A. Graph statistics

Table 4: Graph statistics used to measure graph properties in this work.

Metric name	Computation	Description
Maximum degree	$\max_{v \in V} d(v)$	Maximum degree of all nodes in a graph.
Assortativity	$\rho = \frac{\mathrm{cov}(X,Y)}{\sigma_X \sigma_Y}$	Pearson correlation of degrees of connected nodes, where the (x_i, y_i) pairs are the degrees of connected nodes.
Triangle count	$\frac{\lvert\{\{u,v,w\} \mid \{(u,v),(v,w),(u,w)\} \subseteq E\}\rvert}{6}$	Number of triangles in the graph, where $u \sim v$ denotes that u and v are connected.
Power law exponent	$1 + n \left(\sum_{u \in V} \log \frac{d(u)}{d_{\min}} \right)^{-1}$	Exponent of the power law distribution, where d_{min} denotes the minimum degree in a network.
Inter-community density	$\frac{1}{K} \sum_{j=1}^{K} \sum_{\substack{k=1 \\ k \neq j}}^{K} \frac{1}{\binom{\lvert C_k \rvert}{\lvert C_j \rvert}} \sum_{u \in C_j} \sum_{v \in C_k} A_{uv}$	Fraction of possible inter-community edges present in graph.
Intra-community density	$\frac{1}{K} \sum_{j=1}^{K} \frac{1}{\binom{\lvert C_j \rvert}{2}} \sum_{u,v \in C_j} A_{uv}$	Fraction of possible intra-community edges present in graph.
Wedge count	$\sum_{v \in V} \binom{d(v)}{2}$	Number of wedges (2-stars), i.e. two-hop paths in an undirected graph.
Rel. edge distr. entropy	$\frac{1}{\ln \lvert V \rvert} \sum_{v \in V} -\frac{d(v)}{\lvert E \rvert} \ln \frac{d(v)}{\lvert E \rvert}$	Entropy of degree distribution, 1 means uniform, 0 means a single node is connected to all others.
LCC	$N_{max} = \max_{f \subseteq F} \lvert f \rvert$	Size of largest connected component, where F are all connected components of the graph.
Claw count	$\sum_{v \in V} \binom{d(v)}{3}$	Number of claws (3-stars)
Gini coefficient	$\frac{2 \sum_{i=1}^{\lvert V \rvert} i \hat{d}_i}{\lvert V \rvert \sum_{i=1}^{\lvert V \rvert} \hat{d}_i} - \frac{\lvert V \rvert + 1}{\lvert V \rvert}$	Common measure for inequality in a distribution, where $\hat{\boldsymbol{d}}$ is the sorted list of degrees in the graph.
Community distribution	$c_i = \frac{\sum_{v \in C_i} d(v)}{\sum_{v \in V} d(v)}$	Share of in- and outgoing edges of community C_i, normalized by the number of edges in the graph.

B. Baselines

- **Configuration model.** In addition to randomly rewiring *all* edges in the input graph, we also generate random graphs with similar overlap as graphs generated by NetGAN using the configuration model. For this, we randomly select a share of edges (e.g. 39%) and keep them fixed, and shuffle the remaining edges. This leads to a graph with the specified edge overlap; in Table 2 we show that with the same edge overlap, NetGAN's generated graphs in general match the input graph better w.r.t the statistics we measure.

- **Exponential random graph model.** We use the R implementation of ERGM from the `ergm` package (Handcock et al., 2017). We used the following parameter settings: `edge count, density, degree correlation, deg1.5, and gwesp`. Here, `deg1.5` is the sum of all degrees to the power of 1.5, and `gwesp` refers to the geometrically weighted edgewise shared partner distribution.

- **Degree-corrected stochastic blockmodel.** We use the Python implementation from the `graph-tool` package (Peixoto) using the recommended hyperparameter settings.

- **Variational graph autoencoder.** We use the implementation provided by the authors (`https://github.com/tkipf/gae`). We construct the graph from the predicted edge probabilities using the same protocol as in Sec. 3.3 of our paper. To ensure a fair comparison we perform early stopping, i.e. select the weights that achieve the best validation set performance.

