

1 We thank all the reviewers for their constructive comments. We first explain the update of our model which resolves
 2 some of R2’s questions. Then we address the common concern on the large scale graph generation followed by
 3 individual responses. We will release our PyTorch implementation upon acceptance.

4 **Model Simplification & Extension:** After submission, we tried [removing the RNN, thus simplifying our model to a](#)
 5 [single GNN](#). With more layers and slightly more training epochs, a single GNN performs similarly to the model with
 6 an RNN. Moreover, we [extended the output distribution from a single Bernoulli distribution \(Eq.7 in the paper\) to a](#)
 7 [mixture of Bernoullis](#), which can further model the dependency between the entries in a block without sacrificing the
 8 speed, and achieves state-of-the-art performance on most metrics of Protein dataset (*c.f.* Table 1).

Protein Dataset	Deg.	Clus.	Orbit	Spec.	SBM Graphs	Deg.	Clus.	Orbit
GraphVAE	$6.12e^{-2}$	0.17	0.82	$9.28e^{-2}$	$G(5000, 0.05)$	0.78	1.93	0.37
GraphRNN	$5.51e^{-3}$	0.15	0.40	$2.56e^{-2}$	$G(5000, 0.1)$	0.78	1.83	0.37
GRAN-50	$1.08e^{-2}$	0.08	0.30	$8.97e^{-3}$	GRAN-20	0.54	0.99	0.34

Table 1: GRAN-X: our model (RNN removed) with a mixture of X Bernoulli distributions. $G(n, p)$: Erdos Renyi model with number of nodes n and edge probability p .

9 **Large Scale Graph Generation [R2, R3]:** We experimented with random graphs from stochastic block models
 10 (SBM) to show that our model is able to generate large graphs. In particular, we generated 10 random graphs from an
 11 SBM with 3 communities, 5000 nodes and inter/intra-community probabilities $[[0.2, 0.001, 0.002], [0.001, 0.15, 0.004],$
 12 $[0.002, 0.004, 0.1]]$. We randomly choose 8 graphs for training and test our model on the remaining 2 graphs. The
 13 performances of our model and several Erdos-Renyi random graph baselines are shown in Table 1. The run-time of
 14 generating 3 graphs (5K-nodes) in parallel on a GeForce 1080Ti GPU is around 279s. We experimented with the public
 15 implementation of GraphRNN which runs out-of-memory under this graph size. We are investigating the suggested
 16 real-world graphs and will include the results in the final version.

17 **How to determine graph size [R2, R5]:** We set a maximum graph size which determines when the GNN stops. Then
 18 we sample the graph size from the empirical distribution of graph sizes (collected from the training set). Finally, we
 19 crop the generated adjacency matrix from top-left using this sampled graph size. Although learning to predict the graph
 20 size is arguably better, we found this simple sampling strategy works well in practice.

21 **To R1:**

22 - **Evaluation.** We first compute the statistics, *e.g.*, degree distribution, of the generated and observed graphs respectively.
 23 Then we measure the similarity between them using the maximum mean discrepancy (MMD) with a Gaussian kernel
 24 using total variation distance. For spectra, *i.e.*, eigenvalue distributions, we discretize the range and treat it as a
 25 histogram. In general, it is hard to measure differences using a single metric, as each statistic is just one particular
 26 summarization of the graph. Low MMD scores are necessary but not sufficient for two graphs to be similar. In our
 27 experience, small differences of MMD scores in degree distribution and spectrum, *e.g.*, 0.01 vs. 0.02, seem not to imply
 28 a significant difference in the quality of the visualized graphs whereas clustering coefficients and orbits correlate better
 29 with the visual quality. We will discuss this point further, and show examples along with the metrics.

30 **To R2:**

31 - **How Eq. (11) is trained?** $q(\pi|G)$ is a 2-layer GCN which takes rows in the adjacency matrix as the input node
 32 representation. We also tried random node features and obtained similar performance. There is no ground truth for
 33 training this network; the only training signal comes from maximizing the ELBO.

34 - **Maximum graph size.** Please refer to the section above. - **a in Eq.(5).** They are scalars and normalized by the
 35 sigmoid. - **Remove RNN.** We have since done this; see above. - **Overall Training.** It is trained by teacher forcing. All
 36 subgraphs are trained in parallel as in a pixel CNN. - **Model Size.** On the protein dataset, the size of the best-performing
 37 GraphVAE vs. GraphRNN vs. GRAN-20 is roughly 128MB vs. 2MB vs. 4MB. We will include it in the final version.

38 **To R3:** Please see our response above on large scale graph generation.

39 **To R5:**

40 - **Multiple splits.** Thanks for the suggestion! We will cite that paper and add the results in the final version. -
 41 **Optimization of the GNN.** We used the Adam optimizer with learning rate $1.0e-2$ and decay by 0.3 at 100, 500, and
 42 1000 epochs. The norm of the gradient is clipped at 5. We will discuss it in the final version. - **Long range dependency.**
 43 We believe our model with the GNN and less number of generation steps do have a positive impact on this. But it is not
 44 easy to demonstrate this explicitly through an experiment. - **Graph size.** Please refer to the graph size section above.