Residual2Vec: Debiasing graph embedding with random graphs

Anonymous Author(s)
Affiliation
Address
email

Abstract

Graph embedding maps a graph into a convenient vector-space representation for graph analysis and machine learning applications. Many graph embedding methods hinge on a sampling of context nodes based on random walks. However, random walks can be a biased sampler due to the structural properties of graphs. Most notably, random walks are biased by the degree of each node, where, at each step, a node is sampled proportionally to its degree. The implication of such biases has not been clear, particularly in the context of graph representation learning. Here, we investigate the impact of the random walks' bias on graph embedding and propose *residual2vec*, a general graph embedding method that can debias various structural biases in graphs by using random graphs. We demonstrate that this debiasing not only improves link prediction and clustering performance but also allows us to explicitly model salient structural properties in graph embedding.

1 Introduction

On average, your friends tend to be more popular than you. This is a mathematical necessity known as the *friendship paradox*, which arises due to a sampling bias, i.e., popular people have many friends and thus are likely to be on your friend list [1]. Beyond being a fun trivia, the friendship paradox is a fundamental property of graphs: following an edge is a biased sampling that preferentially samples nodes based on nodes' degree (i.e., the number of neighbors). The fact that random walk is used as the default sampling paradigm across many graph embedding methods raises important questions: what are the implications of this sampling bias in graph embedding? If it is undesirable, how can we debias it?

Graph embedding maps a graph into a dense vector representation, enabling a direct application of many machine learning algorithms to graph analysis [2]. A widely used framework is to turn a graph into a "sentence of nodes" and then feed the sentence to `word2vec` [3–6]. A crucial difference from word embedding is that, rather than using given sentences, graph embedding methods generate synthetic "sentences" from a given graph. In other words, the generation of synthetic "sentences" in graph embedding is an implicit modeling decision, which most methods take for granted. A common approach to generating sentences from a graph is based on random walks, which randomly traverse nodes by following edges. The friendship paradox comes into play when a walker follows an edge (e.g., friendship tie): it is more likely to visit a node with many neighbors (e.g., popular individual). As an example, consider a graph with a core-periphery structure, where core nodes have more neighbors than periphery (Fig. 1A). Although core nodes are in the minority, they become the majority in the sentences generated by random walks (Fig. 1B). This is because core nodes have more neighbors than periphery and thus are likely to be a neighbor of other

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

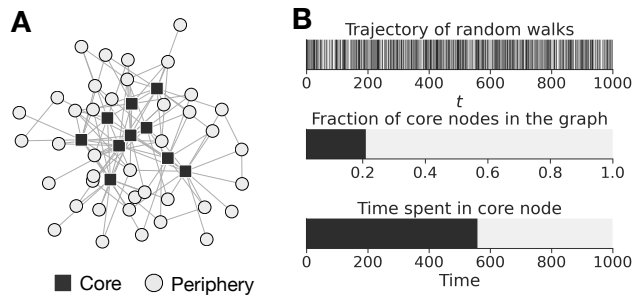

Figure 1: Random walks have a strong preference towards hubs. (**A**) A toy graph generated by a stochastic block model with core-periphery structure, where core nodes have more neighbors than peripheral nodes [7]. (**B**) Random walkers preferentially visit nodes with many neighbors, generating a trajectory that overrepresents the core nodes.

nodes, which is a manifestation of the friendship paradox. Then, how does the sampling bias affect the embedding?

Here, we show that `word2vec` by itself has an implicit bias arising from the optimization algorithm—skip-gram negative sampling (SGNS)—which happens to negate the bias due to the friendship paradox. To leverage this debiasing feature further, we propose a more general framework, `residual2vec`, that can also compensate for other systematic biases in random walks. We show that `residual2vec` performs better than conventional embedding methods in link prediction and community detection tasks. Using a citation network of 260k journals, we demonstrate that the biases from random walks overshadow the salient features of graphs. By removing the bias, `residual2vec` better captures the characteristics of journals such as the impact factor and journal subject.

2 Built-in debiasing feature of SGNS word2vec

2.1 Background: SGNS word2vec

Consider a sentence of words $(x_1, x_2, x_3, \ldots)$ composed of N unique words. `word2vec` associates the tth word x_t with words in its surrounding $x_{t-T}, \ldots, x_{t-1}, x_{t+1}, \ldots, x_{t+T}$, which are referred to as context words, determined by a prescribed window size T. For a center-context word pair (i, j), SGNS `word2vec` models conditional probability

$$P_{\mathrm{w2v}}(j|i) = \frac{\exp(\boldsymbol{u}_i^\top \boldsymbol{v}_j)}{\sum_{j'=1}^{N} \exp(\boldsymbol{u}_i^\top \boldsymbol{v}_{j'})}, \tag{1}$$

where $\boldsymbol{u}_i, \boldsymbol{v}_i \in \mathbb{R}^{K\times 1}$ are *embedding vectors* representing word i as center and context words, respectively, and K is the embedding dimension. An approach to fit P_{w2v} is the maximum likelihood estimation, which is computationally expensive because P_{w2v} involves the sum over all words. Alternatively, several heuristics have been proposed, among which *negative sampling* is the most widely used [3, 4, 6].

Negative sampling trains `word2vec` as follows. Given a sentence, a center-context word pair (i, j) is sampled and labeled as $Y_j = 1$. Additionally, one samples k random context word ℓ from a noise distribution $p_0(\ell)$, and then labels (i, ℓ) as $Y_\ell = 0$. In general, a popular choice of the noise distribution $p_0(\ell)$ is based on word frequency, i.e., $p_0(\ell) \propto P_{\mathrm{d}}(\ell)^\gamma$, where $P_{\mathrm{d}}(\ell)$ is the fraction of word ℓ in the given sentence, and γ is a hyper-parameter. Negative sampling trains $\boldsymbol{u}_i$ and $\boldsymbol{v}_j$ such that its label Y_j is well predicted by a logistic regression model

$$P_{\mathrm{NS}}(Y_j = 1; \boldsymbol{u}_i, \boldsymbol{v}_j) = \frac{1}{1 + \exp(-\boldsymbol{u}_i^\top \boldsymbol{v}_j)}, \tag{2}$$

by maximizing its log-likelihood.