C. Properties of generated graphs

Table 5: Comparison of graph statistics between the CITESEER/CORA-ML graph and graphs generated by GraphGAN and DC-SBM, averaged after 5 trials.

Graph		Max. degree Avg.	Std.	Assortativity Avg.	Std.	Triangle count Avg.	Std.	Power law exponent Avg.	Std.	Avg. Inter-community density Avg.	Std.	Avg. Intra-community density Avg.	Std.	Clustering coefficient Avg.	Std.
CITESEER		77		-0.022		451		2.239		4.9e-4		9.3e-4		1.08e-2	
Conf. model		*	*	-0.017	± 0.006	20	± 6.50	*	*	1.1e-3	± 1e-5	2.3e-4	± 2e-5	5.80e-4	± 1.29e-4
Conf. model	(42% EO)	*	*	-0.020	± 0.009	54	± 8.8	*	*	8.4e-4	± 1e-5	5.1e-4	± 1e-5	1.33e-3	± 6.15e-5
Conf. model	(76% EO)	*	*	-0.024	± 0.006	207	± 11.8	*	*	6.3e-4	± 1e-5	7.6e-4	± 1e-5	5.00e-3	± 2.57e-4
DC-SBM	(6.6% EO)	53	± 5.6	0.022	± 0.018	257	± 30.9	2.066	± 0.014	7.6e-4	± 2e-5	5.3e-4	± 3e-5	1.00e-2	± 2.63e-3
ERGM	(27% EO)	66	± 1	0.052	± 0.005	415.6	± 8	2.0	± 0.01	9.3e-4	± 2e-5	4.8e-4	± 6e-6	1.49e-2	± 5.68e-4
BTER	(2% EO)	70	± 7.2	0.065	± 0.014	449	± 33	2.049	± 0.01	1.1e-3	± 2e-5	2.8e-4	± 6e-6	1.22e-2	± 2.31e-3
VGAe	(0.2% EO)	9.2	± 0.7	-0.057	± 0.016	2	± 1	2.039	± 0.00	1.2e-3	± 1e-5	2.5e-4	± 2e-5	1.35e-3	± 9.96e-4
NetGAN VAL	(42% EO)	54	± 4.2	-0.082	± 0.009	316	± 11.2	2.154	± 0.003	6.5e-4	± 2e-5	8.0e-4	± 2e-5	1.99e-2	± 3.48e-3
NetGAN EO	(76% EO)	63	± 4.3	-0.054	± 0.006	227	± 13.3	2.204	± 0.003	5.9e-4	± 2e-5	8.6e-4	± 1e-5	7.71e-3	± 2.43e-4
CORA-ML		240		-0.075		2,814		1.86		4.3e-4		1.7e-3		2.73e-3	
Conf. model		*	*	-0.030	± 0.003	322	± 31	*	*	1.6e-3	± 1e-5	2.8e-4	± 1e-5	3.00e-4	± 2.88e-5
Conf. model	(39% EO)	*	*	-0.050	± 0.005	420	± 14	*	*	1.1e-3	± 1e-5	8.0e-4	± 1e-5	4.10e-4	± 1.40e-5
Conf. model	(52% EO)	*	*	-0.051	± 0.002	626	± 19	*	*	9.8e-4	± 1e-5	9.9e-4	± 2e-5	6.10e-4	± 1.85e-5
DC-SBM	(11% EO)	165	± 9.0	-0.052	± 0.004	1,403	± 67	1.814	± 0.008	6.7e-4	± 2e-5	1.2e-3	± 4e-5	3.30e-3	± 2.71e-4
ERGM	(56% EO)	243	± 1.94	-0.077	± 0.000	2,293	± 23	1.786	± 0.003	6.9e-4	± 2e-5	1.2e-3	± 1e-5	2.17e-3	± 5.44e-5
BTER	(2% EO)	199	± 13	0.033	± 0.008	3060	± 114	1.787	± 0.004	1.1e-3	± 1e-5	7.5e-4	± 1e-5	4.62e-3	± 5.92e-4
VGAe	(0.3% EO)	13.1	± 1	-0.010	± 0.014	14	± 3	1.674	± 0.001	1.4e-3	± 2e-5	3.2e-4	± 1e-5	1.17e-3	± 2.02e-4
NetGAN VAL	(39% EO)	199	± 6.7	-0.060	± 0.004	1,410	± 30	1.773	± 0.002	6.5e-4	± 1e-5	1.3e-3	± 2e-5	2.33e-3	± 1.75e-4
NetGAN EO	(52% EO)	233	± 3.6	-0.066	± 0.003	1,588	± 59	1.793	± 0.003	6.0e-4	± 1e-5	1.4e-3	± 1e-5	2.44e-3	± 1.91e-4

