
NetGAN without GAN: From Random Walks to Low-Rank Approximations

Luca Rendsburg¹ Holger Heidrich¹ Ulrike von Luxburg^{1,2}

Abstract

A graph generative model takes a graph as input and is supposed to generate new graphs that “look like” the input graph. While most classical models focus on few, hand-selected graph statistics and are too simplistic to reproduce real-world graphs, NetGAN recently emerged as an attractive alternative: by training a GAN to learn the random walk distribution of the input graph, the algorithm is able to reproduce a large number of important network patterns simultaneously, without explicitly specifying any of them. In this paper, we investigate the implicit bias of NetGAN. We find that the root of its generalization properties does not lie in the GAN architecture, but in an inconspicuous low-rank approximation of the logits random walk transition matrix. Step by step we can strip NetGAN of all unnecessary parts, including the GAN, and obtain a highly simplified reformulation that achieves comparable generalization results, but is orders of magnitudes faster and easier to adapt. Being much simpler on the conceptual side, we reveal the implicit inductive bias of the algorithm — an important step towards increasing the interpretability, transparency and acceptance of machine learning systems.

1. Introduction

A graph generative model is a mechanism to achieve the following task: for a given input graph (or a set of input graphs), generate new graphs that have a similar structure as the input graph. The mechanism is supposed to slightly perturb the graph, but should not change its characteristic

structure (such as the community structure, the characteristic path lengths, etc). Being able to create perturbed copies of a graph is useful in many different scenarios, for example: comparing a small sample of brain networks for Alzheimer patients (just one, in the extreme case) to a large population of healthy subjects, making robustness statements about a climate network by running a sensitivity analysis on perturbed copies, or performing a generic bootstrap analysis.

A recent graph generative model that has received a lot of attention is NetGAN (Bojchevski et al., 2018). First, it samples a set of random walks from the input graph to train a GAN (Goodfellow et al., 2014), whose generator learns to produce node-sequences that resemble random walks over the input graph. Generated graphs are then obtained as reconstructions based on these sequences. The inherent assumption of this approach is that random walks describe graphs in a reasonably holistic way: local statistics such as motifs are observable in the individual random walks, while global statistics such as cluster structure and diameter are encoded in the distribution over random walk sequences. As opposed to other approaches, NetGAN does not make any explicit model assumptions; rather, it is supposed to implicitly learn many local and global graph statistics simultaneously by reproducing random walk statistics. However, if the goal is to generalize (perturb) the input graph, there has to be an implicit bias as to which type of generalization (perturbation) is preferred (no free lunch). The goal of our paper is to characterize this bias of NetGAN, which we will achieve by reformulating it in terms of a distance function between graphs. This formulation provides insights on the influence of design choices and model parameters, such as the length of the random walks. Scrutinizing the NetGAN architecture, we observe that many of its components can be considerably simplified. Step by step we strip all the unnecessary parts until we are left with the only crucial ingredient, a low-rank approximation of the logits random walk transition matrix. Our main contributions are:

- **Reformulation of NetGAN.** We reformulate NetGAN as a low-rank approximation with respect to the Kullback-Leibler divergence between transition matrices, which requires neither a GAN nor any sampling.
- **Huge speedup.** Our algorithm retains the generalization

¹Department of Computer Science, University of Tübingen, Germany ²Max Planck Institute for Intelligent Systems, Tübingen, Germany. Correspondence to: Luca Rendsburg <luca.rendsborg@uni-tuebingen.de>.

NetGAN without GAN




Figure 1. Pipelines for NetGAN (upper path) and our proposed method CELL (lower path). CELL is a condensed version of NetGAN that bypasses the expensive sampling steps and replaces the GAN with an optimization problem.

performance of NetGAN, but runs in seconds instead of hours. See Table 1 for a comparison of training times.

- **Transparency.** Our algorithm is conceptually much simpler than NetGAN. This opens the possibility to analyze it theoretically, and allows for application-specific adaptations.

2. Background: NetGAN

2.1. Graph and random walk notation

Like NetGAN, we consider an unweighted, undirected, and connected graph $G = (V, E)$ with nodes $V = [N] = \{1, \dots, N\}$, edges $E \subseteq V \times V$, and number of edges $e(G) = |E|$. It has adjacency matrix $A \in \{0, 1\}^{N \times N}$, degree vector $d \in \mathbb{R}^N$, degree matrix $D = \text{diag}(d) \in \mathbb{R}^{N \times N}$, and the transition matrix for unbiased random walks on G is given by $P = D^{-1}A \in \mathbb{R}^{N \times N}$. We additionally assume G to be non-bipartite so that the random walk described by P has a unique stationary distribution $\pi \in \mathbb{R}^N$. A single random walk of length T is an ordered tuple $R = (v_0, \dots, v_T) \in V^{T+1}$, and a set of n random walks is denoted by $\mathcal{R} = \{R_1, \dots, R_n\}$. The score matrix $S(\mathcal{R}) \in \mathbb{R}^{N \times N}$ counts the transitions in \mathcal{R} , that is, $S_{v,w}$ equals the total number of times random walks in \mathcal{R} transition from v to w . If clear from the context, we drop the dependency on \mathcal{R} and write S instead of $S(\mathcal{R})$. An edge-independent random graph model, sometimes also called inhomogeneous Erdős-Rényi model, is a symmetric matrix $A^\dagger \in [0, 1]^{N \times N}$ of edge probabilities. Graphs on the same vertices $[N]$ are sampled from this model by drawing $e(G)$ edges $\{v, w\}$ with probability $A_{v,w}^\dagger$ independently and without replacement. We use bold symbols, if we consider an object as a random variable (e. g. \mathbf{R} instead of R).