2.2 Implicit debiasing by negative sampling

Negative sampling efficiently produces a good representation [6]. However, an often overlooked fact is that it is a biased estimator unless $\gamma = 0$ [8, 9]. We show that this statistical bias gives rise to a debiasing feature of SGNS `word2vec`.

To see this, let us consider the unbiased variant of negative sampling, *Noise Contrastive Estimation* (NCE) [8]. NCE is a generic estimator for probability model P_m of the form [8]

$$P_m(x) = \frac{f(x;\theta)}{\sum_{x' \in \mathcal{X}} f(x';\theta)}, \tag{3}$$

where f is a non-negative function of data x in all data $\mathcal{X}$. `word2vec` (Eq. (1)) is a special case of P_m, where $f(x) = \exp(x)$ and $x = \boldsymbol{u}_i^\top \boldsymbol{v}_j$. NCE estimates P_m by training a binary logistic regression model, as is the case for negative sampling. However, NCE is unbiased estimator of P_m. Having sampled word pairs, we know that 1 in $1+k$ pairs are taken from the given data, which can be expressed as prior probabilities

$$P(Y_j = 1) = \frac{1}{k+1}, \quad P(Y_j = 0) = \frac{k}{k+1}. \tag{4}$$

Assuming that the given data is generated from P_m, we have

$$P(j|Y_j = 1) = P_m(\boldsymbol{u}_i^\top \boldsymbol{v}_j), \quad P(j|Y_j = 0) = p_0(j). \tag{5}$$

Substituting Eqs. (4) and (5) into the Bayes rule yields posterior

$$P_{\mathrm{NCE}}(Y_j = 1|j) = \frac{P(j|Y_j = 1)\,P(Y_j = 1)}{\sum_{y \in \{0,1\}} P(j|Y_j = y)\,P(Y_j = y)} = \frac{P_m(\boldsymbol{u}_i^\top \boldsymbol{v}_j)}{P_m(\boldsymbol{u}_i^\top \boldsymbol{v}_j) + kp_0(j)}, \tag{6}$$

which can be rewritten with a sigmoid function as

$$P_{\mathrm{NCE}}(Y_j = 1|j) = \frac{1}{1 + kp_0(j)/P_m(\boldsymbol{u}_i^\top \boldsymbol{v}_j)} = \frac{1}{1 + \exp\left[-\ln f(\boldsymbol{u}_i^\top \boldsymbol{v}_j) + \ln p_0(j) + c\right]}, \tag{7}$$

where $c = \ln k + \ln \sum_{x' \in \mathcal{X}} f(x')$ is a constant. Now, let us revisit negative sampling from the perspective of NCE. We rewrite P_{NS} (i.e., Eq. (2)) in the form of Eq. (7) as

$$\begin{aligned} P_{\mathrm{NS}}(Y_j = 1; \boldsymbol{u}_i, \boldsymbol{v}_j) &= \frac{1}{1 + \exp(-\boldsymbol{u}_i^\top \boldsymbol{v}_j)} = \frac{1}{1 + \exp\left[-\left(\boldsymbol{u}_i^\top \boldsymbol{v}_j + \ln p_0(j) + c\right) + \ln p_0(j) + c\right]} \\ &= \frac{1}{1 + \exp\left[-\ln f(\boldsymbol{u}_i^\top \boldsymbol{v}_j) + \ln p_0(j) + c\right]}, \end{aligned} \tag{8}$$

where $f(\boldsymbol{u}_i^\top \boldsymbol{v}_j) = \exp\left(\boldsymbol{u}_i^\top \boldsymbol{v}_j + \ln p_0(j) + c\right)$. By substituting $f(\boldsymbol{u}_i^\top \boldsymbol{v}_j)$ into Eq. (3), SGNS `word2vec` is an unbiased estimator for probability model [10]:

$$P_{\mathrm{w2v}}^{\mathrm{SGNS}}(j \mid i) = \frac{p_0(j)\exp(\boldsymbol{u}_i^\top \boldsymbol{v}_j)}{Z_i'}, \quad \text{where } Z_i' := \sum_{j'=1}^{N} p_0(j')\exp(\boldsymbol{u}_i^\top \boldsymbol{v}_{j'}). \tag{9}$$

Equation (9) clarifies the role of noise distribution p_0. Noise probability p_0 serves as a baseline for $P_{\mathrm{w2v}}^{\mathrm{SGNS}}$, and word similarity $\boldsymbol{u}_i^\top \boldsymbol{v}_j$ represents the *deviation* from $p_0(j)$, or equivalently the characteristics of words *not* captured in $p_0(j)$. Notably, baseline $p_0(j) \propto P_{\mathrm{d}}(j)^\gamma$ is determined by word frequency $X_{\mathrm{d}}(j)$ and thus negates the word frequency bias. This realization—that we can explicitly use a noise distribution to obtain "residual" information—is the motivation for our method, *residual2vec*.

3 Residual2vec graph embedding

In this section, we describe `residual2vec` graph embedding. We assume that the given graph is undirected and weighted although our results can be generalized to directed graphs.

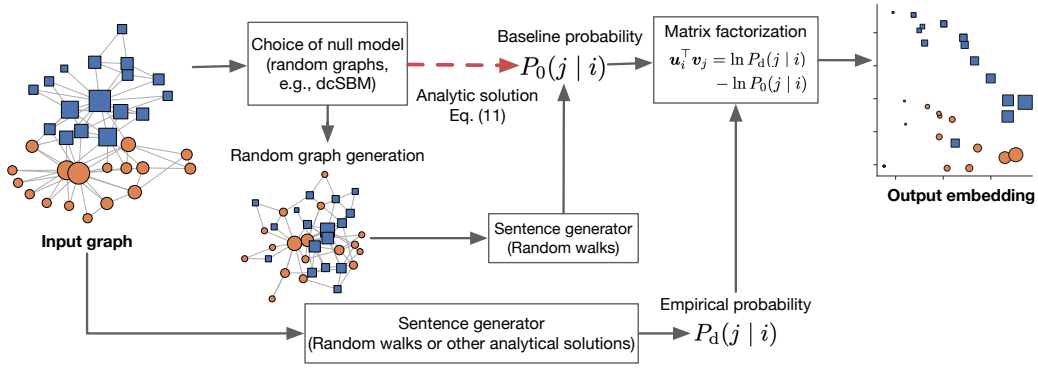

Figure 2: `residual2vec` framework. We first choose an explicit "null" model (e.g., dcSBM random graph). By running random walks in the input and random graphs, "sentences of nodes" are generated and then compared to extract the residual information not in the random graphs. The final embedding encodes the residual information. Probabilities $P_{\mathrm{d}}(j \mid i)$ and $P_0(j \mid i)$ can be computed analytically in many cases. For instance, we use the analytical solution for dcSBM (i.e., $P_{\mathrm{d}}(j \mid i)$ and $P_0(j \mid i)$) instead of simulating actual random walks.