Graph		Wedge count Avg.	Std.	Rel. edge distr. entr. Avg.	Std.	Largest conn. comp Avg.	Std.	Claw count Avg.	Std.	Gini coeff. Avg.	Std.	Edge overlap Avg.	Std.	Characteristic path length Avg.	Std.
CITESEER		16,824		0.959		2,110		125,701		0.404		1		10.33	
Conf. model		*	*	0.955	± 0.001	2,011	± 6.8	*	*	*	*	0.008	± 0.001	5.95	± 0.03
Conf. model	(42% EO)	*	*	0.956	± 0.001	2,045	± 12.5	*	*	*	*	0.42	± 0.002	6.14	± 0.03
Conf. model	(76% EO)	*	*	0.957	± 0.001	2,065	± 10.2	*	*	*	*	0.76	± 0.0	6.85	± 0.04
DC-SBM	(6.6% EO)	15,531	± 592	0.938	± 0.001	1,697	± 27	69,818	± 11,969	0.502	± 0.005	0.066	± 0.011	7.75	± 0.26
ERGM	(27% EO)	16,346	± 101	0.945	± 0.001	1,753	± 15	80,510	± 1,337	0.474	± 0.003	0.27	± 0.01	5.92	± 0.01
BTER	(2% EO)	18,193	± 661	0.940	± 0.001	1,708	± 14	113,425	± 19,737	0.491	± 0.007	0.02	± 0.002	5.66	± 0.07
VGAe	(0.2% EO)	8,141	± 47	0.986	± 0.000	2,110	± 0	6,611	± 144	0.256	± 0.003	0.002	± 0.001	7.75	± 0.04
NetGAN VAL	(42% EO)	12,998	± 84.6	0.969	± 0.000	2,079	± 12.6	57,654	± 4,226	0.354	± 0.001	0.42	± 0.006	8.28	± 0.11
NetGAN EO	(76% EO)	15,202	± 378	0.963	± 0.000	2,053	± 23	94,149	± 11,926	0.385	± 0.002	0.76	± 0.01	7.68	± 0.13
CORA-ML		101,872		0.941		2,810		3.1e6		0.482		1		5.61	
Conf. model		*	*	0.928	± 0.002	2,785	± 4.9	*	*	*	*	0.013	± 0.001	4.38	± 0.01
Conf. model	(39% EO)	*	*	0.931	± 0.002	2,793	± 2.0	*	*	*	*	0.39	± 0.0	4.41	± 0.02
Conf. model	(52% EO)	*	*	0.933	± 0.001	2,793	± 6.0	*	*	*	*	0.52	± 0.0	4.46	± 0.02
DC-SBM	(11% EO)	73,921	± 3,436	0.934	± 0.001	2,474	± 18.9	1.2e6	± 170,045	0.523	± 0.003	0.11	± 0.003	5.12	± 0.04
ERGM	(56% EO)	98,615	± 385	0.932	± 0.001	2,489	± 11	3,1e6	± 57,092	0.517	± 0.002	0.56	± 0.014	4.59	± 0.02
BTER	(2% EO)	91,813	± 3,546	0.935	± 0.000	2,439	± 19	2.0e6	± 280,945	0.515	± 0.003	0.02	± 0.001	4.59	± 0.03
VGAe	(0.3% EO)	31,290	± 178	0.990	± 0.000	2,810	± 0	46,586	± 937	0.223	± 0.003	0.003	± 0.001	5.28	± 0.01
NetGAN VAL	(39% EO)	75,724	± 1,401	0.959	± 0.000	2,809	± 1.6	1.8e6	± 141,795	0.398	± 0.002	0.39	± 0.004	5.17	± 0.04
NetGAN EO	(52% EO)	86,763	± 1,096	0.954	± 0.001	2,807	± 1.6	2.6e6	± 103,667	0.42	± 0.003	0.52	± 0.001	5.20	± 0.02

