, \varphi}(E, V, \mathcal{S}) - \mathbb{E}_{\mathcal{z} \sim p(\mathcal{z}), (V, \mathcal{S}) \sim p^*} f_{E, \varphi}(g_{E, \theta}(\mathcal{Z}, V, \mathcal{S}), V, \mathcal{S}) , \quad (4)$$

Note that the conditioning on the skeleton $\mathcal{S} = (\mathcal{V}, \mathcal{E})$ acts on the computational graph of the critic f and of the generator g , i.e. on the structure of the Message Passing Neural Network (MPNN, explained in detail in the section 2.4). This special way of conditioning through the computational graph is one of the novelties of our method.

2.3. Implementation

In practice, we start by sampling the skeleton $\mathcal{S} = \{\mathcal{V}, \mathcal{E}\}$ from the data. For all node $\nu_i \in \mathcal{V}$, we sample a random noise vector \mathbf{z}_i . So, we get $\mathcal{Z} = \{\mathbf{z}_i\}_1^n$ and, by construction $|\mathcal{V}| = |\mathcal{Z}|$. Doing so, we obtain a graph with random noise vectors as initial latent representation of the node features $\mathcal{G}_0 = \{\mathcal{V}, \mathcal{E}, \mathcal{Z}\}$ (and without edge feature). This graph is the input of the *node-annotation* generator. Similarly, for the edge generation, we sample a graph from the data without edge feature. We also sample $\mathcal{Z} = \{\mathbf{z}_i\}_1^m$, where $m = |\mathcal{E}|$. The input of the *edge-annotation* generator is the graph $\mathcal{G}_0 = \{\mathcal{V}, \mathcal{E}, \mathcal{V}, \mathcal{Z}\}$ with latent random noise vectors as initial latent edge representation. The generators are permutation-equivariant functions. They output node or edge features following the order of the input.

The critics receive alternatively real or generated graphs (without edge feature during the annotation step). The critics are permutation-invariant functions. We provide details the generators and the critics in the next subsection.

The two steps (the two conditional GANs) are trained independently. Using teacher forcing, we use real data for the conditioning during training. For inference, we generate the edges features by conditioning on the previously generated node features. The Figure 1 illustrates the model architecture.

The software implementation of our method together with instructions for replicating our experiments are available at <https://github.com/yoboget/GrannGAN>. In this section we provide some details of the implementation to help the reader understand important design decisions.

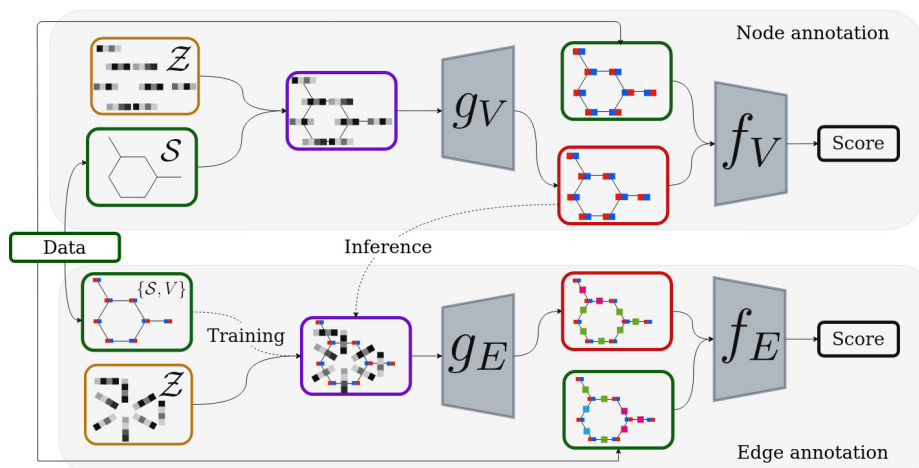


Figure 1: GrannGAN architecture: grayscale boxes represent random noise, colored boxes represent node and edge features.

2.4. Message passing neural network

In both the node- and the edge-annotation phases described in section 2.2 all GrannGAN generator g and critic f functions are message passing neural networks (MPNN) Gilmer et al.. As mentioned here-above, the node-neighbourhood structure of the MPNNs is derived

from the conditioning on the graph skeleton $\{\mathcal{V}\mathcal{E}\}$. Therefore, the skeleton acts on the i and j indices.