2.2. NetGAN

In this section, we give a high-level overview of the NetGAN algorithm; for more details, we refer the reader to [Bojchevski et al. \(2018\)](#). NetGAN is a graph generative model: given a single input graph G , it returns graphs G' on the same set of nodes by proceeding in two main steps. First, it

Table 1. Training time (in seconds) for NetGAN and our proposed method CELL on a variety of networks. NetGAN requires a GPU for training, while CELL runs on a CPU.

DATA SET (NODES/ EDGES)	NETGAN	CELL
CORA-ML (2,810/ 7,981)	7,478	21
CITeseer (2,110/ 3,668)	4,654	10
POLBLOGS (1,222/ 16,779)	55,276	15
RT-GOP (4,687/ 5,529)	14,800	23
WEB-EDU (3,031/ 6,474)	11,000	16

learns the distribution over random walks drawn from the input graph in the *learning step*. It then reconstructs the graph based on “synthetic” random walks sampled from this learned distribution in the *reconstruction step*. See Figure 1 for a schematic overview.

Learning step. Given an input graph G , NetGAN samples a large set \mathcal{R} of random walks of fixed length T with randomly chosen start nodes. These random walks form the training set for a GAN: the generator tries to produce node sequences of length T that resemble the observed random walks in \mathcal{R} , while the discriminator tries to distinguish real from generated sequences. Both generator and discriminator use the Long short-term memory architecture (LSTM) ([Hochreiter & Schmidhuber, 1997](#)), and they are trained with the Wasserstein loss ([Arjovsky et al., 2017](#)). Training finishes once an early stopping criterion is met, after which the generator is used to sample synthetic random walks.

During and after training, the generator constructs each synthetic random walk (v_0, \dots, v_T) in a step-by-step procedure. First, random noise z is used to initialize the memory state m_0 of the LSTM architecture and the start node v_0 of the sequence. A function f_θ with learnable parameters θ then repeatedly updates the two values: given the current memory state m_t and node v_t , it outputs the next memory state m_{t+1} and the distribution p_{t+1} over the next node v_{t+1} in form of logits. The next node v_{t+1} is then obtained as a sample from this distribution. In equations, this update is

described by

$$\begin{aligned} (m_{t+1}, p_{t+1}) &= f_{\theta}(m_t, v_t), \\ v_{t+1} &\sim \text{Cat}(\sigma(p_{t+1})), \end{aligned} \quad (1)$$

where Cat denotes the categorical distribution and σ the softmax function, which converts the logits into a probability distribution on $[N]$. This procedure is repeated until the sequence has the desired length T .

Reconstruction step. After training is finished, NetGAN uses the generator to generate a large set of n synthetic random walks. Their transitions are counted in a joint score matrix S , which is then converted into an edge-independent random graph model A^{\dagger} by symmetrizing and then normalizing it, that is,

$$A_{k,l}^{\dagger} = \frac{\max\{S_{k,l}, S_{l,k}\}}{\sum_{k',l'=1}^N \max\{S_{k',l'}, S_{l',k'}\}}. \quad (2)$$

To obtain the new graph G' , NetGAN samples $e(G)$ edges independently and without replacement from A^{\dagger} while preventing self-loops and isolated nodes.

3. What causes the generalization?

In this section, we identify those parts of NetGAN that we believe to be absolutely necessary to achieve the two goals of producing new graphs that (i) resemble the input graph by mimicking its graph statistics, but (ii) also generalize the input graph by sharing only a certain amount of its edges. The complicated GAN- and LSTM-based architecture used by NetGAN disguises its underlying bias and makes a direct analysis difficult. Therefore, we examine all the individual steps of NetGAN, not in terms of *how* they work, but *what* they aim to achieve.

The random walks? The intuition of NetGAN is that graphs with a similar random walk distribution also share many of their topological properties. In fact, as we observe in Section 5.3, learning the transition matrix of random walks by counting their transitions is sufficient for perfectly reconstructing the input graph. This excludes the possibility that by reducing graphs to their random walk statistics, we introduce an irreversible systematic bias.

The GAN? The role of the GAN is to learn the random walk distribution of the input graph. We prove in Section 5.3 that if the GAN perfectly learns the random walk distribution, NetGAN will simply reproduce the input graph instead of generalizing it. However, the results reported by [Bojchevski et al. \(2018\)](#) show that even if NetGAN is trained for a long time, it produces graphs that are considerably different from the input graph as measured by edge overlap. Consequently, there has to be another mechanism that prevents the GAN from memorizing the input graph.

The LSTM? As the authors of NetGAN pointed out themselves, the LSTM architecture, which is supposed to capture long-term dependencies, seems to be an odd choice for learning Markov sequences that by construction do not have any such dependencies. It is possible that this architecture choice injects noise into the learning process, which prevents memorization of the input graph. Yet, this type of noise seems to be rather uncontrolled, and we consider it unlikely that this aspect of the LSTM cannot be replaced by a simpler, more direct mechanism.

