
Disentangling Interpretable Generative Parameters of Random and Real-World Graphs

Niklas Stoehr, Emine Yilmaz
University College London
n.stoehr@outlook.com
emine.yilmaz@ucl.ac.uk

Marc Brockschmidt, Jan Stuehmer
Microsoft Research Cambridge
mabrocks@microsoft.com
jastueh@microsoft.com

Abstract

While a wide range of interpretable generative procedures for graphs exist, matching observed graph topologies with such procedures and choices for its parameters remains an open problem. Devising generative models that closely reproduce real-world graphs requires domain knowledge and time-consuming simulation. While existing deep learning approaches rely on less manual modelling, they offer little interpretability. This work approaches graph generation (decoding) as the inverse of graph compression (encoding). We show that in a disentanglement-focused deep autoencoding framework, specifically β -Variational Autoencoders (β -VAE), choices of generative procedures and their parameters arise naturally in the latent space. Our model is capable of learning disentangled, interpretable latent variables that represent the generative parameters of procedurally generated random graphs and real-world graphs. The degree of disentanglement is quantitatively measured using the *Mutual Information Gap* (MIG). When training our β -VAE model on *ER random graphs*, its latent variables have a near one-to-one mapping to the ER random graph parameters n and p . We deploy the model to analyse the correlation between graph topology and node attributes measuring their mutual dependence without handpicking topological properties. To allow experimenting with the code, we provide an interactive notebook¹.

1 Introduction

Motivation and Related Work Conventional network analysis aims at finding interpretable models that explain interaction dynamics by examining graphs as discrete objects [1]. Random graph generator models [2] like *Erdős–Rényi random graphs* (ER graphs) [3] are usually too generic to accurately represent the versatile linking patterns of real-world graphs [2, 4, 5]. Devising models that reproduce characteristic topologies prevalent in social [6], biological [7], internet [8] or document [9] graphs typically requires a thorough understanding of the domain and time-consuming graph simulations, thereby imposing strong assumptions and modelling bias. Recently, deep learning on non-euclidean data such as graphs has received substantial attention [10]. As these techniques require little or no explicit modelling and capture complex graph structure [11, 12], we propose to use them as a tool to obtain interpretable generative parameters of graphs. As a limiting factor, most existing models generate graphs sequentially based on concatenations of node embeddings. These are not only non-interpretable but also impose an artificial node ordering instead of considering a global representation of the entire graph [13–18]. *DisenGCN* [19] focuses on interpretability, but is limited to node-level linking mechanisms. The latent space of *NetGAN* [20] reveals topological properties instead of generative parameters. Some recent works on interpretable graph embeddings [12, 21–24] provide visualizations for inspection, but no parameters suitable for a generative model.

¹https://colab.research.google.com/drive/1M--YX4d0St3imDPdecPbjVX-T6Ae0_0G




Figure 1: **Architecture Overview:** We seek a continuous function h mapping the disentangled latent variables z_j into mutually independent, interpretable generative parameters v_k .

In other domains, interest in model interpretability has caused a focus on the latent space of neural models [25]. Intuitively, the aim is to shape the latent space such that the euclidean distance between the latent representations of two data points corresponds to a “distance” between the actual data points [26]. Latent variables describe probability distributions over the latent space. The goal of latent variable disentanglement can be understood as wanting to use each latent variable to encode one and only one data property in a one-to-one mapping [27], making the latent space more interpretable. Varying one latent variable should then correspond to a change in one observable factor of variation in the data, while other factors remain invariant [26]. Most work in this field has been focused on visual and sequential data [27–30].