We use the following equations to perform L update steps over the hidden states $\mathbf{h}_i^{(l)}$ and $\mathbf{r}_{ij}^{(l)}$ of each node and edge in the graph respectively

$$\mathbf{r}_{ij}^{(l+1)} = \phi_r^{(l)}([\mathbf{h}_i^{(l)}, \mathbf{h}_j^{(l)}, \mathbf{r}_{ij}^{(l)}]) \quad (5)$$

$$\tilde{\mathbf{h}}_i^{(l+1)} = \sum_{j \in \mathcal{N}_i} \frac{1}{\sqrt{d_i} \sqrt{d_j}} \mathbf{r}_{ij}^{(l+1)} \quad (6)$$

$$\mathbf{h}_i^{(l+1)} = \phi_h^{(l)}([\mathbf{h}_i^{(l)}, \tilde{\mathbf{h}}_i^{(l+1)}]) . \quad (7)$$

Here ϕ are learned differentiable functions (we use small feedforward networks), \mathcal{N}_i is the first order neighbourhood of node \mathbf{v}_i as given by the graph skeleton \mathcal{S} , d_i is its degree, and $[\mathbf{a}, \mathbf{b}]$ is the concatenation of vectors \mathbf{a} and \mathbf{b} .

We outline the whole generative pipeline in figure 1. In the generator of the node-annotation phase $g_{\theta_{\mathcal{V}}}$, we initiate the node hidden states with a latent random noise vectors $\mathbf{h}_i^{(0)} = \mathbf{z}_i, \forall \mathbf{v}_i \in \mathcal{V}$ (and drop $\mathbf{r}_{ij}^{(0)}$ from inputs to $\phi_r^{(0)}$). The generated node features are the hidden states of the last update step $\hat{\mathbf{v}}_i = \mathbf{h}_i^{(L)}$. In contrast, in the edge-annotation generator $g_{\theta_{\mathcal{E}}}$ the latent random noise is used for initiating the edge hidden states $\mathbf{r}_{ij}^{(0)} = \mathbf{z}_{ij}, \forall \mathbf{e}_{ij} \in \mathcal{E}$ while the node hidden states are the node features generated from the previous phase $\mathbf{h}_i^{(0)} = \hat{\mathbf{v}}_i$. The generated edge features are the hidden edge states of the last update step $\hat{\mathbf{e}}_{ij} = \mathbf{r}_{ij}^{(L)}$.

The critic functions MPNNs are initiated from the real data examples as $\mathbf{h}_i^{(0)} = \mathbf{v}_i$ and $\mathbf{r}_{ij}^{(0)} = \mathbf{e}_{ij}$ or from the synthetic examples of the generators respectively. To produce the critic scores, the last-step update functions ($\phi_h^{(L)}$ for the node critic and $\phi_r^{(L)}$ for the edge critic) have scalar outputs that are averaged to enter the loss in equation (2). Therefore, the critic evaluates each node, aggregating information from all the nodes included in a radius of $2(L - 1)$. While this receptive field may not cover the whole graph, we assume that the node and edge features can be evaluated locally.

It has been shown Arvind et al. (2020) that MPNN cannot capture graph substructures other than forests of stars. To alleviate this issue, we further embody the graph topology into the MPNNs by extending the node representation by a set of skeleton-related features. Similarly to Bouritsas et al. (2020), we complement the node hidden states at the initial step $\mathbf{h}_i^{(0)}$ of all generators and critics by the node degree d_i and the number of k -cycles (cycles of length k the node is part of) extracted from the graph skeleton.

2.5. Graph representations

An important property of graph representation \mathcal{G} are their non-uniqueness. There are in general $n!$ possible permutations π determining the ordering of the nodes. The particular choice of the representation ordering from the complete set $\Pi = \{\pi_i\}_{i=1}^n$ is therefore another source of stochasticity so that

$$p(\mathcal{G}) = p\left(\bigcup_{\pi \in \Pi} \pi(\mathcal{G})\right) = \sum_{\pi \in \Pi} p(\pi(\mathcal{G})) . \quad (8)$$

From equation (8) we observe that when relying on the ordered graph representations \mathcal{G} , one shall in principle model the complete set of distributions $p(\pi(\mathcal{G}))$ for all the $n!$ permutations π to capture the unordered-set graph distribution $p(\mathcal{G})$. However, operating over the individual permutations $\pi(\mathcal{G})$ would lead to inefficient use of the model capacity. As an alternative, previous methods often fallback to operating over a single representation of each graph using some heuristic to fix the canonical ordering (such as the minimum depth-first-search Goyal et al. (2020), the unique breadth-first-search You et al. (2018) or the further sequentialed domain driven SMILES Daylight Chemical Information Systems (2011)). In result, the model has to learn not only the distribution of the graphs but also the ordering heuristic which is again inefficient from the perspective of generating the unordered graph sets \mathcal{G} .