Computational trick: low-rank approximation. What is left? In our opinion, the only component that explains why NetGAN successfully generalizes graphs is a computational trick: the LSTM is not operating on the high-dimensional space \mathbb{R}^N directly. In order to reduce computational complexity, it uses learnable down- and up-projections $W_{\text{down}} \in \mathbb{R}^{N \times H}$ and $W_{\text{up}} \in \mathbb{R}^{H \times N}$ with $H \ll N$. As we derive in Section A of the supplementary, these projections force the update rule of a node and memory state pair (v_t, m_t) with v_t as one-hot vector to be of the form

$$\begin{aligned} p_{t+1} &= v_t^{\top} W(m_t), \\ v_{t+1} &\sim \text{Cat}(\sigma(p_{t+1})), \end{aligned} \quad (3)$$

where $W(m_t) \in \mathbb{R}^{N \times N}$ depends on m_t and has rank at most H . Because $W(m_t)$ is the transition matrix after applying σ , we refer to it as the *logit transition matrix*. NetGAN forces this matrix to have low rank, which leads us the following conjecture:

Conjecture: The key ingredient of NetGAN is to learn the random walk distribution by performing a low-rank approximation of the logit transition matrix.

To validate this conjecture, we derive a simplified method that applies this low-rank approximation directly and demonstrate its comparable performance in experiments.

4. Stripping NetGAN

We now gradually simplify NetGAN by stripping it of all unnecessary components in Section 4.1. Additionally, we observe in Section 4.2 that sampling random walks can be circumvented with a limit argument. This leads to our new, highly simplified method called Cross-Entropy Low-rank Logits (CELL), see Section 4.3 for a summary and Figure 1 for a schematic outline.

4.1. Low-rank approximation replaces the GAN

Motivated by the above conjecture, we now prune the update rule in Eq. (3) until we arrive at a rank-constrained optimization problem. Justified by the Markov property of unbiased random walks, we first drop the LSTM and the memory state m_t . In Section A of the supplementary, we derive

that the GAN learns the random walk distribution by choosing its transition matrix directly from the parametric family $\mathcal{P} = \{\sigma_{\text{rows}}(W) \in \mathbb{R}^{N \times N} : W \in \mathbb{R}^{N \times N}, \text{rank}(W) \leq H\}$, where σ_{rows} denotes the function that applies the softmax σ to each row of a matrix. The training set for this problem consists only of the transitions of random walks in \mathcal{R} , and the noise random variable z plays the subordinate role of choosing the first node. This parametric family formulation defeats the purpose of using a GAN at all, which is why instead we revert to the classical maximum likelihood approach (or, equivalently, the cross-entropy loss) on \mathcal{P} : using the notation $(k, l) \in \mathcal{R}$ to denote all transitions (with multiple counting) of random walks in \mathcal{R} , the resulting problem is given by

$$\begin{aligned} \min_{W \in \mathbb{R}^{N \times N}} & -\sum_{(k,l) \in \mathcal{R}} \log \sigma_{\text{rows}}(W)_{k,l}, \\ \text{s. t.} & \text{rank}(W) \leq H. \end{aligned} \quad (4)$$

In short: instead of learning the random walk distribution by training a GAN, we approximate its transition matrix directly by solving a rank-constrained optimization problem.

4.2. Bypassing random walk sampling

There is another aspect of NetGAN that is somewhat puzzling: even to learn a graph of moderate size, for example the graph CORA-ML with about 3,000 vertices and 8,000 edges, NetGAN needs to sample 7,500,000 random walks of length 15 from the input graph, which are worth 112,500,000 edges. In other words, we see every edge of the input graph about 14,000 times on average — with which any edge-frequency statistic would be very close to its expected value. The same order of magnitude applies to the sampling of random walks from the generator in the reconstruction step. With that observation, a natural question is whether we can circumvent the random walk sampling, and the answer is yes. Since the random walks are only used in form of the score matrix that contains the frequency of node transitions, and this matrix converges for a large number of random walks, we can substitute the actual score matrix with its limit value. The remainder of this section formalizes this idea in Eq. (7) and applies it to NetGAN at both sampling steps.

Convergence of the score matrix S . First, we consider a single random walk $\mathbf{R} = (v_0, \dots, v_T)$ of length T as a random variable, whose distribution depends on the distribution $q_0 \in \mathbb{R}^N$ of the first node v_0 and the transition matrix P . For $t \in \{1, \dots, T\}$, let $Q_t \in \mathbb{R}^{N \times N}$ denote the distribution of the t -th transition (v_{t-1}, v_t) in \mathbf{R} . Its marginal v_{t-1} is distributed as $q_{t-1} \in \mathbb{R}^N$ and its conditional $v_t | v_{t-1}$ is distributed as P , which yields the matrix decomposition

$$Q_t = \text{diag}(q_{t-1})P. \quad (5)$$

From this perspective, counting the transitions of a single random walk R in a score matrix $S(R) \in \mathbb{R}^{N \times N}$ can be expressed as $S(R) = \sum_{t=1}^T \hat{Q}_t(R)$, where $\hat{Q}_t(R)$ is the empirical version of Q_t based on one sample. The score matrix $S = S(R_1, \dots, R_n)$ based on n random walks R_1, \dots, R_n decomposes into $S = \sum_{j=1}^n S(R_j)$, and with the above considerations we have $S = \sum_{j=1}^n \sum_{t=1}^T \hat{Q}_t(R_j)$. By the Glivenko-Cantelli theorem for empirical distributions, we can compute the limit of S/n for $n \rightarrow \infty$ as

$$\frac{S}{n} = \sum_{t=1}^T \frac{1}{n} \sum_{j=1}^n \hat{Q}_t(R_j) \xrightarrow[n \rightarrow \infty]{\text{a.s.}} \sum_{t=1}^T Q_t. \quad (6)$$