Contributions We assume that graphs are generated by superposition of interpretable, generative procedures parameterized by generative parameters v_k such as n and p in ER graphs. We hypothesize that these generative parameters v_k can be encoded by a minimal set of disentangled latent variables z_j in an unsupervised machine learning model. To this end, we apply the idea of β -Variational Autoencoders (β -VAE) [31] in the context of graphs. Intuitively, our autoencoder tries to compress (encode) a graph into a latent variable representation suitable for generating (decoding) it back into the original graph as outlined in figure 1. If the number of latent variables is lower than the dimensionality of the input data, they force a compressed representation that prioritizes the most salient data properties. In this article, we

- (1) discuss how to adapt the β -VAE model to graphs in section 2,
- (2) apply it to recover parameters for topology-generating procedures in section 3.1, and
- (3) leverage it to quantify dependencies between graph topology and node attributes in section 3.2.

2 Model

We instantiate the idea of β -VAEs [31] with graph-specific encoders and decoders. Our encoder model $q_\phi(z | x)$ is a *Graph Convolutional Network (GCN)* [32] and the decoder $p_\theta(x | z)$ is a deconvolutional neural network. Hence, in our setting the encoder is operating on the graph structure, whereas the decoder produces a graph by computing an adjacency matrix. We train this autoencoder in the β -VAE setting, in which the loss to minimize is $\mathbb{E}_{z \sim q_\phi(z|x)} [\log p_\theta(x | z)] + \beta (\text{KL}(q_\phi(z | x) \parallel p(z)))$. In the loss term, the reconstruction loss is balanced with the KL regularization term using a parameter $\beta \geq 1$. A higher value of β yields stricter alignment to the Gaussian prior $p(z) = \mathcal{N}(0, 1)$, leading to an orthogonalization of the encoding in z [27–29]. To further enforce disentangled representations of v_k , we attach an additional *parameter decoder* h to the latent space that learns a direct mapping $h(z) \rightarrow v$ between latent variables z_j and generative parameters v_k . If h is implemented as a linear mapping, the latent space needs to align with the generative parameters v_k , hence further favoring the result of the encoder $q_\phi(z | x)$ to be disentangled. If the latent space is perfectly disentangled, there should exist a one-to-one, bijective mapping $h(z) \rightarrow v$ between latent variables z_j and generative parameters v_k .

For graphs of which we know the ground truth generative parameters v_k , we use the metric *Mutual Information Gap (MIG)* [27] to quantify the degree of correlation between z_j and v_k . MIG measures both, the extent to which latent variables z_j share mutual information with generative parameters v_k , and the mutual independence of the latent variables from each other. The metric ranges between 0 and 1, where 1 represents a perfectly disentangled scenario in which there exists a deterministic, invertible one-to-one mapping between z_j and v_k . MIG is computed by first identifying the two latent variables z_j of highest *mutual information (MI)* with each generative parameter v_k . The MIG score is then defined as the difference (gap) between the highest and second highest MI, averaged over the generative factors v_k .

3 Evaluation

3.1 Modelling Graph Topology with Latent Variables

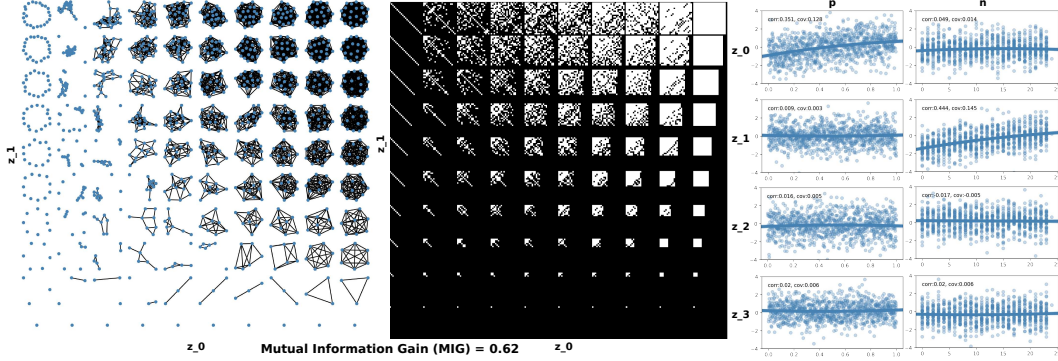


Figure 2: **Disentangled latent representation of ER graphs** The latent space appears axis-aligned with z_0 and z_1 orthogonally representing p and n . Changing one latent variable z_0 or z_1 corresponds to a change in one generative parameter p or n respectively, while being relatively invariant to changes in other parameters. z_2 and z_3 are not utilized by the model.