A fundamental property of GrannGAN is that it is completely insensitive to the ordering $\pi(\mathcal{G})$ that can in result be chosen arbitrarily. This permutation equivariance¹ of the message passing operations is pivotal for our method as it solves the problem of the graph representation non-uniqueness while allowing for efficient use of the model capacity.

3. Related work

GrannGAN is closely related to the broad category of graph generative models, an area which has attracted significant attention of the research community resulting in a flurry of papers in the last several years. A systematic review of the major advancements in the field can be found for example in the excellent survey Faez et al. (2021).

Unlike previous work on that field, our method focuses on the particular problem of generating node and edge features conditionally over a given graph skeleton. To the best of our knowledge, we are the first to investigate such structural-based conditional graph feature generation.

The closest to our settings are methods modelling the distribution of annotated graphs. These can be categorized into two large families: the ones generating graphs sequentially and those generating in the whole graph in one go.

3.1. Sequential graph generation

The first are methods generating the graphs sequentially, starting from small structures to which they gradually connect new graph components (nodes, edges, or complete sub-structures). Most of these focus on generating chemical molecules and adopt various measures to improve the performance on this particular generative problem. CharacterVAE Gómez-Bombarelli et al. (2018) and GrammarVAE Kusner et al. (2017) rely directly on sequential SMILES representation of the molecules. JT-VAE Jin et al. (2018) operates over hand-crafted vocabulary of chemically valid sub-structures. MolecularRNN Popova et al. (2019) and GraphAF Shi et al. (2019) generate the graphs by successive node and edge sampling steps and ensure validity of the generated molecules by valency checking and the possibility for resampling at each of these steps.

Thanks to these domain-motivated components, the sequential methods often achieve excellent performance on molecular graphs datasets. However, those directly relying on

1. A function f is equivariant with respect to permutation π if $f(\pi(\mathbf{x})) = \pi(f(\mathbf{x}))$.

chemistry-specific data representations (e.g SMILES) cannot be easily extended to beyond the molecular problems. On the other hand, methods not using the SMILES representation need to establish their own heuristic for the sequential traversal of the graph such as the breadth-first-search in GraphRNN You et al. (2018) which in turn needs to be learned by the generative model together with the graph data distribution. As any other sequential models, the graph autoregressive generators need to be particularly careful about capturing long-term dependencies, and suffer from slow sampling process.

Different from the previous models NetGAN Bojchevski et al. (2018) uses GANs to learn the distribution of random walks over one big graph.

3.2. Annotated graph generation in one go

The other large family of models, into which GrannGAN can also be related, are those generating the graphs in one go, that is all the nodes and edges and their features together. GraphVAE Simonovsky and Komodakis (2018) proposes to sidestep the problem of graph linearization characteristic for sequential models by relying on the variational auto-encoder framework Kingma and Welling (2014). Due to modelling the graph in the ordered representation of annotation and adjacency matrix, it needs to employ an expensive (though inexact) graph-matching procedure in the training loss calculation, which significantly hampers its scalability. The flow-based graphNVP Madhawa et al. (2019) adopts the coupling strategy of Dinh et al. (2017) applying it to the rows of the graph annotation and adjacency matrices and thus preserving the ordering of the nodes and edges through the flow. GraphNVP is the only model for annotated graph using permutation equivariant generative function. In general, annotated graph generation in one go is still an open field of research. Our contribution can be seen as a proposition to tackle this challenging issue.

Misc-GAN Zhou et al. (2019) aims to translate graphs from a source-domain graph into a target-domain graph using a multiscale cycle-GAN. It operates on unannotated graphs.

The closest to our GrannGAN is the MolGAN De Cao and Kipf (2018) model, which uses GAN for molecule generation. Unlike in GrannGAN, the MolGAN generator is a feed-forward network sampling the ordered annotation and adjacency matrices of the graph representation. The authors of MolGAN observe that the model tends to suffer from mode collapse resulting in insufficient variation in the generated samples. We suppose that this issue comes from the loss of capacity by learning to generate various permutations of the same graph and from an additional Reinforcement Learning module encouraging the generator to produce graphs with some specific properties. We never experienced mode collapse during our experiments. As we show it in the experiment section, our model presents excellent uniqueness and novelty rates. We compare to MolGAN and other methods discussed here in our experiments in section 4.