D. Graph statistics during the training process

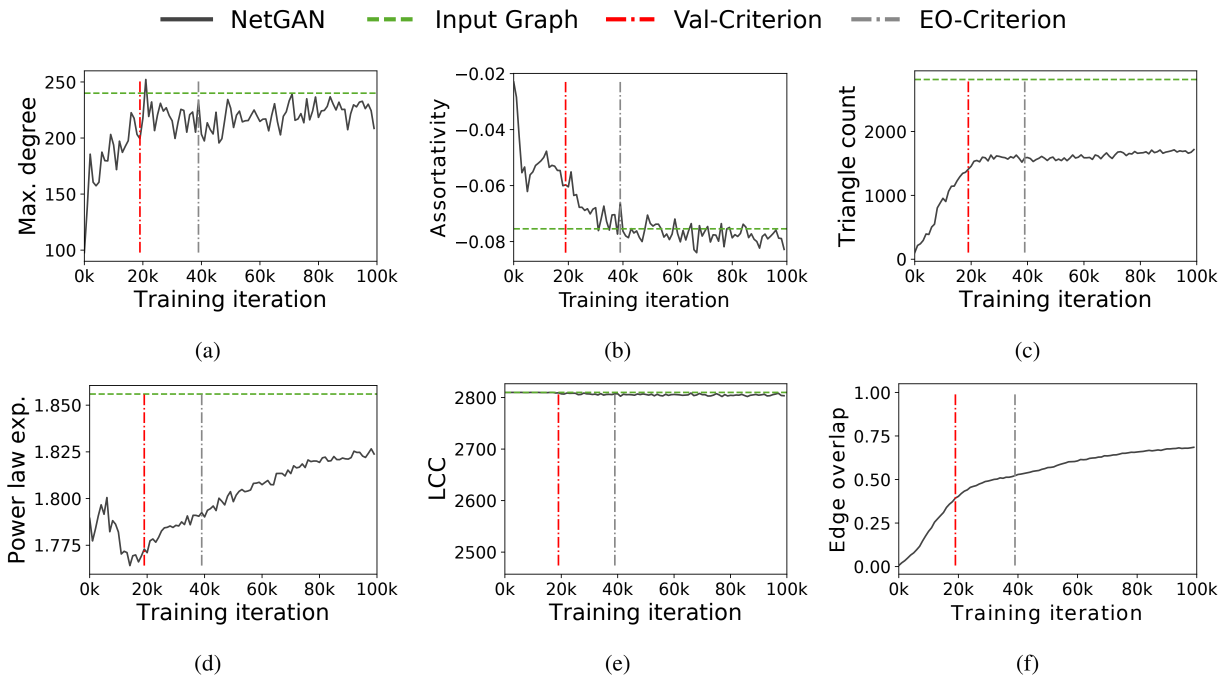

Figure 7: Evolution of graph statistics during training on CORA-ML

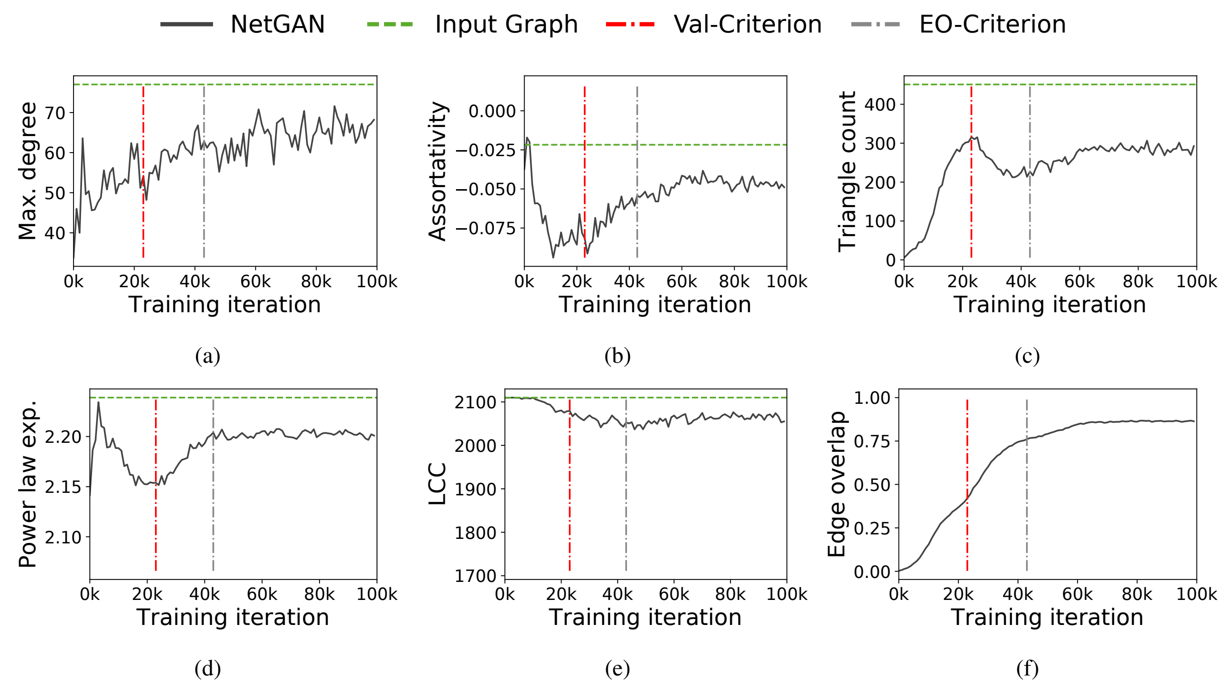

Figure 8: Evolution of graph statistics during training on CITESEER

E. Latent space interpolation heatmaps

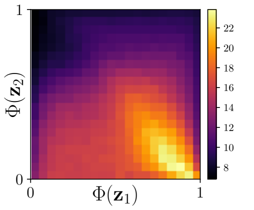
(a) Avg. degree of start node

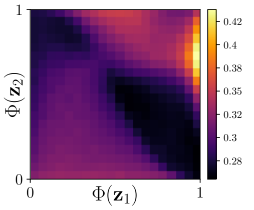
(b) Avg. share of nodes in the same comm. as the starting node

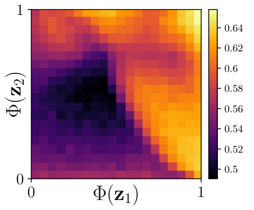
(c) Gini coefficient (input graph: 0.48)

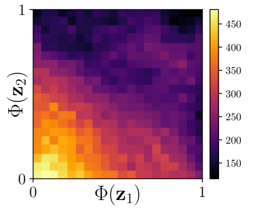
(d) Max. degree (input graph: 240)

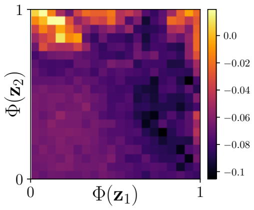
(e) Assortativity (input graph: -0.075)

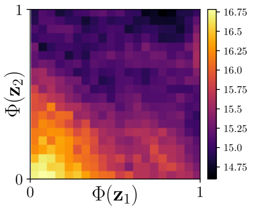
(f) Claw count (input graph: 3.1×10^6)