3.1 Model

The presence of $p_0(j)$ effectively negates the bias in random walks due to degree. This bias dictates that, for a sufficiently long trajectory of random walks in undirected graphs, the frequency $P_{\mathrm{d}}(j)$ of node j is proportional to degree d_j (i.e., the number of neighbors) irrespective of the graph structure [11]. Now, if we set $\gamma = 1$, baseline $p_0(j) = P_{\mathrm{d}}(j)$ matches exactly with the node frequency in the trajectory, negating the bias due to degree. But, we are free to choose any p_0. This consideration leads us to *residual2vec* model:

$$P_{\mathrm{r2v}}(j \mid i) = \frac{P_0(j \mid i) \exp(\boldsymbol{u}_i^\top \boldsymbol{v}_j)}{Z'_i}, \tag{10}$$

where we explicitly model baseline transition probability denoted by $P_0(j \mid i)$. In doing so, we can obtain the *residual* information that is not captured in P_0. Figure 2 shows the framework of `residual2vec`. To negate a bias, we consider "null" graphs, where edges are randomized while keeping the property inducing the bias intact [12, 13]. Then, we compute P_0 either analytically or by running random walks in the null graphs. The P_0 is then used as the noise distribution to train SGNS `word2vec`.

Random graph models Among many models for random graph [14–20], here we focus on the degree-corrected stochastic block model (dcSBM), which can be reduced to many fundamental random graph models with certain parameter choices [17]. With the dcSBM, one partitions nodes into B groups and randomizes edges while preserving (i) the degree of each node, and (ii) the number of inter-/intra-group edges. Preserving such group connectivity is useful to negate biases arising from less relevant group structure such as bipartite and multilayer structures. The dcSBM can be mapped to many random graph models. In fact, when $B = 1$, the dcSBM is reduced to *the configuration model* that preserves the degree of each node [20]. Furthermore, by setting $B = 1$ and $d_i =$ constant, the dcSBM is reduced to *the Erdős-Rényi random graph* that only preserves the number of edges.

With the dcSBM, we can calculate $P_0(j \mid i)$ without explicitly generating random graphs:

$$P_0(j \mid i) = \frac{d_j}{D_{g_j}} \left(\frac{1}{T} \sum_{t=1}^{T} \mathbf{P}_{\mathrm{SBM}}^t \right)_{g_i, g_j}, \tag{11}$$

where node j has degree d_j and belongs to group g_j. There are $D_g = \sum_{\ell=1}^{N} d_\ell \delta(g_\ell, g)$ edges emanating from group g. The entry $P_{g,g'}^{\mathrm{SBM}}$ of matrix $\mathbf{P}_{\mathrm{SBM}} = (P_{g,g'}^{\mathrm{SBM}}) \in \mathbb{R}^{B \times B}$ is the fraction of edges to group g' in D_g. Table 1 lists $P_0(j \mid i)$ for the special classes of the dcSBM. See Supplementary Information for the step-by-step derivation.

Table 1: Baseline probability $P_0(j \mid i)$ for the dcSBM along with its special cases.

	Random graph models		
	Erdős-Rényi model [14]	The configuration model [20]	dcSBM [17]
$P_0(j \mid i)$	$1/N$	$d_j / \sum_{\ell=1}^{N} d_\ell$	$d_j D_{g_j}^{-1} \left(T^{-1} \sum_{t=1}^{T} \mathbf{P}_{\text{SBM}}^t \right)_{g_i, g_j}$
Algorithm	`DeepWalk` [3]	`node2vec` ($\gamma = 0.75$) [4], `NetMF` ($\gamma = 1$) [21]	

Other graph embedding as special cases of residual2vec `residual2vec` can be considered as a general framework to understand structural graph embedding methods because many existing graph embedding methods are special cases of `residual2vec`. `node2vec` and `NetMF` use SGNS `word2vec` with $p_0(j) \propto P_{\text{d}}(j)^\gamma$. This $p_0(j)$ is equivalent to the baseline for the configuration model [20], where each node i has degree d_i^γ. `DeepWalk` is also based on `word2vec` but trained with an unbiased estimator (i.e., the hierarchical softmax). Because negative sampling with $p_0(j) = 1/N$ is unbiased [8, 9], `DeepWalk` is equivalent to `residual2vec` with Erdős-Rényi random graphs as the null model.

3.2 Residual2vec as matrix factorization

Many structural graph embedding methods implicitly factorize a matrix $\mathbf{R}$ where entry R_{ij} corresponds to dot product $\boldsymbol{u}_i^\top \boldsymbol{v}_j$ [21]. `residual2vec` can also be described as factorizing a matrix $\mathbf{R}$ which captures residual pointwise mutual information. Just like a Kth order polynomial function can be fit to K points without errors, we can fit `word2vec` to the given data without errors when the embedding dimension K is equal to the number of unique words N [19]. In other words, $P_{\text{d}}(j \mid i) = P_{\text{r2v}}(j \mid i)$, $\forall i, j$ if $K = N$. By substituting $P_{\text{r2v}}(j \mid i) = P_0(j \mid i) \exp(\boldsymbol{u}_i^\top \boldsymbol{v}_j)/Z_i'$, we obtain

$$P_{\text{d}}(j \mid i) = \frac{P_{\text{d}}(j \mid i) \exp\left(\boldsymbol{u}_i^\top \boldsymbol{v}_j + \ln P_0(j \mid i) - \ln P_{\text{d}}(j \mid i)\right)}{\sum_{j'=1}^{N} P_{\text{d}}(j' \mid i) \exp\left(\boldsymbol{u}_i^\top \boldsymbol{v}_{j'} + \ln P_0(j' \mid i) - \ln P_{\text{d}}(j' \mid i)\right)}. \tag{12}$$