4. Experiments

In this section, we present experiments using our model as a graph generative model. There currently exists no other models conditioning on the graph skeleton and, therefore, no other method to directly compare with. Instead, we use generative models presented in the previous section as baselines. However, we underline that these models do not have

the same purpose. In the second part of this section, we also present the results of the conditional generation by fixing the skeleton.

We evaluate our method in a set of experiments over three datasets containing (node- and edge-) annotated graphs². Two of the datasets, QM9 Ramakrishnan et al. (2014) and ZINC Sterling and Irwin (2015), are from the chemical domain that is frequently used as a test bed for annotated graph modeling. The third is the fingerprint dataset Riesen and Bunke (2008) included in the TUDataset collection Morris et al. (2020).

As it has recently been discussed, for example in O’Bray et al. (2021) and Thompson et al. (2021), evaluating generative models of graph structures is particularly challenging due to the impracticality of visual or perceptual comparisons of the data examples. The comparison with existing graph generative models is therefore difficult and can provide only crude indication of the method capabilities.

For the chemical datasets the most common evaluation metrics are the following

- *validity* is the proportion of chemically valid molecules in the total generated examples and measures the ability of the method to understand the chemical constraints differentiating general graphs from chemically valid molecules
- *uniqueness* is the proportion of unique samples within the *valid* generated samples and measures the variability of the generated data
- *novelty* is the proportion of examples within the *valid* and *unique* set that do not exist in the training dataset and measures the ability of the method to go beyond data memorization

We complement these by an **overall** score of *valid-unique-novel* molecules calculated simply as the product of the three above metrics and measuring the overall quality of the generated examples as a proportion of samples with all three of these desirable properties in the set of generated data. So, the overall score gives the rate of valid molecules that are neither in the dataset nor already generated. Note that in most cases the overall score is the metric of interest.

While these evaluation metrics may be useful indicators for chemical downstream tasks, they do not provide any measure of the distance between the distribution of reference and the generated distributions. Various indicators calling on the framework of maximum mean discrepancy (MMD) Gretton et al. (2012) have been proposed in the literature. These are used very inconsistently and, as demonstrated in O’Bray et al. (2021), they are highly sensitive to the specific choice of the graph statistics, the kernel and the hyperparameters used for the MMD calculation. Instead of calculating the MMDs on graphs statistics, we report directly the distance between the generated feature distributions, using the Jensen-Shannon Distance for each feature. We show that with respect to these metrics, our model outperforms by far GraphAF, model considered as the state-of-the-art.

4.1. QM9

QM9 is one of the most commonly used datasets for testing models for annotated graph generation. It consists of ~ 134 k stable organic molecules with up to 9 atoms of four types.

2. Unfortunately, we did not find other public dataset with enough node- and edge-annotated graphs.

The 4 atom types and 4 bond types are encoded as one-hot vectors in the node feature descriptions \mathbf{v} .

In table 1 we present an overview of the generative performance of our GrannGAN method in comparison with a set of well-established graph generative methods: MolGAN De Cao and Kipf (2018), GraphVAE Simonovsky and Komodakis (2018), GraphNVP Madhawa et al. (2019), GraphAF Shi et al. (2019), CharacterVAE Gómez-Bombarelli et al. (2018), GrammarVAE Kusner et al. (2017), and JT-VAE Jin et al. (2018).

The GrannGAN results are calculated from 1000 new data examples generated by conditioning on skeletons randomly sampled from the training data. The results of the baseline methods are those reported by the authors in the original papers³. Despite the frailty of such comparisons due to inconsistencies in the original experimental protocols (for example, some of the methods, such as GraphAF, work over the *kekulized* versions of the molecules reducing the number of edge categories to three), our GrannGAN achieves excellent results. It outperforms all one-go generative models and is on par with the best auto-regressive model.