Using Eq. (5), the normalized right-hand side is given by $\sum_{t=1}^T Q_t/T = \text{diag}(\rho_T)P$, where $\rho_T = \sum_{t=1}^T q_{t-1}/T$. In that sense, using a large number of random walks reduces to node weights ρ_T . Since the underlying graph is by assumption connected and non-bipartite, the stationary distribution π of P exists and is unique, that is, $\pi = \lim_{t \rightarrow \infty} q_t$. Hence the Cesàro mean ρ_T also converges to π as $T \rightarrow \infty$. Or, in other words: for any initial distribution q_0 , the node weights induced by sufficiently long random walks are given by π . In conjunction with Eq. (6), we obtain the limit of the normalized score matrix

$$\frac{S}{nT} \xrightarrow[n, T \rightarrow \infty]{\text{a.s.}} \text{diag}(\pi)P. \quad (7)$$

Note that we take two limits to approximate S . We take the first limit with respect to the amount of random walks n , because NetGAN samples many random walks. The second limit with respect to the length T dilutes the influence of the initial distribution ($\rho_1 = q_0$) in favor of the stationary distribution ($\lim_{T \rightarrow \infty} \rho_T = \pi$). This is appropriate because most real-world networks have small diameter, and the length $T = 15$ used in NetGAN already ensures that ρ_T is close to its limit distribution. Furthermore, the authors of NetGAN already observed that taking longer random walks increases performance. Finally, in Section 5.4 we will see that the information encoded in the start distribution of the random walk can be more directly incorporated by the node weights.

Replacing random walks from the input graph. The objective in Eq. (4) sums over all node transitions in \mathcal{R} . We count the transitions in a corresponding score matrix $S = S(\mathcal{R})$ to rewrite the objective function as

$$-\sum_{k,l=1}^N S_{k,l} \log \sigma_{\text{rows}}(W)_{k,l}. \quad (8)$$

Normalizing S does not change the minimum, and allows us to approximate it with the limit $\text{diag}(\pi)P$ in Eq. (7). Since we consider unbiased random walks according to $P = D^{-1}A$, the stationary distribution π is proportional to the degrees d , hence $\text{diag}(\pi)P \propto \text{diag}(d)D^{-1}A = A$.

This means that we observe every edge in every direction with the same frequency, see Lovász et al. (1993) for a survey on random walks on graphs. We use this new weighting A to define our final objective function

$$F(W) = - \sum_{k,l=1}^N A_{k,l} \log \sigma_{\text{rows}}(W)_{k,l} \quad (9)$$

and our final objective

$$\begin{aligned} \min_{W \in \mathbb{R}^{N \times N}} F(W), \\ \text{s. t. } \text{rank}(W) \leq H, \end{aligned} \quad (10)$$

whose solution is denoted as W^* . Note that the sum in Eq. (9) grows only as $\mathcal{O}(e(G))$, and we can enforce the rank-constraint in Eq. (10) with the factorization $W = W_{\text{down}}W_{\text{up}}$, where $W_{\text{down}} \in \mathbb{R}^{N \times H}$, $W_{\text{up}} \in \mathbb{R}^{H \times N}$, resulting in $\mathcal{O}(NH)$ trainable parameters and a non-convex optimization problem.

Replacing random walks from the generator. In principle, we could use the synthetic transition matrix $P^* = \sigma_{\text{rows}}(W^*)$ defined with the solution W^* of Eq. (10) in place of the generator: we produce synthetic random walks of length T , with transition matrix P^* , and with the same distribution over the first node as in the training set, and then count their transitions in a score matrix. But since the score matrix is needed only up to proportionality for the edge-independent model, we can use the limit in Eq. (7) instead, which replaces sampling random walks with solving the eigenvector problem $\pi^{*\top} P^* = \pi^{*\top}$. That is, we skip sampling random walks and simply set $S = \text{diag}(\pi^*)P^*$.

4.3. Our proposed algorithm:

Cross-Entropy Low-rank Logits (CELL)

In the previous section, we have shown how to (i) replace the LSTM and GAN architecture with a low-rank approximation of the logit transition matrix with respect to the cross-entropy loss, (ii) replace sampling random walks from the input graph with using its adjacency matrix directly, and (iii) replace sampling random walks from the generator with solving an eigenvector problem. The result of this analysis is our simplified algorithm Cross-Entropy Low-rank Logits (CELL), summarized in Algorithm 1. It takes the adjacency matrix A of a graph G as input and returns a symmetric matrix A^\dagger of edge probabilities, from which new graphs G' can be sampled. For solving optimization problem (10), we factorize $W = W_{\text{down}}W_{\text{up}}$ with $W_{\text{down}} \in \mathbb{R}^{N \times H}$ and $W_{\text{up}} \in \mathbb{R}^{H \times N}$ to satisfy the rank constraint, and optimize with Adam (Kingma & Ba, 2014). Training continues until a stopping criterion is met, for which we pause at regular intervals and generate new graphs to evaluate the stopping criterion. In this paper, we consider the criterion of reaching a predefined edge overlap of generated graphs and input graph, see Section 6.1.