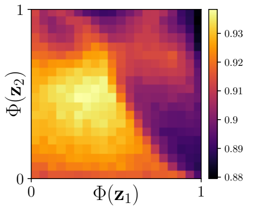
(g) Rel. edge distr. entropy (input graph: 0.94)

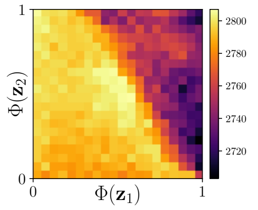
(h) Largest conn. comp. (input graph: 2,810)

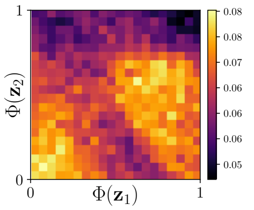
(i) Edge overlap

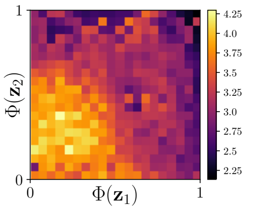
(j) Power law exponent (input graph: 1.86)

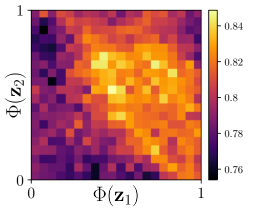
(k) Avg. precision link prediction

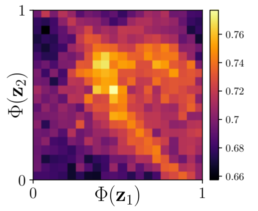
(l) ROC AUC link prediction

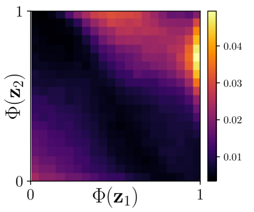
(m) Share of walks in single community

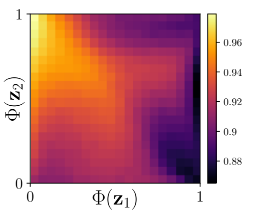
(n) tAvg. start node entropy

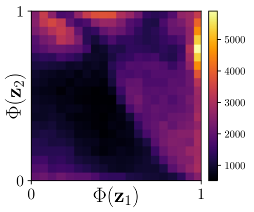
(o) Triangle count (input graph: 2,814)

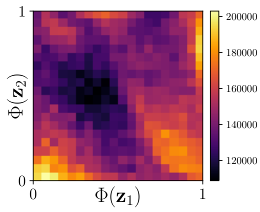
(p) Wedge count (input graph: 101,872)

Figure 9: Properties of the random walks as well as the graphs sampled from the 20×20 latent space bins, trained on CORA-ML.