Table 1: QM9: performance comparison

| | Model | valid | unique | novel | overall |
|--------------|------------------------|-------|--------|-------|---------|
| no chemistry | GrannGAN | 82.5 | 99.9 | 64.9 | 53.4 |
| | MolGAN (wo. RL) | 87.7 | 2.9 | 97.7 | 2.5 |
| | GraphVAE | 55.7 | 76.0 | 61.6 | 26.1 |
| | GraphNVP | 83.1 | 58.2 | 99.2 | 46.8 |
| | GraphAF (wo. validity) | 67.0 | 94.5 | 88.8 | 56.3 |
| chemistry | MolGAN-RL | 99.8 | 2.3 | 97.9 | 2.2 |
| | CharacterVAE | 10.3 | 90.0 | 67.5 | 11.9 |
| | GrammarVAE | 60.2 | 80.9 | 9.3 | 11.9 |
| | GraphAF | 100 | 94.5 | 88.8 | 83.9 |

As explained above, the metrics in the table are very much domain specific. Methods in the lower part of table 1 expressly focus on these introducing into the models various expert-designed modules promoting chemical validity of the generated molecular graphs (e.g., valid sub-structures vocabulary, valency constraints, rejection sampling, etc.). While these improve the targeted metrics in table 1, they also bias the generative process and, therefore, the model distribution towards the validity metric in exchange for the distribution approximation goal. We list these methods in table 1 for completeness despite them not being really comparable to our approach which focuses on the distribution approximation of general graphs.

Figure 2 documents the ability of the GrannGAN method to learn the feature distribution by comparing the empirical real data distribution with the distribution of the generated examples. We have replicated the results of the GraphAF method, which is the most competitive with GrannGAN according to table 1), to extract similar statistics. GrannGAN

3. CharacterVAE and GrammarVAE results (not reported in the original papers) are taken over from the method replications in GraphNVP.




Figure 2: QM9: empirical distributions of real and generated features (samples of 1000 graphs)

Table 2: Jensen-Shannon distances between real and generated distributions for node and edge features of the QM9 and ZINC datasets (the lower the better)

| Model | nodes - QM9 | edges - QM9 | nodes - ZINC | edges - ZINC |
|----------|-----------------------|----------------------|----------------------|----------------------|
| GrannGAN | $1.87 \cdot 10^{-3}$ | $0.04 \cdot 10^{-4}$ | $1.50 \cdot 10^{-3}$ | $0.60 \cdot 10^{-3}$ |
| GraphAF | $15.80 \cdot 10^{-3}$ | $11.5 \cdot 10^{-4}$ | $15.5 \cdot 10^{-3}$ | $3.39 \cdot 10^{-3}$ |

matches the feature distributions of the reference dataset almost exactly while we observe a less closer match in the GraphAF⁴ generated data. This observation is numerically confirmed in table 2 listing the Jensen-Shannon distances (JSD) between the real and generated data samples, where GrannGAN is systematically better by one order of magnitude.

4.2. ZINC

We use here the ZINC250k version used previously for generative modeling evaluations. It consists of 250k randomly selected molecules from the complete ZINC set with up to 38 atoms of 9 types and 4 bond types.

Table 3 summarizes the generative performance of GrannGAN over 1000 newly synthesized examples compared to results reported in the literature. MolGAN and GraphVAE did not experiment (or did not provide the results) on this more challenging dataset of larger graphs. GrannGAN scales to this dataset easily and performs competitively with respect to the chemically motivated metrics (keeping in mind the difficulties of such comparisons related to differences in experimental protocols and evaluation procedures).

Figure 3 documents the excellent performance of GrannGAN in learning the feature distributions, matching it closely for both the node and edge features. The differences in

4. GraphAF uses somewhat different training set than we. For example, by kekulizing the molecules it removes all aromatic edges and thus reduces the number of edge types to 3. This is true also for the ZINC dataset.

the true and GraphAF generated data distributions are more pronounced as is also clear from the numerical evaluation via the JSD presented in table 2.

Table 3: ZINC: performance comparison

| Model | | valid | unique | novel | overall |
|--------------|------------------------|-------|--------|-------|---------|
| no chemistry | GrannGAN | 56.5 | 100 | 100 | 56.5 |
| | MolGAN (wo. RL) | — | — | — | — |
| | GraphVAE | 13.5 | — | — | — |
| | GraphNVP | 42.6 | 94.8 | 100 | 40.4 |
| | GraphAF (wo. validity) | 68.0 | 99.1 | 100 | 67.3 |
| chemistry | MolGAN-RL | — | — | — | — |
| | CharacterVAE | 7.2 | 9.0 | 100 | 0.6 |
| | GrammarVAE | 0.7 | 67.5 | 100 | 0.5 |
| | JT-VAE | 100 | 100 | 100 | 100 |
| | GraphAF | 100 | 99.1 | 100 | 99.1 |