First, we evaluate our approach on synthetically generated graphs, concretely, ER graphs [3]. The ER generation procedure takes two parameters: the number of nodes n and a uniform linking probability p . Ideally, our model should be able to single out these independent generative parameters by utilizing only two latent variables that describe a one-to-one mapping. To test this hypothesis, we generate 10,000 ER graphs, n varying between 1 and 24 and p between 0 and 1. We use these to train our model with a latent space of size $J = 4$ and $\beta = 5.0$.

To inspect the latent space of the trained model, we sample from z_j in fixed-size steps and decode the sample through $p_\Theta(x | z)$. Figure 2 shows graphs (on the left) and adjacency matrices (in the center) sampled from the latents z_0 and z_1 while keeping other latents z_3, z_4 fixed. The adjacency matrices allow reading off topological properties of the graphs such as degree distribution and assortativity since nodes are sorted according to the extended *BOSAM* [33] algorithm. Instead of decoding samples, we may also encode graph instances x with known ground truth generative parameters v_k and observe the latents z_j . We generate a new set of 1,000 ER graphs with varying v_k and feed these graphs to the trained model. In figure 2 (on the right), each row displays samples from one latent variable z_j and the columns represent generative parameters p and n . We find that a change in p or n results in a change in z_0 or z_1 respectively, while being invariant to changes in other variables. This is manifested in a MIG of 0.62, denoting moderate to strong disentanglement. z_2 and z_3 do not show correlations with either p or n , emphasizing their "non-utilization". This shows that the latent variables of our model correctly discover the dimensionality 2 of the underlying generative procedure of ER graphs.

We repeat the experiment on a uni-, bi- and tri-parametric random graph model and two real-world graphs presented in the appendix. The selected graphs are complete binary tree graphs, *BA graphs* [34], *Small-World graphs* [35] as well as the *CORA* [36] and *Wikipedia Hyperlink* [37] graph.

3.2 Measuring Graph Topology-Node Attribute Dependence

In addition to pure graph topology τ , we consider node-level attributes Ω and measure the degree to which τ and Ω are mutually dependent. For example in a co-authorship graph where nodes represent authors and undirected links represent joint papers between authors, each node may hold additional information about the author’s overall citation count. We denote this additional information as node attributes Ω . Intuitively, more collaborations and therefore a higher node degree encourage a higher citation count, though there may be numerous other hidden correlations between graph topology and node attributes. Most existing topology-based approaches cannot make a general statement to what extent graph topology and node attributes are correlated without hand-picking particular topological properties such as the node degree [38].