The equality holds if $\exp(\boldsymbol{u}_i^\top \boldsymbol{v}_j + \ln P_0(j \mid i) - \ln P_{\text{d}}(j \mid i)) = c_i$ for all i and j, where c_i is a constant. The solution is not unique because c_i can be any real value. We choose $c_i = 1$ to obtain a solution in the simplest form, yielding matrix $\mathbf{R}$ that `residual2vec` factorizes:

$$R_{ij} := \boldsymbol{u}_i^\top \boldsymbol{v}_j = \ln P_{\text{d}}(j \mid i) - \ln P_0(j \mid i). \tag{13}$$

Matrix $\mathbf{R}$ has an information-theoretic interpretation. We rewrite

$$R_{ij} = \ln \frac{P_{\text{d}}(i,j)}{P_{\text{d}}(i)} - \ln \frac{P_0(i,j)}{P_0(i)} = \ln \frac{P_{\text{d}}(i,j)}{P_{\text{d}}(i)P_{\text{d}}(j)} - \ln \frac{P_0(i,j)}{P_0(i)P_0(j)} + \ln P_{\text{d}}(j) - \ln P_0(j). \tag{14}$$

The dcSBM preserves the degree of each node and thus has the same the degree bias with the given graph, i.e., $P_{\text{d}}(i) = P_0(i)$. Therefore, $\ln P_{\text{d}}(j) - \ln P_0(j) = 0$ and thus we obtain

$$R_{ij} = \text{PMI}_{\sim P_{\text{d}}}(i,j) - \text{PMI}_{\sim P_0}(i,j), \text{ where } \text{PMI}_{\sim P}(i,j) = \ln \frac{P(i,j)}{P(i)P(j)}. \tag{15}$$

$\text{PMI}_{\sim P}(i,j)$ is the pointwise mutual information that measures the correlation between center i and context j under joint distribution $P(i,j)$, i.e., $\text{PMI}_{\sim P}(i,j) = 0$ if i and j appear independently, and $\text{PMI}_{\sim P}(i,j) > 0$ otherwise. In sum, R_{ij} reflects residual pointwise mutual information of i and j from the null model.

3.3 Efficient matrix factorization

Although we assume that $N = K$ above, in practice, we want to find a compact vector representation (i.e., $K \ll N$) that still yields a good approximation [19, 21]. There are several computational challenges in factorizing $\mathbf{R}$. First, $\mathbf{R}$ is ill-defined for any node pair (i,j) that never appears because $R_{ij} = \ln 0 = -\infty$. Second, $\mathbf{R}$ is often a dense matrix with $\mathcal{O}(N^2)$ space complexity. For these issues, a common remedy is a truncation [19, 21]:

$$\tilde{R}_{ij} := \max(R_{ij}, 0). \tag{16}$$

Table 2: List of empirical graphs and structural properties. Variables N and M indicate the number of nodes and edges, respectively.

Network	N	M	Assortativity	Max. degree	Clustering coef.	Ref.
Airport	3,188	18,834	-0.02	911	0.493	[22]
Protein-Protein	3,852	37,840	-0.10	593	0.148	[23]
Wikipedia vote	7,066	100,736	-0.08	1065	0.142	[24]
Coauthorship (HepTh)	8,638	24,827	0.24	65	0.482	[25]
Citation (DBLP)	12,494	49,594	-0.05	713	0.118	[26]
Coauthorship (AstroPh)	17,903	197,031	0.20	504	0.633	[25]

This truncation discards negative node associations ($R_{ij} < 0$) while keeping the positive associations ($R_{ij} > 0$) based on the idea that negative associations are common, and thus are less informative [21]. In both word and graph embeddings, the truncation substantially reduces the computation cost of the matrix factorization [19, 21].

We factorize $\tilde{\mathbf{R}} = (\tilde{R}_{ij})$ into embedding vectors $\boldsymbol{u}_i$ and $\boldsymbol{v}_j$ such that $\boldsymbol{u}_i^\top \boldsymbol{v}_j \simeq \tilde{R}_{ij}$ by using the truncated singular value decomposition (SVD). Specifically, we factorize $\tilde{\mathbf{R}}$ by $\tilde{\mathbf{R}} = \mathbf{\Phi}_{\text{left}} \cdot \text{diag}(\sigma_1, \sigma_2, \ldots, \sigma_K) \cdot \mathbf{\Phi}_{\text{right}}^\top$, where $\mathbf{\Phi}_{\text{left}} = (\phi_{ik}^{\text{left}})_{ik}$ and $\mathbf{\Phi}_{\text{right}} = (\phi_{ik}^{\text{right}})_{ik} \in \mathbb{R}^{N \times K}$ are the left and right singular vectors of $\tilde{\mathbf{R}}$ associated with the K largest singular values $(\sigma_1, \ldots \sigma_K)$ in magnitude, respectively. Then, we compute $\boldsymbol{u}_i$ and $\boldsymbol{v}_j$ by $u_{ik} = \sigma_k^\alpha \phi_{ik}^{\text{left}}$, $v_{ik} = \sigma_k^{1-\alpha} \phi_{ik}^{\text{right}}$ with $\alpha = 0.5$ following the previous studies [19, 21].

Computing $P_{\text{d}}(j \mid i)$ analytically is computaionally expensive because it scales as $\mathcal{O}(TN^3)$, where T is the window size [21]. Alternatively, one can simulate random walks to estimate $P_{\text{d}}(j \mid i)$. Yet, both approaches require $\mathcal{O}(N^2)$ space complexity. Here, we reduce the time and space complexity by approximating the given graph by the dcSBM with $\hat{B}$ groups (we use $\hat{B} = 1,000$ in this paper) and then compute an approximated $P_{\text{d}}(j \mid i)$. This approximation reduces the time and space complexity to $\mathcal{O}(MN\hat{B} + T\hat{B}^3)$ and $\mathcal{O}(\hat{B}^2 + N)$, respectively, where M is the number of edges, with a high accuracy (the average Pearson correlation of 0.85 for the graphs in Table 2). See Supplementary Information for details.