Figure 3: ZINC: empirical distributions of real and generated features (samples of 1000 graphs)

4.3. Fingerprints

Fingerprint is a dataset of graphs representing the skeletonized regions of interest in human fingerprints [Riesen and Bunke \(2008\)](#). The graphs consist of up to 26 nodes. Both nodes and edges are described by 2 continuous features related to their position and orientation in the 2d fingerprint image. We have linearly re-scaled both node and edge features between -1 and 1.

We use the non-annotated graphs as the skeletons for new feature generations. Each row displays graphs arranged according to the new features generated by GrannGAN.

In table 4 we list the JS distance (with the logarithm in base 2) between the empirical distributions of the training data and newly generated features starting from 500 randomly selected skeletons. These have been calculated by discretizing the features into 200 even




Figure 4: Fingerprints: empirical distributions \mathbf{v}_1 , \mathbf{v}_2 and \mathbf{d} sampled from our GrannGAN and from the dataset (samples of 1000 graphs)




Figure 5: Fingerprints: empirical distributions \mathbf{e}_1 and \mathbf{e}_2 sampled from our GrannGAN and from the dataset (samples of 1000 graphs)

sized bins. The columns in the table correspond to the two node \mathbf{v} and edge \mathbf{e} features respectively. Following the same methodology, we also report the JS distance in the distribution of the euclidean distances between connected nodes \mathbf{d} . Unlike the feature distribution, this statistic consider the relation between node features. In figures 4 and 5, we show the node and edge features distribution for the data and newly generated examples from our model as well as the distribution of distances between node \mathbf{d} ⁵.

Table 4: Jensen-Shannon distances between real and generated distributions for node and edge features of the Fingerprint dataset (the lower the better)

| | $\mathbf{v1}$ | $\mathbf{v2}$ | \mathbf{d} | $\mathbf{e1}$ | $\mathbf{e2}$ |
|---------------------|----------------------|----------------------|-----------------------|-----------------------|----------------------|
| GrannGAN | $5.02 \cdot 10^{-2}$ | $4.55 \cdot 10^{-2}$ | $32.84 \cdot 10^{-2}$ | $16.78 \cdot 10^{-2}$ | $9.74 \cdot 10^{-2}$ |
| Data samples | $1.13 \cdot 10^{-2}$ | $1.27 \cdot 10^{-2}$ | $1.49 \cdot 10^{-2}$ | $1.37 \cdot 10^{-2}$ | $1.07 \cdot 10^{-2}$ |

Again, our model shows excellent performance in closely matching the node and edge features and their relations.

These experiments show the quality of our GrannGAN for the generative task. It outperforms all previous one-go generative models. The sequential GraphAF is the only model presenting slightly better results using the molecular metrics. However, we show that our model is much better at capturing the feature distributions.

5. We assume that the oscillation of the curve is just an artifact coming from the relatively small data points by bin

4.4. Conditional generation

In this section, we show that our model can be used for conditional generation by fixing the skeleton and still producing a high level of diversity.

We want to guarantee that the sample diversity does not come mainly from the skeleton. To test the generative capacity of our GrannGAN, we randomly sampled 100 skeletons from the ZINC dataset. We generate 1000 new instances conditioned on each sampled skeleton. We compute the average rate of uniqueness, i.e. the rate of unique instances over all the valid molecules.

Table 5 reports the results of the experiment. As expected, the validity rate is similar to the one by sampling a new skeleton for each instance. Even by fixing the skeleton, we reached a rate of 92.5% unique instances on average. On average, the conditional generation of 1000 instances over the same skeleton produces 527.0 unique molecules. So, we show that the performance of our model does not depend critically on the sampling of multiple skeletons and can be properly used as a conditional generative model.

Table 5: Conditional generation of 1000 instances from the same skeleton (average)

| Model | Validity (%) | Uniqueness (%) | Valid and unique (‰) |
|----------|--------------|----------------|----------------------|
| GrannGAN | 57.0 | 92.5 | 527.0 |

We have demonstrated that the GrannGAN is competitive with state-of-the-art model for graph generation even though it is first meant for conditional generation. GraphAF is slightly better on the molecular metric, but much worse at matching the feature distribution. In addition, our model generates instances in one go and is therefore much faster during inference. More fundamentally, our model preserves its quality while generating new instances conditioned on the skeleton, a task that no other model can do.