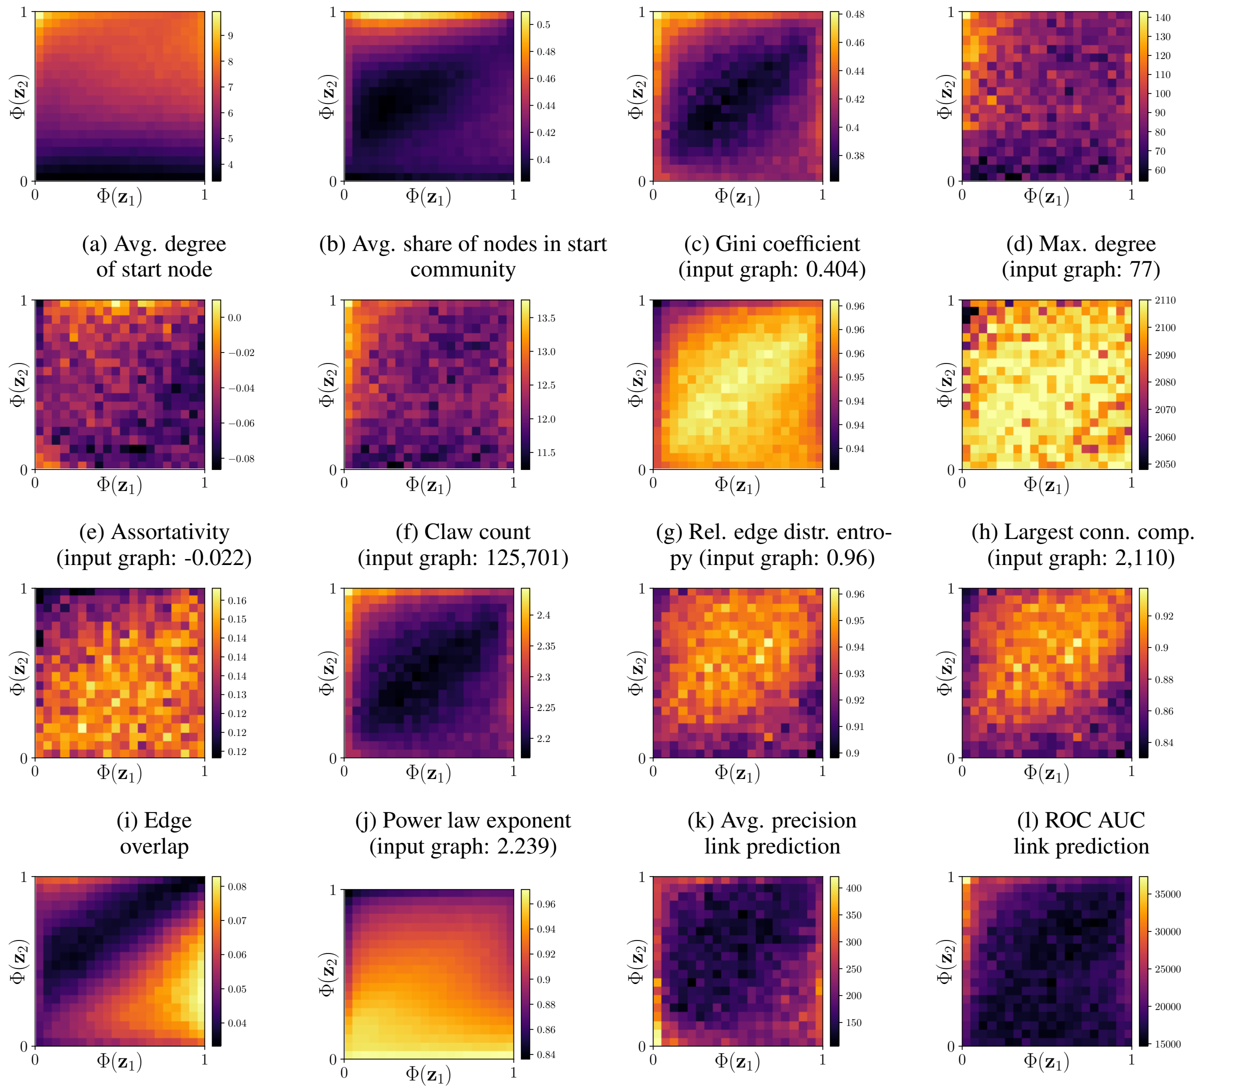

(a) Avg. degree of start node

(b) Avg. share of nodes in start community

(c) Gini coefficient (input graph: 0.404)

(d) Max. degree (input graph: 77)

(e) Assortativity (input graph: -0.022)

(f) Claw count (input graph: 125,701)

(g) Rel. edge distr. entropy (input graph: 0.96)

(h) Largest conn. comp. (input graph: 2,110)

(i) Edge overlap

(j) Power law exponent (input graph: 2.239)

(k) Avg. precision link prediction

(l) ROC AUC link prediction

(m) Share of walks in single community

(n) Avg. start node entropy

(o) Triangle count (input graph: 451)

(p) Wedge count (input graph: 16,824)

Figure 10: Properties of the random walks as well as the graphs sampled from the 20×20 latent space bins, trained on CITESEER.