4 Results

We test `residual2vec` using link prediction and community detection benchmarks [4, 12, 27–29]. We use the configuration model [20] as the null graph for `residual2vec`, denoted by `r2v-config`, which yields a degree-debiased embedding. As baselines, we use (i) two random-walk-based methods, `node2vec` [4] and `DeepWalk` [3], (ii) two matrix-factorization based methods, `Glove` [30] and Laplacian eigenmap (`LEM`) [31], and (iii) the graph convolutional network (`GCN`) [32] and `GraphSAGE` [33]. For all random-walk-based methods, we run 10 walkers per node for 80 steps and set $T = 10$ and training iterations to 5. We set the parameters of `node2vec` by $p = 1$ and $q \in \{0.5, 1, 2\}$. For `Glove`, we input the sentences generated by random walks. We use two-layer `GCN` and `GraphSAGE` implemented in StellarGraph package [34] with the parameter sets (e.g., the number of layers and activation function) used in [32, 33]. Because node features are not available in the benchmarks, we alternatively use degree and eigenvectors for the $\hat{K}$ smallest eigenvalues of the normalized Laplacian matrix because they are useful for link prediction and clustering [35, 36]. We set $\hat{K}$ to embedding dimension K (i.e., $\hat{K} = K$). Increasing $\hat{K}$ does not improve performance much (Supplementary Information).

Link prediction Link prediction task is to find missing edges based on graph structure, a basic task for various applications such as recommending friends and products [4, 29, 37]. The link prediction task consists of the following three steps. First, given a graph, a fraction ($\rho = 0.5$) of edges are randomly removed. Second, the edge-removed graph is embedded using a graph embedding method. Third, the removed edges are predicted based on a likelihood score calculated based on the generated embedding. In the edge removal process, we keep edges in a minimum spanning tree of the graph to ensure that the graph is a connected

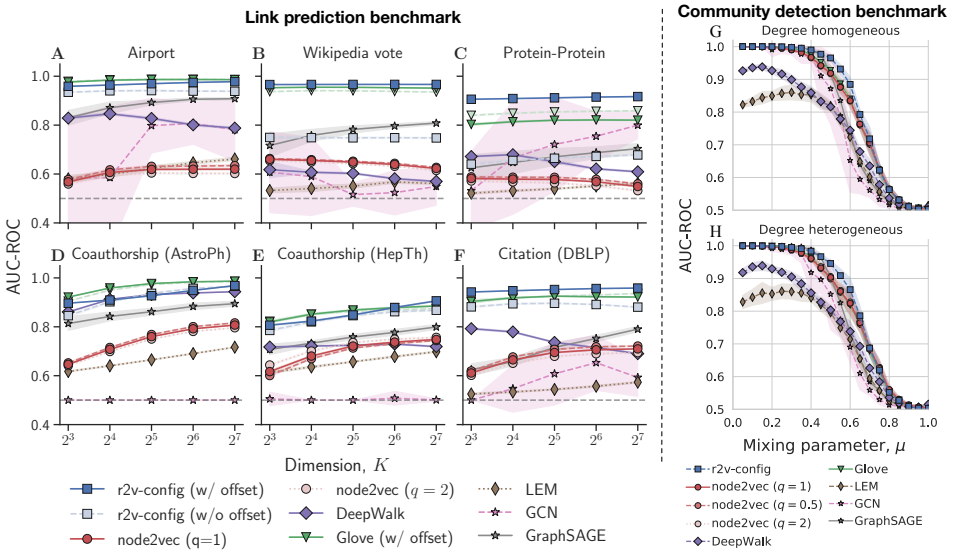

Figure 3: Performance for link prediction and community detection. We average AUC-ROC values over 30 runs with different random seeds and compute the 95% confidence interval by a bootstrapping with 10^4 resamples. `r2v-config` consistently performs the best or nearly the best for all graphs in both benchmarks, as indicated by the higher AUC-ROC values.

component [4, 29]. This is because predicting edges between disconnected graphs is an ill-defined task because each disconnected component has no relation to the other.

We compute the prediction score by $\boldsymbol{u}_i^\top \boldsymbol{u}_j + z_i + z_j$, where z_i represents an offset for node similarity. Only `Glove` and `r2v-config` have offset z_i, which corresponds to the bias term for `Glove` and $z_i = \ln d_i$ for `r2v-config`. Other methods do not have offset similarity and thus $z_i = 0$ for all i. We measure the performance by the area under the curve of the receiver operating characteristics (AUC-ROC) for the prediction scores, with the removed edges and the same number of randomly sampled non-existent edges being the positive and negative classes, respectively. We perform the benchmark for the graphs in Table 2.

`r2v-config` performs the best or nearly the best for all graphs (Figs. 3A–F). It consistently outperforms other random walk-based methods in all cases despite the fact that `node2vec` and `r2v-config` train the same model. The two methods have two key differences. First, `r2v-config` uses baseline $P_0(\ell \mid i) = P_{\mathrm{d}}(\ell)$, whereas `node2vec` uses $P_0(\ell \mid i) \propto P_{\mathrm{d}}(\ell)^{3/4}$ that does not exactly fit to the degree bias. Second, `r2v-config` optimizes the model based on a matrix factorization, which often yields a better embedding than the stochastic gradient descent algorithm used in `node2vec` [19, 21]. The performance of `residual2vec` is substantially improved when incorporating offset z_i for prediction, demonstrating that understanding and leveraging inherent bias in methods are critical for link prediction.

Community detection We use the Lancichinetti–Fortunato–Radicchi (LFR) community detection benchmark [27]. The LFR benchmark generates graphs having groups of densely connected nodes (i.e., communities) with a power-law degree distribution with a prescribed exponent τ. In the LFR benchmark, each node has, on average, a specified fraction μ of neighbors in different communities. We generate graphs of $N = 1,000$ nodes with the parameters used in Ref. [27] and embed the graphs to $K = 64$ dimensional space. We evaluate the performance by randomly sampling $10,000$ node pairs and calculate the AUC-ROC for their cosine similarities, with nodes in the same and different communities being the positive and negative classes, respectively. A large AUC value indicates that nodes in the same community have a higher similarity than those in different communities.