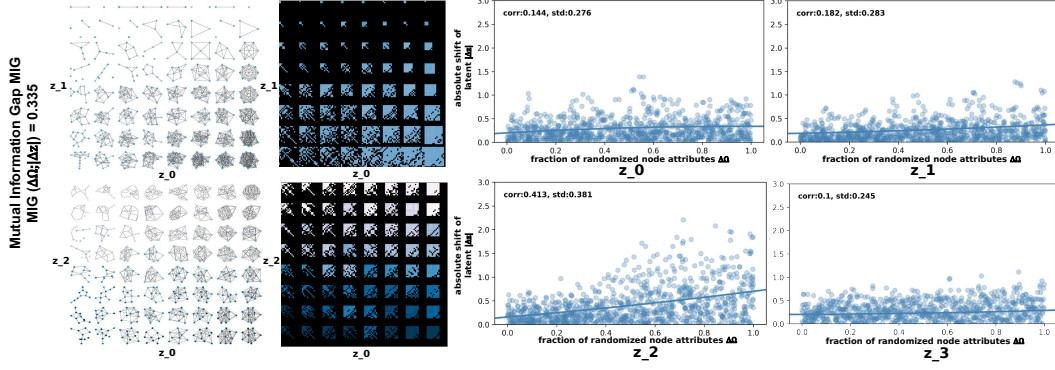


Figure 3: **Latent representation of ER graphs with uniform node attributes** Node attribute values are indicated by the shade of blue. Traversing z_0 and z_1 while keeping other latent variables fix reveals a change in the topology τ , as p and n vary. z_0 and z_1 are invariant to node attributes Ω . Since z_2 is most volatile to $\Delta\Omega$, it presumably models Ω .

We claim that the dependence between topological structure τ and attributes Ω is encoded in the latent variables. If τ and Ω are generated by independent generative procedures, they may be described by two disentangled sets of latent variables [29]. Proposing a node attribute randomization approach, we work with two data sets, the original graphs X and their attribute-randomized versions $X^{\Delta\Omega}$. Since random graph generators such as ER graphs [3] do not cover node attributes, we first have to generate synthetic node attributes. Independent from n and p , and hence from the topology τ , all nodes of an ER graph are uniformly at random assigned the same node attribute which is a value between 0 and 1. We train the modified β -VAE on this graph data set X . After training, we randomize the node attributes, ending up with the randomized graph data $X^{\Delta\Omega}$. We vary the randomization degree $\Delta\Omega$ between 0 and 1, which denotes the fraction of randomized nodes. Finally, we present X and $X^{\Delta\Omega}$ to the trained model in order to observe how the randomization affects the latent variables z .

In the case of τ - Ω independence, randomizing node attributes causes a shift in only those latent variables modelling Ω . To indirectly quantify the dependence between τ and Ω , we measure the correlation between $\Delta\Omega$ and $|\Delta z|$. $|\Delta z|$ describes the absolute change of z_j due to $\Delta\Omega$. If only one latent variable changes while others are invariant, τ and Ω are generated from a fixed number of independent factors of variation [27]. Disentanglement between latent variables serves as a proxy for the dependence of generative parameters v_k . Figure 3 (left and center) displays manifolds of samples from latent space. Traversing z_0 and z_1 while fixing other latent variables reveals a change in τ , as p and n change, but invariance to Ω . z_2 is modelling Ω , which is supported by figure 3 (right) showing absolute shifts Δz_j in the latents depending on the fraction of randomized nodes $\Delta\Omega$.

Treating the randomization degree $\Delta\Omega$ as a generative parameter, we calculate the mutual information (MI) between $\Delta\Omega$ and the absolute change in every latent z_j . $MIG(\Delta\Omega; |\Delta z|)$ then computes the gap between the first and second highest MI, normalized by the entropy $H(\Delta\Omega)$. In the equation below, $j^{max} = \arg\max_j MI(\Delta\Omega; |\Delta z_j|)$ denotes the index of latent $z_{j^{max}}$ with highest MI regarding $\Delta\Omega$.

$$MIG(\Delta\Omega; |\Delta z_j|) = \frac{1}{H(\Delta\Omega)} \left(MI(\Delta\Omega; |\Delta z_{j^{max}}|) - \max_{j \neq j^{max}} MI(\Delta\Omega; |\Delta z_j|) \right)$$

The latent variable reacting most strongly to $\Delta\Omega$ is z_2 . $MIG(\Delta\Omega; |\Delta z|)$ corresponding to figure 3 is 0.335, indicating moderate disentanglement of Ω and τ as $z_{j/2}$ are mostly invariant to $\Delta\Omega$. We repeat the experiment on the *Microsoft Academic Graph (MAG)* [39] and *Amazon Co-Purchasing Graph* [40], presented in the appendix. In particular for the Microsoft Academic Graph, the analysis reveals a strong impact of the collaboration patterns (graph topology) on the citation count (node attributes).