Algorithm 1 Cross-Entropy Low-rank Logits (CELL)¹

input adjacency matrix $A \in \{0, 1\}^{N \times N}$, rank $H \ll N$

output matrix of edge probabilities $A^\dagger \in [0, 1]^{N \times N}$

- 1: Solve optimization problem (10) for W^*
- 2: Compute transition matrix: $P^* \leftarrow \sigma_{\text{rows}}(W^*)$
- 3: Solve eigenvalue problem $\pi^{*\top} P^* = \pi^{*\top}$ for π^*
- 4: Compute score matrix: $S \leftarrow \text{diag}(\pi^*)P^*$
- 5: Convert score matrix S to edge-independent model A^\dagger :
 $S^\dagger \leftarrow \max\{S, S^\top\}$; $A^\dagger \leftarrow S^\dagger / \text{sum}(S^\dagger)$

return A^\dagger

5. Conceptual analysis

Our simple reformulation of NetGAN now opens the possibility to formally analyze the inductive bias associated with its components and allows for user-specific adaptations.

5.1. Inductive bias of NetGAN

Our analysis has shown that the graphs produced by NetGAN come from the class of graphs whose logit transition matrix has a low rank. Note that this does *not* imply that the transition matrix itself has low rank. Even if W^* is trained to have low rank, the corresponding synthetic transition matrix $P^* = \sigma_{\text{rows}}(W^*)$ can have full rank, as is visualized by the eigenvalues in Figure 2. Additionally, our experiments in Section 6.2 suggest that approximating the transition matrix with a low rank matrix and the Frobenius norm as loss function does not achieve good generalization performance. However, minimizing the cross-entropy




Figure 2. Portion of absolute eigenvalues (sorted and rescaled) for CORA-ML with CELL trained to 50% edge overlap for $H = 9$.

loss for approximation instead yields generalization performance comparable to the one of NetGAN and CELL. Since the cross-entropy corresponds to the Kullback-Leibler (KL) divergence (see Section B of the supplementary), this suggests that **using the KL divergence as distance measure for approximating transition matrices is the reason for the good generalization performance. On a high level, NetGAN generalizes a graph by choosing new graphs, whose transition matrix is similar in terms of KL-divergence, from a restricted set of graphs.** Whether this restriction is realized by a low-rank assumption on the logits or on the transition matrices itself is not essential, although the former is computationally more feasible.

¹Code available at <https://github.com/hheidrich/CELL>

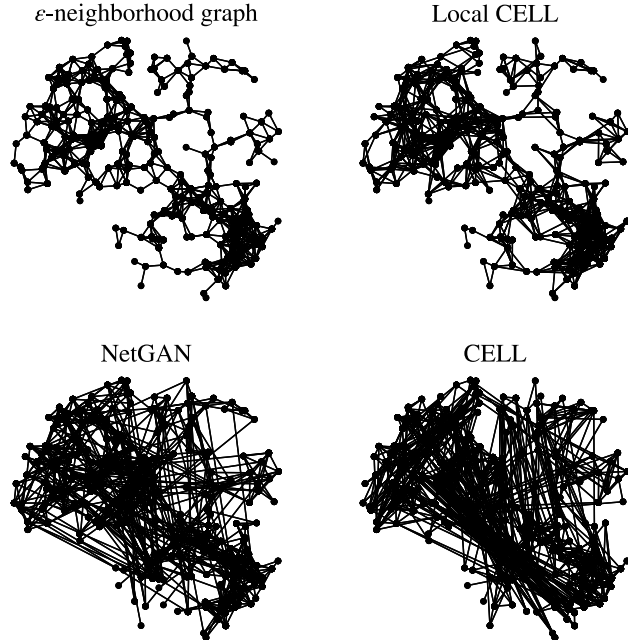


Figure 3. Comparison of an ε -neighborhood graph (top left) with graphs generated by Local CELL, a version of our method biased towards short edges, NetGAN, and our method CELL. Only Local CELL does not generate edges between distant points.

5.2. Bias of the optimization objective and resulting hard examples

We optimize the objective function in Eq. (9) to learn the random walk distribution in form of its transition matrix. By inspecting the objective, we can understand how this is achieved: the synthetic transition matrix $\sigma_{\text{rows}}(W)$ is rewarded *directly* for putting mass on edges of the input graph ($A_{k,l} = 1$). But because the total mass is limited ($\sum_l \sigma_{\text{rows}}(W)_{k,l} = 1$), it is only penalized *indirectly* for wasting mass on non-edges ($A_{k,l} = 0$). In particular, there is no distinction between different non-edges. This hints towards **poor performance for graphs with strong restrictions on the set of edges we deem realistic**, because there is no notion of “bad” edges that could prevent their generation; a possible remedy to this problem is extending the objective function with such a notion.