5. Conclusions

We have presented here a new method, GrannGAN, for implicit distribution learning and generating features of nodes and edges conditioned on graph skeletons. The promising results of the method, as documented in our experiments, indicate that the method can learn the complex high-dimensional distributions and generate new data examples with features coherent with the underlying graph structure. When the generation of new skeletons is not crucial, GrannGAN can be used as a full generative model by sampling the skeleton from the data. It is also the first method for graph generation conditioned on the skeleton.

These favorable results provide a solid starting point for multiple possible extensions in future work. For example, in some applications (e.g. bio-chemistry) a reasonable starting point for the feature generations may be a partly annotated graph (in place of a completely non-annotated graph). Our model, which is already conditional, could be easily adapted to this setting by re-using the ideas of our second edge-annotation phase. Another interesting extension is that of property optimization. An additional conditioning variable could be included in the generators and critics to guide the model in generating property-conditioned features.

Acknowledgments

The contribution of Magda Gregorova has been supported by High-Tech Agenda Bayern.

References

- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein Generative Adversarial Networks. In *Proceedings of the 34th International Conference on Machine Learning*, pages 214–223. PMLR, July 2017. URL <https://proceedings.mlr.press/v70/arjovsky17a.html>. ISSN: 2640-3498.
- V. Arvind, Frank Fuhlbrück, Johannes Köbler, and Oleg Verbitsky. On Weisfeiler-Leman invariance: Subgraph counts and related graph properties. *Journal of Computer and System Sciences*, 113:42–59, nov 2020. ISSN 10902724.
- Aleksandar Bojchevski, Oleksandr Shchur, Daniel Zügner, and Stephan Günnemann. NetGAN: Generating graphs via random walks. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 610–619. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/bojchevski18a.html>.
- Giorgos Bouritsas, Fabrizio Frasca, Stefanos Zafeiriou, and Michael M. Bronstein. Improving Graph Neural Network Expressivity via Subgraph Isomorphism Counting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, jun 2020. URL <http://arxiv.org/abs/2006.09252>.
- Inc. Daylight Chemical Information Systems. Daylight Theory Manual, August 2011. URL <https://www.daylight.com/dayhtml/doc/theory/>.
- Nicola De Cao and Thomas Kipf. MolGAN: An implicit generative model for small molecular graphs. *arXiv:1805.11973 [cs, stat]*, May 2018. URL <http://arxiv.org/abs/1805.11973>. arXiv: 1805.11973.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=HkpbnH9lx>.
- Faezeh Faez, Yassaman Omimi, Mahdieh Soleymani Baghshah, and Hamid R. Rabiee. Deep Graph Generators: A Survey. *IEEE Access*, 9:106675–106702, 2021. ISSN 2169-3536. Conference Name: IEEE Access.
- Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1263–1272. PMLR. URL <https://proceedings.mlr.press/v70/gilmer17a.html>. ISSN: 2640-3498.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial

- Nets. In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. URL <https://papers.nips.cc/paper/2014/hash/5ca3e9b122f61f8f06494c97b1afccf3-Abstract.html>.
- Nikhil Goyal, Harsh Vardhan Jain, and Sayan Ranu. GraphGen: A Scalable Approach to Domain-agnostic Labeled Graph Generation. 2020. URL <https://doi.org/10.1145/3366423.3380201>.
- Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A Kernel Two-Sample Test. *Journal of Machine Learning Research*, 13(25): 723–773, 2012. ISSN 1533-7928. URL <http://jmlr.org/papers/v13/gretton12a.html>.
- Rafael Gómez-Bombarelli, Jennifer N. Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D. Hirzel, Ryan P. Adams, and Alán Aspuru-Guzik. Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules. *ACS Central Science*, 4(2):268–276, February 2018. ISSN 2374-7943. URL <https://doi.org/10.1021/acscentsci.7b00572>. Publisher: American Chemical Society.
- Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction Tree Variational Autoencoder for Molecular Graph Generation. In *Proceedings of the 35th International Conference on Machine Learning*, pages 2323–2332. PMLR, July 2018. URL <https://proceedings.mlr.press/v80/jin18a.html>. ISSN: 2640-3498.
- Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and Yann LeCun, editors, *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. URL <http://arxiv.org/abs/1312.6114>.
- Matt J. Kusner, Brooks Paige, and José Miguel Hernández-Lobato. Grammar Variational Autoencoder. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1945–1954. PMLR, July 2017. URL <https://proceedings.mlr.press/v70/kusner17a.html>. ISSN: 2640-3498.
- Renjie Liao, Yujia Li, Yang Song, Shenlong Wang, Will Hamilton, David K Duvenaud, Raquel Urtasun, and Richard Zemel. Efficient Graph Generation with Graph Recurrent Attention Networks. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/d0921d442ee91b896ad95059d13df618-Abstract.html>.
- Kaushalya Madhawa, Katsuhiko Ishiguro, Kosuke Nakago, and Motoki Abe. GraphNVP: an Invertible Flow-based Model for Generating Molecular Graphs. September 2019. URL <https://openreview.net/forum?id=ryxQ6T4YwB>.
- Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral Normalization for Generative Adversarial Networks. February 2018. URL <https://openreview.net/forum?id=B1QRgziT-&foo=bar>.