As μ increases from zero, the AUC for all methods decreases because nodes have more neighbors in different communities. `DeepWalk`, `LEM`, and `GCN` have a small AUC value even at $\mu = 0.05$. For all μ values, `r2v-config` achieves the highest AUC.

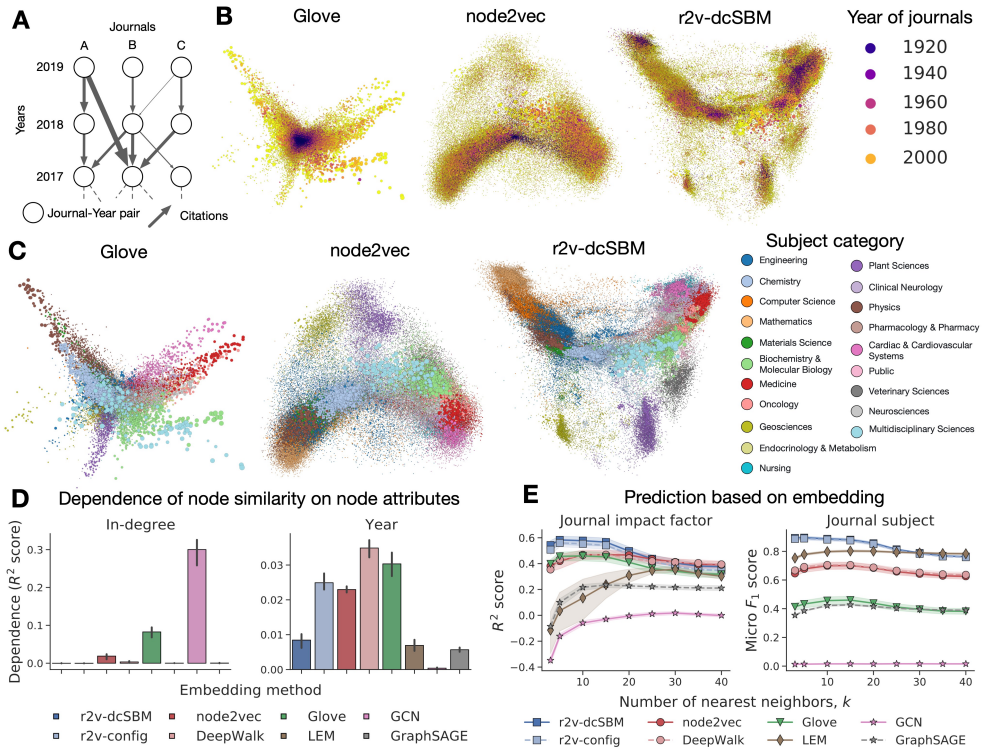

Figure 4: Embedding of the WOS journal citation graph. (**A**) Each node represents a pair (j, t) of journal j and year t, and each edge is weighted by citations. (**B**, **C**) A 2d projection of 128-dimensional embedding by the LDA. (**D**) `r2v-dcSBM` produces the embedding that is less dependent on degree and time. (**E**) By using the k-nearest neighbor algorithm, the embedding by `r2v-dcSBM` best predicts the impact factor and subject category.

Case study Can debiasing reveal the salient structure of graphs more clearly? We construct a journal citation graph using citation data between 1900 and 2019 indexed in the Web of Science (WoS) (Fig. 4A). Each node represents a pair (i, t) of journal i and year t. Each undirected edge between (i, t) and (j, t') is weighted by the number of citations between journal i in year t and journal j in year t'. The graph consists of $242,789$ nodes and $254,793,567$ undirected and weighted edges. Because the graph has a high average degree (i.e., $1,049$), some algorithms are computationally demanding. For this reason, we omit `node2vec` with $q = 0.5$ and $q = 2$ due to memory shortage. Furthremore, for `GCN`, we set $\hat{K} = 32$, which still took more than 18 hours. We also perform the whole analysis for the directed graph to respect the directionality of citations. Although all methods perform worse in predicting impact factor and subject category, we find qualitatively the same results. See Supplementary Information for the results for the directed graph.

Here, in addition to the degree bias, there are also temporal biases, e.g., there has been an exponential growth in publications, older papers had more time to accumulate citations, and papers tend to cite those published in prior few years [38]. To remove both biases, we use `residual2vec` with the dcSBM (denoted by `r2v-dcSBM`), where we group journals by year to randomize edges while preserving the number of citations within and between years. We generate $K = 128$ dimensional embeddings with $T = 10$.

Figures 4B and C show the 2d projection by the Linear Discriminant Analysis (LDA) with journals' subject categories as the class labels. `Glove` and `node2vec` suffer from the temporal structure, capturing temporal information prominently. By contrast, `r2v-dcSBM` is less affected by the temporal structure, delineating journals' subjects more clearly.

We quantify the effect of the bias by randomly sampling $50,000$ node pairs (i, j) and then fitting a linear regression model $y_{ij} = w_0 + w_1(x_i + x_j) + w_2|x_i - x_j| + w_3 x_i x_j$ that predicts

cosine similarity y_{ij} for node pair (i, j) with attributes x_i and x_j, where x_i is either the degree or the year of node i. We perform 5-cross validations and compute R^2-score (Fig. 4D). A smaller R^2-score indicates that node similarity is less dependent on node attributes and thus less biased. `LEM` has the smallest R^2-score for both degree and year. `r2v-dcSBM` has a smaller R^2-score than `r2v-config` for year, respectively, suggesting that `r2v-dcSBM` successfully negates the biases due to time.

Is debiasing useful to capture the more relevant structure of graphs? We use embedding vectors to predict journal's impact factor (IF) and subject category. By employing the k-nearest neighbor algorithm, we carry out 5-cross validations and measure the prediction performance by R^2-score and the micro-F1 score. `r2v-dcSBM` best captures the disciplines and impact of journals (Fig. 4E). This result demonstrates that debiasing embedding can reveal the salient structure of graphs that is overshadowed by other systematic biases.