We illustrate this effect with the example of ε -neighborhood graphs in Figure 3. Here, we want to avoid the generation of edges between nodes with large distance in the Euclidean space, which is not taken into account by NetGAN and CELL. However, a simple adaptation of our method, denoted as “Local CELL”, can prevent long edges without loss of generalization performance. The corresponding experiment is provided in Section C.4 of the supplementary.

5.3. No bias in the reconstruction step

A natural question is whether our method might be able to generalize even without the rank constraint in the learning step. Or, phrased differently, whether the reconstruction step introduces a generalization bias. This is not the case. We derive in Section B of the supplementary that without any rank constraint, we exactly recover the input transition matrix as $P^* = P$. Because of $\text{diag}(\pi)P \propto A$, this also holds for the score matrix

$$S = \text{diag}(\pi^*)P^* \propto A. \quad (11)$$

Since A is already symmetric, the edge-independent model is given by $A^\dagger \propto A$. This model is equivalent to uniformly sampling edges from the input graph G , and sampling $e(G)$ edges from this model without replacement means sampling all of them. Hence it simply returns the input graph with zero variance. Therefore, **reconstructing the graph with an edge-independent model does not contribute to generalization**. Another interpretation of this observation is that random walks are sufficient to learn a graph in principle.

5.4. Influence of the random walk parameters

For NetGAN it is still unclear how the length T and the start distribution for the first node q_0 of the random walks influence the generated graphs. We derived in Section 4.2 that it does not exploit any complicated patterns in the random walk paths, but simply counts the transitions, which comes down to a weighting of the nodes. When translating NetGAN to our approach, we observe that **the random walk length controls how much influence the start distribution has on the node weights**: instead of taking the limit $T \rightarrow \infty$ in the derivation of Section 4.2, we could have completed the analysis with the node weights $\rho = \sum_{t=1}^T q_{t-1}/T$ to arrive at the parametrized objective function

$$F_\rho(W) = - \sum_{k,l=1}^N \frac{\rho_k}{d_k} A_{k,l} \log \sigma_{\text{rows}}(W)_{k,l}. \quad (12)$$

This allows for further interpretation and adaptation:

Random walks of length one are sufficient. The random walk parameters T and q_0 are relevant for Eq. (12) only because they determine the node weights ρ . On the other hand, all possible node weights ω can be realized by choosing $q_0 = \omega$ and $T = 1$. This implies for NetGAN that **only using random walks of length one imposes no restriction**, if the distribution of the start node is considered as a hyperparameter instead.

An example that is now readily explained is the setting in Jalilifard et al. (2019). They observe empirically that using short random walks in NetGAN reduces the performance, and propose to counteract by choosing the start

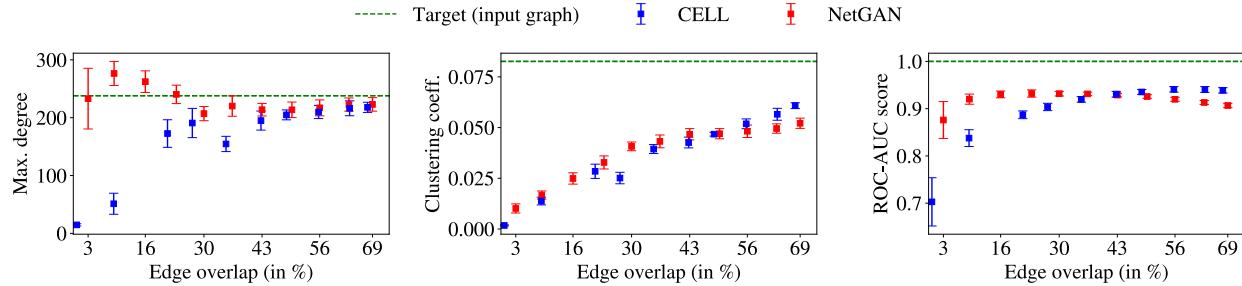


Figure 4. Mean and standard deviation for five trials on CORA-ML, plotted against edge overlap with the input graph. Aside from different initializations, NetGAN and CELL display similar behavior. Additional experiments are provided in the supplementary, Section C.6.

distribution as the density function described in Zhou et al. (2009). Within our framework, this is explained by the node weights: for short random walks, they are close to the uniform distribution (the start distribution of NetGAN), which overemphasizes nodes with low degree and results in bad performance. Choosing the start distribution closer to the stationary distribution instead has the same effect on the node weights as using long random walks.

Node weights ρ as a hyperparameter. Instead of indirectly setting the node weights ρ through the random walk parameters q_0 and T as is done in NetGAN, we can treat ρ as a hyperparameter directly to incorporate beliefs about the graph. Weighting nodes according to the stationary distribution assigns equal weight to all edges in Eq. (12). In general, increasing the weight of a node encourages generated graphs to include its adjacent edges. This enables us, for example, to “protect” a certain set of nodes in the sense of preserving their neighborhoods in the generated graphs by increasing their weight.

6. Experiments

The purpose of this section is to (i) verify that CELL has performance comparable to NetGAN while being much faster, and (ii) demonstrate the importance of the cross-entropy loss and benefit of the logit-transformation by comparing with other low-rank approximation baselines.