Conclusion This work demonstrates the potential of latent variable disentanglement in graph deep learning for unsupervised discovery of generative parameters of random and real-world graphs. Experiments have largely confirmed our hypotheses, but also revealed shortcomings. Future work should advance node order-independent graph decoders and target interpretability by exploiting generative models that do not sacrifice reconstruction fidelity for disentanglement.

References

- [1] Albert-László Barabási and Márton Pósfai. *Network Science*. Cambridge University Press, Cambridge, 2016.
- [2] Deepayan Chakrabarti and Christos Faloutsos. Graph mining: Laws, generators, and algorithms. *ACM Computing Survey*, 38, 2006.
- [3] Paul Erdős and Alfred Rényi. On random graphs. *Publicationes Mathematicae Debrecen*, 6, 1959.
- [4] Jure Leskovec, Deepayan Chakrabarti, Jon Kleinberg, Christos Faloutsos, and Zoubin Ghahramani. Kronecker graphs: An approach to modeling networks. *Journal of Machine Learning Research*, 11, 2010.
- [5] S.N. Dorogovtsev and J.F.F. Mendes. Evolution of networks. *Advances in Physics*, 51(4), 2002.
- [6] Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. Graph evolution: Densification and shrinking diameters. *ACM Transactions on Knowledge Discovery from Data*, 8:56–68, 2007.
- [7] Alexei Vazquez, Andrea Flammini, Andrea Maritan, and Alessandro Vespignani. A global protein function prediction in protein-protein interaction networks. *Nature Biotech*, pages 697–700, 2003.
- [8] Shi Zhou and Raúl J. Mondragón. Accurately modeling the internet topology. *Physical Review E*, 128:578–586, 2004.
- [9] Filippo Menczer. Evolution of document networks. *Proceedings of the National Academy of Sciences*, 101:5261–5265, 2004.
- [10] Michael M. Bronstein, Joan Bruna Estrach, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: Going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.
- [11] Ivan Brugere, Brian Gallagher, and Tanya Y. Berger-Wolf. Network structure inference, a survey: Motivations, methods, and applications. *ACM Computing Survey*, 51(2), 2018.
- [12] Daixin Wang, Peng Cui, and Wenwu Zhu. Structural deep network embedding. In *Knowledge Discovery and Data Mining (KDD)*.
- [13] Daniel D. Johnson. Learning graphical state transitions. In *International Conference on Learning Representations (ICLR)*, volume 34, pages 370–378, 2017.
- [14] Jiaxuan You, Rex Ying, Xiang Ren, William L. Hamilton, and Jure Leskovec. Graphrnn: A deep generative model for graphs. *CoRR*, abs/1802.08773, 2018.
- [15] Yujia Li, Oriol Vinyals, Chris Dyer, Razvan Pascanu, and Peter Battaglia. Learning deep generative models of graphs. In *International Conference on Machine Learning (ICML)*, 2018.
- [16] Qi Liu, Miltiadis Allamanis, Marc Brockschmidt, and Alexander Gaunt. Constrained graph variational autoencoders for molecule design. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [17] Yujia Li, Richard Zemel, Marc Brockschmidt, and Daniel Tarlow. Gated graph sequence neural networks. In *International Conference on Learning Representations (ICLR)*, 2016.
- [18] Martin Simonovsky and Nikos Komodakis. Graphvae: Towards generation of small graphs using variational autoencoders. In *Artificial Neural Networks and Machine Learning (ICANN)*, pages 412–422. Springer International Publishing, 2018.
- [19] Jianxin Ma, Peng Cui, Kun Kuang, Xin Wang, and Wenwu Zhu. Disentangled graph convolutional networks. In *International Conference on Machine Learning (ICML)*.
- [20] Aleksandar Bojchevski, Oleksandr Shchur, Daniel Zügner, and Stephan Günnemann. NetGAN: Generating graphs via random walks. In *International Conference on Machine Learning (ICML)*, 2018.