5 Discussion

In this paper, starting from the insight that `word2vec` with SGNS has a built-in debiasing feature that cancels out the bias due to the degree of nodes, we generalize this debiasing feature further, proposing a method that can selectively remove any structural biases that are modeled by a null random graph. By exposing the bias and explicitly modeling it, we provide a new way to integrate prior knowledge about graphs into graph embedding, and a unifying framework to understand structural graph embedding methods. Under our residual2vec framework, other structural graph embedding methods that use random walks can be understood as special cases with different choices of null models. Through empirical evaluations, we demonstrate that debiasing improves link prediction and community detection performances, and better reveals the characteristics of nodes, as exemplified in the embedding of the WoS journal citation graph.

Our method is highly flexible because any random graph model can be used. Although we focus on two biases arising from degree and group structure in a graph, one can remove other biases such as the degree-degree correlation, clustering, and bipartitivity by considering appropriate null graphs. Beyond these statistical biases, there have been growing concerns about social bias (e.g., gender stereotype) as well as surveillance and privacy in AI applications, which prompted the study of gender and frequency biases in word embedding [39–42]. The flexibility and power of our selective and explicit debiasing approach may also be useful to address such biases that do not originate from common graph structures.

There are several limitations in `residual2vec`. We assume that random walks have a stationary distribution, which may not be the case for directed graphs. One can ensure the stationarity in random walks by randomly teleporting walkers [43]. Second, it is not yet clear to what extent debiasing affects downstream tasks (e.g., by losing information about the original graph). Nevertheless, we believe that the ability to understand and control systematic biases is critical to model graphs through the prism of embedding.

Broader Impact

There has been an ever-increasing concern on inappropriate social stereotyping and the leak of sensitive information in word and graph embeddings [39, 44]. Although we have not studied social biases in this paper, given the wide usage of graph embedding methods to model social data, our approach may lead to methods and studies that expose and mitigate social biases that manifest as structural properties in graph datasets. Our general idea and approach may also be applied to modeling natural language and may contribute to the study of biases in language models. At the same time, by improving the accuracy of graph embedding, our method may also have negative impacts such as privacy attacks and exploitation of personal data (surveillance capitalism) [44, 45]. Nevertheless, we believe that our approach contributes to the effort to create transparent and accountable machine learning methods, especially because our method enables us to explicitly model what is structurally expected.

References

1. Feld, S. L. Why Your Friends Have More Friends Than You Do. *American Journal of Sociology* **96,** 1464–1477. https://doi.org/10.1086/229693 (1991).
2. Cai, H., Zheng, V. W. & Chang, K. C.-C. A Comprehensive Survey of Graph Embedding: Problems, Techniques, and Applications. *IEEE Transactions on Knowledge and Data Engineering* **30,** 1616–1637 (2018).
3. Perozzi, B., Al-Rfou, R. & Skiena, S. *DeepWalk: Online Learning of Social Representations* in *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (ACM Press, New York, New York, USA, 2014), 701–710. ISBN: 9781450329569.
4. Grover, A. & Leskovec, J. *Node2Vec: Scalable Feature Learning for Networks* in *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (ACM, San Francisco, California, USA, 2016), 855–864.
5. Dong, Y., Chawla, N. V. & Swami, A. *metapath2vec* in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* **Part F1296** (ACM, New York, NY, USA, 2017), 135–144. ISBN: 9781450348874.
6. Mikolov, T., Sutskever, I., Chen, K., Corrado, G. & Dean, J. Distributed Representations of Words and Phrases and their Compositionality. *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing,* 1389–1399. eprint: 1310.4546 (2013).
7. Zhang, X., Martin, T. & Newman, M. E. J. Identification of core-periphery structure in networks. *Physical Review E* **91,** 032803. ISSN: 1539-3755. http://link.aps.org/doi/10.1103/PhysRevE.91.032803 (2015).
8. Gutmann, M. & Hyvärinen, A. *Noise-contrastive estimation: A new estimation principle for unnormalized statistical models* in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics* (eds Teh, Y. W. & Titterington, M.) **9** (PMLR, Chia Laguna Resort, Sardinia, Italy, 2010), 297–304.
9. Dyer, C. Notes on Noise Contrastive Estimation and Negative Sampling. arXiv: 1410.8251 (2014).
10. Murray, D. *et al. Unsupervised embedding of trajectories captures the latent structure of mobility* 2020. arXiv: 2012.02785 [cs.LG].
11. Masuda, N., Porter, M. A. & Lambiotte, R. Random walks and diffusion on networks. *Physics Reports* **716-717,** 1–58. ISSN: 0370-1573. eprint: 1612.03281 (2017).
12. Fortunato, S. Community detection in graphs. *Physics Reports* **486,** 75–174. ISSN: 03701573 (2010).
13. Newman, M. E. J. Networks: An introduction. *Oxford University,* 163–186. ISSN: 18600832. arXiv: 1212.2425 (2014).
14. Erdős, P & Rényi, A. On random graphs. *Transactions of the American Mathematical Society* **286,** 257. ISSN: 00029947 (1984).
15. Karrer, B. & Newman, M. E. J. Random graph models for directed acyclic networks. *Phys. Rev. E* **80,** 46110 (2009).
16. Garlaschelli, D. & Loffredo, M. I. Generalized Bose-Fermi Statistics and Structural Correlations in Weighted Networks. *Phys. Rev. Lett.* **102,** 38701 (2009).
17. Karrer, B. & Newman, M. E. J. Stochastic blockmodels and community structure in networks. *Physical Review E* **83,** 016107. ISSN: 1539-3755. http://link.aps.org/doi/10.1103/PhysRevE.83.016107 (2011).
18. Expert, P., Evans, T. S., Blondel, V. D. & Lambiotte, R. Uncovering space-independent communities in spatial networks. *Proceedings of the National Academy of Sciences* **108,** 7663–7668 (2011).
19. Levy, O. & Goldberg, Y. Neural word embedding as implicit matrix factorization. *Advances in Neural Information Processing Systems* **3,** 2177–2185. ISSN: 10495258 (2014).
20. Fosdick, B, Larremore, D, Nishimura, J & Ugander, J. Configuring Random Graph Models with Fixed Degree Sequences. *SIAM Review* **60,** 315–355 (2018).