6.1. Setup of the experiments

Data sets and preprocessing. We experiment on a variety of graph data sets: the citation networks CORA-ML (McCallum et al., 2000) and CITESEER (Sen et al., 2008), the political blogs network POLBLOGS (Adamic & Glance, 2005), the retweet network RT-GOP, and the web graph WEB-EDU (Gleich et al., 2004). All graphs except for CORA-ML are taken from Rossi & Ahmed (2015). For CORA-ML, we use the same preprocessed version as Bojchevski et al. (2018), an overview of the data sets is given

in Table 3. We preprocess the graphs by removing loops, edge weights, and edge directions. We then restrict them to their largest connected component to ensure that they are connected. For evaluating the link prediction performance during and after training, we split each graph into training-, validation-, and test-set by taking out 10% of the edges for validation and another 5% for testing, while ensuring that the remaining graph stays connected. The validation set is only used for the VAL-criterion, an alternative stopping criterion based on link prediction performance that is described in Section C.3 of the supplementary.

Baselines. We compare our model CELL to NetGAN (Bojchevski et al., 2018) and a number of non-parametric baselines: the configuration model, which simply rewires some randomly chosen edges (Molloy & Reed, 1995), and low-rank approximations of the adjacency matrix (LR-Adj), the random walk transition matrix (LR-Trans), the symmetric normalized Laplacian (LR-Lap), and the modularity matrix (LR-Mod), in a similar framework as described by Baldesi et al. (2018). To investigate the contribution of the logit transformation for CELL, we additionally consider a low-rank approximation of the transition matrix with respect to the cross-entropy loss instead of using the Frobenius norm (LR-CE). The original paper by Bojchevski et al. (2018) also compared to a number of parametric baselines, which have the purpose of explicitly fitting some hand-selected graph parameters, but fail to reproduce others. For brevity we do not report the results of these parametric baselines.

Setup and evaluation metrics. To make the results comparable, we train CELL and NetGAN until the same stopping criterion of 52% edge overlap with the input graph is satisfied. This is done by pausing the training at regular intervals, generating a single graph, and calculating the ratio of shared edges to input edges. While NetGAN is trained on a GPU, only a CPU is required for training CELL.

Our first evaluation metric is a set of common graph statistics for input and generated graphs, whose purpose is to measure the extent to which the newly generated graphs

Table 2. Graph statistics and link prediction performance on CORA-ML for generated graphs from NetGAN, our method CELL, and baselines, averaged over five trials. Statistics that are matched by model design for the configuration model are indicated as *, and cases that are not applicable as -. CELL produces statistics comparable to NetGAN, but is orders of magnitudes faster. This experiment is repeated for all other data sets in Section C.5 of the supplementary.

GRAPH	MAX. DEGREE	ASSORT- ACTIVITY	TRIANGLE COUNT	SQUARE COUNT	POWER LAW EXP.	CLUSTER- ING COEFF.	CHARAC. PATH LEN.	ROC-AUC SCORE	TIME (IN S)
CORA-ML	238	-0.076	2,802	14,268	1.86	8.26e-2	5.63	1	-
CONF. MODEL (52% EO)	*	-0.053	623	3111	*	1.96e-2	4.43	-	1
LR-ADJ (53% EO)	121	-0.042	444	1,128	1.72	2.78e-2	5.17	0.561	32
LR-TRANS (57% EO)	139	-0.058	558	1,617	1.77	2.94e-2	5.07	0.709	33
LR-LAP (52% EO)	167	-0.084	691	1942	1.79	2.79e-2	4.76	0.800	38
LR-MOD (53% EO)	122	-0.043	437	1,135	1.72	2.75e-2	5.17	0.557	48
LR-CE (52% EO)	193	-0.068	1,388	6,284	1.79	5.68e-2	5.37	0.950	73
NETGAN (54% EO)	219	-0.071	1,461	5,555	1.80	5.23e-2	5.13	0.950	7,478
CELL (53% EO)	204	-0.070	1,396	6,880	1.82	5.07e-2	5.26	0.938	21

Table 3. Data sets used. Nodes and edges refer to the largest connected component.

NAME	NODES	EDGES
CORA-ML	2,810	7,981
CITSEER	2,110	3,668
POLBLOGS	1,222	16,779
RT-GOP	4,687	5,529
WEB-EDU	3,031	6,474

reproduce network patterns of the input graph. Since memorizing the input graph trivially reproduces all of its graph statistics, we additionally evaluate the generalization properties in a link prediction task. To do so, we use the edges in the test set and an equal amount of randomly chosen non-edges from the original graph. After training, these are presented to the generative models, which try to classify them as existent or non-existent in the original graph on the basis of the score matrix (or, equivalently, the edge-independent model A^\dagger). This matrix is produced by all considered models except for the configuration model. A high value in the score matrix suggests the existence of the corresponding edge, while a low value suggests that the edge did not exist in the original graph. The performance is measured by the ROC-AUC score (Area Under Curve for Receiver Operating Characteristic curve), applied to the score matrix evaluated at the edges in question.

6.2. Evaluation

CELL vs. NetGAN. The results for graphs generated on CORA-ML are presented in Table 2. Compared to the other baselines, **CELL generates graphs with statistics close to those of NetGAN** and has similar link prediction performance; some of their small differences might be at-

tributed to the noise of the LSTM used by NetGAN, and to the different optimization procedures. The latter can be observed in Figure 4, which shows the evolution of generated graph statistics during training: NetGAN starts off with a different initialization, but as training continues, the generated graph statistics get close to the target well before memorizing the input graph. Further confirmation of this behavior is given in Sections C.6 and C.7 of the supplementary. However, the most striking difference is the training time, for which our method is **orders of magnitudes faster**, see Table 1.