- [21] Aditya Grover and Jure Leskovec. Node2vec: Scalable feature learning for networks. In *Knowledge Discovery and Data Mining (KDD)*, 2016.
- [22] Shaosheng Cao, Wei Lu, and Qiongkai Xu. Deep neural networks for learning graph representations. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [23] Bryan Perozzi, Vivek Kulkarni, and Steven Skiena. Walklets: Multiscale graph embeddings for interpretable network classification. *CoRR*, abs/1605.02115, 2016.
- [24] Emmanuel Noutahi, Dominique Beani, Julien Horwood, and Prudencio Tossou. Towards interpretable sparse graph representation learning with laplacian pooling. *CoRR*, abs/1905.11577, 2019.
- [25] Saket Navlakha. Learning the structural vocabulary of a network. *Neural Computing*, 29(2): 287–312, 2017. ISSN 0899-7667.
- [26] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.
- [27] Tian Qi Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2610–2620, 2018.
- [28] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2172–2180. Curran Associates, Inc., 2016.
- [29] Hyunjik Kim and Andriy Mnih. Disentangling by factorising. *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 80:2649–2658, 2018.
- [30] Jan Stühmer, Richard E. Turner, and Sebastian Nowozin. Independent subspace analysis for unsupervised learning of disentangled representations. *arXiv*, abs/1909.05063, 2019.
- [31] Irina Higgins, Loïc Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations (ICLR)*, volume 1, pages 370–378, 2017.
- [32] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *International Conference on Learning Representations (ICLR)*, 34:34–42, 2017.
- [33] Yuchun Guo, Changjia Chen, and Shi Zhou. Bosam: A topology visualisation tool for large-scale complex networks. *arXiv preprint cs/0602034*, 2006.
- [34] Albert-Laszlo Barabasi and Reka Albert. Emergence of scaling in random networks. *Science*, 286, 1999.
- [35] Duncan J. Watts and Steven H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684):440–442, 1998.
- [36] Andrew McCallum. Cora dataset, 2017.
- [37] Ryan A. Rossi and Nesreen K. Ahmed. The network data repository with interactive graph analytics and visualization. In *AAAI*, 2015. URL <http://networkrepository.com>.
- [38] M. J. Zaki. Scalable algorithms for association mining. *IEEE Transactions on Knowledge and Data Engineering*, 12(3):372–390, 2000.
- [39] Arnab Sinha, Zhihong Shen, Yang Song, Hao Ma, Darrin Eide, Bo-June (Paul) Hsu, and Kuansan Wang. An overview of microsoft academic service (mas) and applications. In *Proceedings of the 24th International Conference on World Wide Web (WWW)*, pages 243–246, New York, NY, USA, 2015. ACM.
- [40] Jure Leskovec, Lada A. Adamic, and Bernardo A. Huberman. The dynamics of viral marketing. *ACM Transactions on the Web*, 5:340–350, 2007.

Code

To allow experimenting with the code, we provide an interactive notebook at https://colab.research.google.com/drive/1M--YX4d0St3imDPdecPbjVX-T6Ae0_OG.