- Christopher Morris, Nils M. Kriege, Franka Bause, Kristian Kersting, Petra Mutzel, and Marion Neumann. Tudataset: A collection of benchmark datasets for learning with graphs. In *ICML 2020 Workshop on Graph Representation Learning and Beyond (GRL+ 2020)*, 2020. URL www.graphlearning.io.
- Leslie O’Bray, Max Horn, Bastian Rieck, and Karsten Borgwardt. Evaluation Metrics for Graph Generative Models: Problems, Pitfalls, and Practical Solutions. September 2021. URL <https://openreview.net/forum?id=tBtoZYKd9n>.
- Mariya Popova, Mykhailo Shvets, Junier Oliva, and Olexandr Isayev. MolecularRNN: Generating realistic molecular graphs with optimized properties. *arXiv:1905.13372 [cs, q-bio, stat]*, May 2019. URL <http://arxiv.org/abs/1905.13372>. arXiv: 1905.13372.
- Raghunathan Ramakrishnan, Pavlo O. Dral, Matthias Rupp, and O. Anatole von Lilienfeld. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific Data*, 1(1):140022, August 2014. ISSN 2052-4463. URL <https://www.nature.com/articles/sdata201422>. Number: 1 Publisher: Nature Publishing Group.
- Kaspar Riesen and Horst Bunke. IAM graph database repository for graph based pattern recognition and machine learning. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 5342 LNCS, pages 287–297. Springer, Berlin, Heidelberg, 2008. ISBN 3540896880. URL <http://www.iam.unibe.ch/fki/databases/iam-graph-database>.
- Chence Shi, Minkai Xu, Zhaocheng Zhu, Weinan Zhang, Ming Zhang, and Jian Tang. GraphAF: a Flow-based Autoregressive Model for Molecular Graph Generation. September 2019. URL <https://openreview.net/forum?id=S1esMkHYPr>.
- Martin Simonovsky and Nikos Komodakis. GraphVAE: Towards Generation of Small Graphs Using Variational Autoencoders. *arXiv:1802.03480 [cs]*, February 2018. URL <http://arxiv.org/abs/1802.03480>. arXiv: 1802.03480.
- Teague Sterling and John J. Irwin. ZINC 15 – Ligand Discovery for Everyone. *Journal of Chemical Information and Modeling*, 55(11):2324–2337, November 2015. ISSN 1549-9596. URL <https://doi.org/10.1021/acs.jcim.5b00559>. Publisher: American Chemical Society.
- Rylee Thompson, Boris Knyazev, Elahe Ghalebi, Jungtaek Kim, and Graham W. Taylor. On Evaluation Metrics for Graph Generative Models. September 2021. URL <https://openreview.net/forum?id=EnwCZixjSh>.
- Jiaxuan You, Rex Ying, Xiang Ren, William Hamilton, and Jure Leskovec. GraphRNN: Generating Realistic Graphs with Deep Auto-regressive Models. In *Proceedings of the 35th International Conference on Machine Learning*, pages 5708–5717. PMLR, July 2018. URL <https://proceedings.mlr.press/v80/you18a.html>. ISSN: 2640-3498.
- Dawei Zhou, Lecheng Zheng, Jiejun Xu, and Jingrui He. Misc-gan: A multi-scale generative model for graphs. *Frontiers in Big Data*, 2:3, 4 2019. ISSN 2624909X.