CELL vs. baselines. Almost all **baselines fail to reproduce most of the graph statistics**, while CELL is reasonably close to all of them. Only LR-CE, the version of our method without the logit space, has performance very similar to CELL. This hints towards the **importance of the cross-entropy loss** rather than the logit space for successfully generalizing a graph. However, using the logit space still has the advantage of requiring only a small rank ($H = 9$ for CELL as compared to $H = 950$ for LR-CE), which results in less trainable parameters and shorter training time.

7. Discussion and future work

We derived a condensed version of NetGAN by identifying its essential steps and performing them directly. We verified experimentally that it retains the generalization performance of NetGAN, but is much faster. Additionally, our simple formulation of the algorithm makes it more accessible for analysis and application-specific extensions.

Analysis. In essence, we revealed the initial random-walk-based approach to be a low-rank approximation of the random walk transition matrix in the logit space. More naive low-rank approximations of matrices related to the input graph do not achieve competitive performance when ap-

proximating with respect to the Frobenius norm, but do so for the cross-entropy loss — a curious fact that we plan to investigate in future work. Based on our new, simplified methods we could analyze the inductive biases of the different components and the role of the parameters of NetGAN. For example, we discover that length and choice of start node of the random walks amount to a weighting of the nodes, which controls their importance in the graph generation process. Based on our better understanding of the bias, we can construct examples which both NetGAN and our algorithm cannot treat in a satisfactory manner.

Extensions. We demonstrated that our method is easily extendable by manipulating the loss function. An additional loss term can prevent the generation of edges we deem undesired, and node weights can emphasize user-specified nodes. Because learning step and reconstruction step are independent, each of them could be replaced by a different procedure. For example, instead of sampling from an edge-independent model, a more general method would sample independent paths to further emphasize locality.

Conclusion. Beyond the particular case of NetGAN, our work is part of a more high-level agenda. Machine learning is used in diverse applications, often not by machine learning experts, and the outcome of algorithms might have considerable impact in science and society. In such a context it is particularly important that our community actively attempts to understand the inherent inductive biases, strengths, and also the weaknesses of algorithms. Finding examples where an algorithm works is important — but maybe even more important is to understand under which circumstances the algorithm produces misleading results. For graph generative models, this might concern medical studies on brain graphs or geoscience studies on climate graphs. We should work hard to make our algorithms as transparent and interpretable as possible. This paper is a small step in that direction.

Acknowledgements

This work has been supported by the German Research Foundation through the Cluster of Excellence “Machine Learning – New Perspectives for Science” (EXC 2064/1 number 390727645), the BMBF Tübingen AI Center (FKZ: 01IS18039A), and the International Max Planck Research School for Intelligent Systems (IMPRS-IS).

References

Adamic, L. A. and Glance, N. The political blogosphere and the 2004 US election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery*, 2005.

Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein gen-

erative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning*, 2017.

Baldesi, L., Butts, C. T., and Markopoulou, A. Spectral graph forge: Graph generation targeting modularity. In *Conference on Computer Communications*, 2018.

Boguná, M., Pastor-Satorras, R., Díaz-Guilera, A., and Arenas, A. Models of social networks based on social distance attachment. In *Physical review E*, 2004.

Bojchevski, A., Shchur, O., Zügner, D., and Günnemann, S. NetGAN: Generating graphs via random walks. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.

Gleich, D. F., Zhukov, L., and Berkhin, P. Fast parallel PageRank: A linear system approach. Technical report, Yahoo! Research Labs, 2004.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27*, 2014.

Hochreiter, S. and Schmidhuber, J. Long short-term memory. In *Neural computation*, 1997.

Jalilifard, A., Caridá, V., Mansano, A., and Cristo, R. Can netgan be improved by short random walks originated from dense vertices? *arXiv preprint arXiv:1905.05298*, 2019.

Kingma, D. and Ba, J. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2014.

Lovász, L. et al. Random walks on graphs: A survey. In *Combinatorics, Paul erdos is eighty*, 1993.

McCallum, A. K., Nigam, K., Rennie, J., and Seymore, K. Automating the construction of internet portals with machine learning. In *Information Retrieval*, 2000.

Molloy, M. and Reed, B. A critical point for random graphs with a given degree sequence. In *Random Structures & Algorithms*, 1995.

Orlitsky, A. Estimating and computing density based distance metrics. In *Proceedings of the 22nd international conference on Machine learning*, 2005.

Rossi, R. A. and Ahmed, N. K. The network data repository with interactive graph analytics and visualization. In *Association for the Advancement of Artificial Intelligence*, 2015.

Sen, P., Namata, G., Bilgic, M., Getoor, L., Galligher, B., and Eliassi-Rad, T. Collective classification in network data. In *AI magazine*, 2008.

Tenenbaum, J. B., De Silva, V., and Langford, J. C. A global geometric framework for nonlinear dimensionality reduction. In *science*, 2000.

Zhou, Y., Cheng, H., and Yu, J. X. Graph clustering based on structural/attribute similarities. In *Proceedings of the VLDB Endowment*, 2009.