Appendix

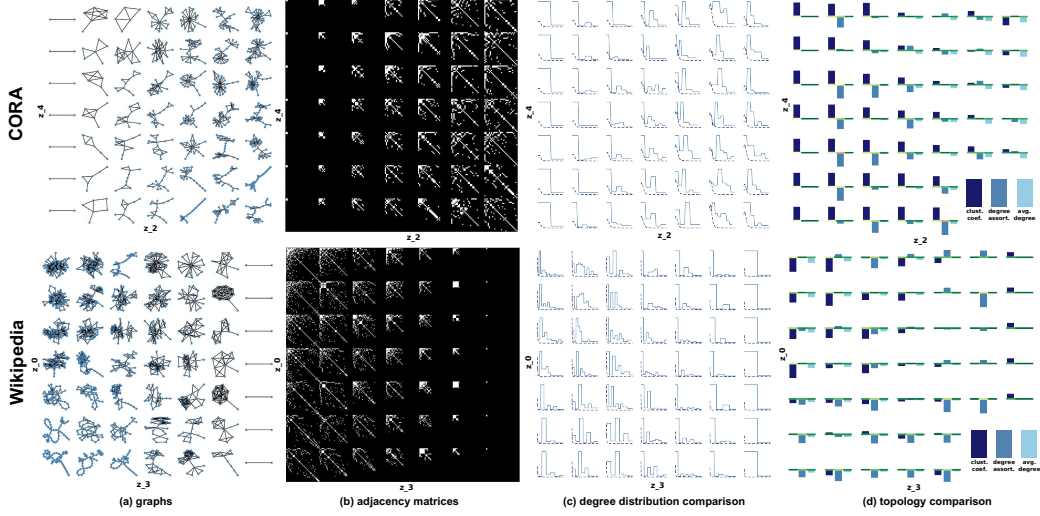


Figure 4: **Latent representations of real-world graphs** Latent space of β -VAE model trained on 10,000 sub-graphs from *CORA* [36] and *Wikipedia* [37] sampled using *Biased Second-Order Random Walks* [21]. Plot (a) and (b) show manifolds of decoded x instances, presented as graphs and adjacency matrices respectively. Plot (c) compares the normalized degree distribution of x with the distribution of the entire, original graph. Similarly, plot (d) shows the difference in clustering coefficient, degree assortativity and average degree.

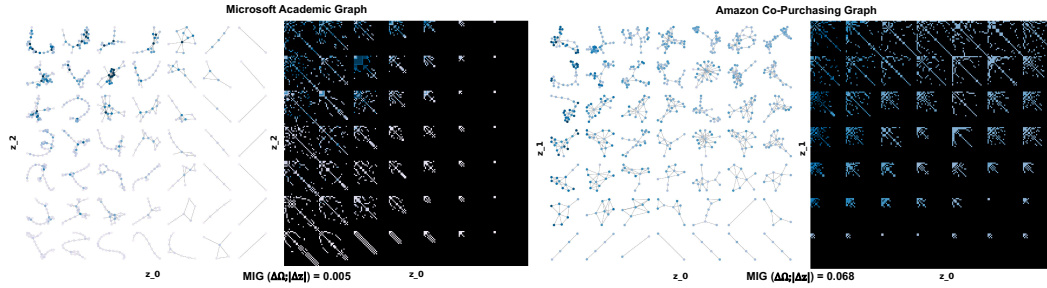


Figure 5: **Latent representations of real-world graphs with node attributes** Manifold of graph instances obtained from traversing latent variables z_j and decoding samples according to $p_{\Theta}(x|z)$. In the *Microsoft Academic Graph*, topology τ and node attributes Ω can hardly be disentangled ($MIG(\Delta\Omega; |\Delta z|) = 0.005$). A reason lies in a strong correlation (0.4662) between the number of collaborations (τ) and citations (Ω).