F. Latent space interpolation community histrograms – CITESEER

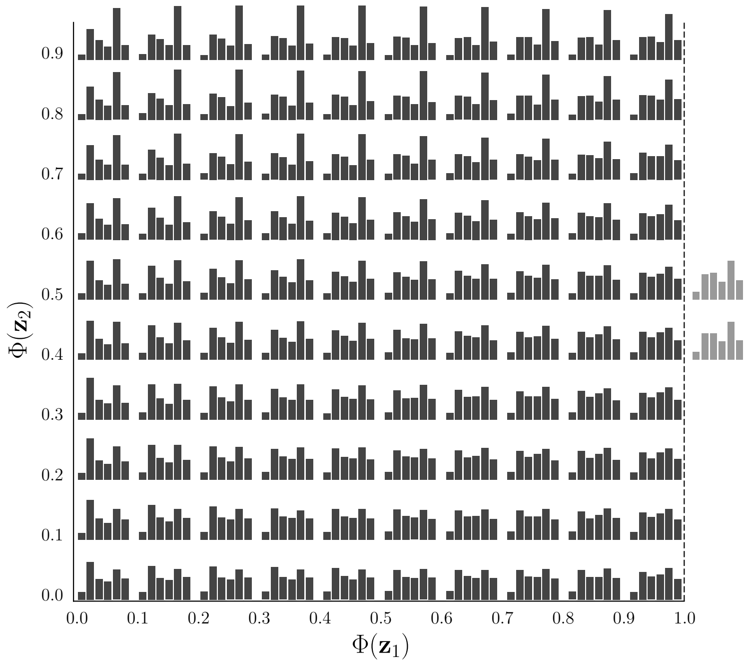

Figure 11: Community distributions of graphs generated by NetGAN on subregions of the latent space $\boldsymbol{z}$, trained on the CITESEER network.

G. Recovering ground-truth edge probabilities

To further investigate the ability of NetGAN to capture the graph structure we perform an additional experiment with the goal of analyzing how well we can recover the ground-truth edge probabilities given a graph generated from a prescribed generative model. Towards that end, first, we generate a graph from DC-SBM ($N = 300$ nodes and 3 communities), then we fit NetGAN on this graph, and finally we compare the ground truth edge probabilities to the edge scores inferred by NetGAN – specifically we compute their ranking correlation. We find a correlation of 0.998 (with EO = 0.42), which shows that NetGAN uncovered the underlying generative process, without overfitting to the input graph.

H. Hyperparameter configuration

As discussed in Sec. 4.2 NetGAN is not sensitive to the choice of most hyperparameters. For completeness, we report here sensible defaults that we used in used in our experiments. The generator and discriminator each have a single hidden layer with 40 and 30 hidden units respectively. The down-projection matrix for the generator is $\boldsymbol{W}_{down,g} \in \mathbb{R}^{N \times H_g}$ with $H_g = 64$, and for the discriminator is $\boldsymbol{W}_{down,d} \in \mathbb{R}^{N \times H_d}$ with $H_d = 32$. The latent code $\boldsymbol{z}$ is drawn from a $d = 16$ dimensional multivariate standard normal distribution. We anneal the temperature from $\tau = 1.0$ down to $\tau = 0.5$ every 500 iterations with a multiplicative decay of 0.995. We tune the parameters p and q (used to bias the generated random walks) for each dataset separately using the procedure in Grover & Leskovec (2016).

We use Adam (Kingma & Ba, 2014) to optimize all the parameters with a learning rate of $1e{-}3$ and we set the regularization strength for the L_2 penalty to $1e{-}6$. We perform five update steps for the parameters of the discriminator for each single update step of the parameters of the generator, and we set the Wasserstein gradient penalty applied to the discriminator to 10 as suggested by Gulrajani et al. (2017). For early stopping, we evaluate the score every 500 iterations, and set the patience to 5 evaluation steps. To calculate the validation score we generate 15M transitions, e.g. for a random walk of length 16 (i.e. 15 transitions per random walk) this equals 1M random walks.

For more details we refer the reader to the provided reference implementation at https://www.kdd.in.tum.de/netgan.