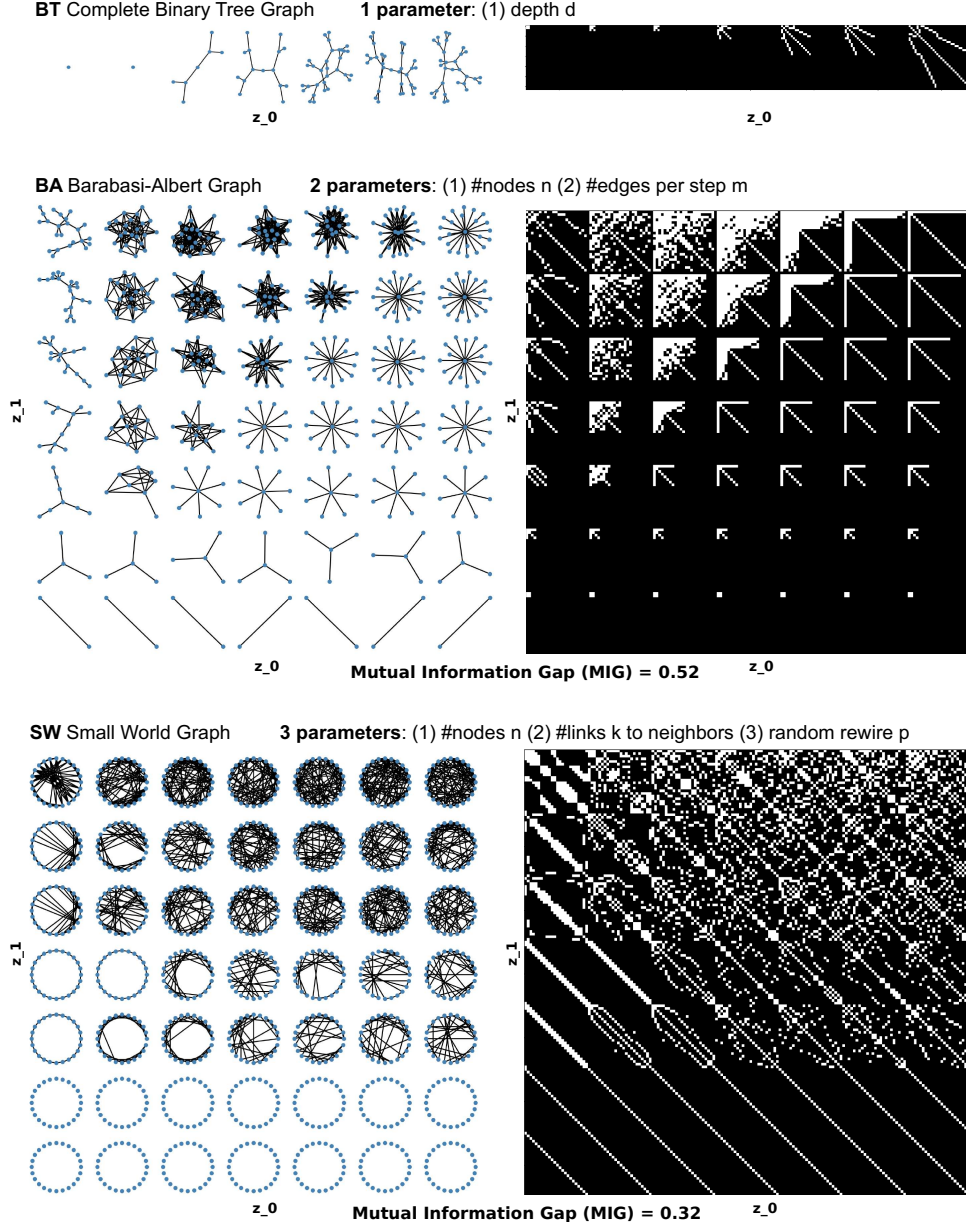


Figure 6: Disentangled latent representation of uni-, bi- and tri-parametric random graph generator models Latent representation of uni-parametric complete binary tree graph, bi-parametric *Barabasi-Albert (BA) graphs* [34] and tri-parametric *Small-World graphs (SW)* [35]. For visualizing the tri-parametric *SW graphs*, we pick a fixed value for z_0 throughout all samples from the latent space. Since z_0 models the number of nodes n , all generated graphs in the manifolds are of fixed size. In compliance with intuition, the higher the degree of freedom in terms of generative parameters, the more difficult their successful disentanglement, manifested in a lower *MIG* value for the tri-parametric *SW graph*. If a uni-parametric model is described by a single latent variable, *MIG* is not informative.