21. Qiu, J. *et al.* Network embedding as matrix factorization: Unifying DeepWalk, LINE, PTE, and Node2vec. *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining. WSDM '18* 459–467. eprint: 1710.02971 (2018).
22. Opsahl, T. *Why Anchorage is not (that) important: Binary ties and Sample selection* (2016).
23. Stark, C. *et al.* BioGRID: a general repository for interaction datasets. *Nucleic Acids Research* **34,** D535–D539. ISSN: 0305-1048. https://doi.org/10.1093/nar/gkj109 (2006).
24. Leskovec, J., Huttenlocher, D. & Kleinberg, J. Governance in Social Media: A Case Study of the Wikipedia Promotion Process. *Proceedings of the International AAAI Conference on Web and Social Media* **4** (2010).
25. Leskovec, J., Kleinberg, J. & Faloutsos, C. Graph Evolution: Densification and Shrinking Diameters. *ACM Trans. Knowl. Discov. Data* **1.** ISSN: 1556-4681 (2007).
26. Ley, M. *The DBLP Computer Science Bibliography: Evolution, Research Issues, Perspectives BT - String Processing and Information Retrieval* in (eds Laender, A. H. F. & Oliveira, A. L.) (Springer Berlin Heidelberg, Berlin, Heidelberg, 2002), 1–10. ISBN: 978-3-540-45735-0.
27. Lancichinetti, A. & Fortunato, S. Community detection algorithms: A comparative analysis. *Phys. Rev. E* **80,** 56117 (2009).
28. Newman, M. E. J. Finding community structure in networks using the eigenvectors of matrices. *Physical Review E* **74,** 036104. ISSN: 1539-3755. eprint: 0605087 (physics) (2006).
29. Abu-El-Haija, S., Perozzi, B. & Al-Rfou, R. *Learning Edge Representations via Low-Rank Asymmetric Projections* in *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management* (ACM, Singapore, Singapore, 2017), 1787–1796.
30. Pennington, J., Socher, R. & Manning, C. *Glove: Global Vectors for Word Representation* in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* **19** (Association for Computational Linguistics, Stroudsburg, PA, USA, 2014), 1532–1543. https://www.aclweb.org/anthology/D14-1162http://aclweb.org/anthology/D14-1162.
31. Belkin, M. & Niyogi, P. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation* **15,** 1373–1396. ISSN: 08997667 (2003).
32. Kipf, T. N. & Welling, M. *Semi-Supervised Classification with Graph Convolutional Networks* in *International Conference on Learning Representations (ICLR)* (2017).
33. Hamilton, W. L., Ying, R. & Leskovec, J. Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems* **2017-Decem,** 1025–1035. ISSN: 10495258. arXiv: 1706.02216 (2017).
34. Data61, C. *StellarGraph Machine Learning Library* https://github.com/stellargraph/stellargraph. 2018.
35. Von Luxburg, U. & Luxburg, U. A tutorial on spectral clustering. *Statistics and Computing* **17,** 395–416. ISSN: 0960-3174. arXiv: 0711.0189. http://link.springer.com/10.1007/s11222-007-9033-z (2007).
36. Kunegis, J. & Lommatzsch, A. Learning spectral graph transformations for link prediction. *Proceedings of the 26th Annual International Conference on Machine Learning - ICML '09,* 1–8. http://portal.acm.org/citation.cfm?doid=1553374.1553447http://doi.acm.org/10.1145/1553374.1553447 (2009).
37. Zhang, Z. *et al. Arbitrary-order proximity preserved network embedding* in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (2018), 2778–2786.
38. Wang, D., Song, C. & Barabási, A. L. Quantifying long-term scientific impact. *Science* **342,** 127–132. ISSN: 10959203 (2013).
39. Bolukbasi, T., Chang, K.-W., Zou, J., Saligrama, V. & Kalai, A. Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings. *30th Conference on Neural Information Processing Systems,* 1–9. ISSN: 10495258. arXiv: 1607.06520. https://code.google.com/archive/p/word2vec/http://arxiv.org/abs/1607.06520 (2016).

40. Ethayarajh, K., Duvenaud, D. & Hirst, G. *Towards Understanding Linear Word Analogies* in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (Association for Computational Linguistics, Florence, Italy, 2019), 3253–3262. https://www.aclweb.org/anthology/P19-1315.

41. Ethayarajh, K., Duvenaud, D. & Hirst, G. *Understanding Undesirable Word Embedding Associations* in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (Association for Computational Linguistics, Florence, Italy, 2019), 1696–1705. https://www.aclweb.org/anthology/P19-1166.

42. Zhou, K., Ethayarajh, K. & Jurafsky, D. Frequency-based Distortions in Contextualized Word Embeddings. arXiv: 2104.08465. http://arxiv.org/abs/2104.08465 (2021).

43. Lambiotte, R & Rosvall, M. Ranking and clustering of nodes in networks with smart teleportation. *Phys. Rev. E* **85,** 56107. https://link.aps.org/doi/10.1103/PhysRevE.85.056107 (2012).

44. Edwards, H. & Storkey, A. Censoring representations with an adversary. *4th International Conference on Learning Representations, ICLR 2016 - Conference Track Proceedings,* 1–14. arXiv: 1511.05897 (2016).

45. Bose, A. J. & Hamilton, W. L. Compositional fairness constraints for graph embeddings. *36th International Conference on Machine Learning, ICML 2019* **2019-June,** 1139–1148. arXiv: 1905.10674 (2019).

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] See Figures 3 and 4.
   (b) Did you describe the limitations of your work? [Yes] See the last paragraph of the Discussion section.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See the Broader Impact section.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Sections 2.2 and 3.1, 3.2, and 3.3.
   (b) Did you include complete proofs of all theoretical results? [Yes] See Sections 2.2, 3.1, 3.2, 3.3 for the main derivation and Sections 1 and 2 in Supplementary Information for the step-by-step derivation.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See Section 3 in Supplementary Information.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 4.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Figures 3 and 4.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Section 3 in Supplementary Information.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] See Table 2.
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We include our code to the Supplementary Information.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] All data except WoS data are publicly available. We discuss how we obtain the consent in Section 3 in Supplementary Information.

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes] Our data did not contain privacy-sensitive information.

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [No] We neither use crowdsourcing nor conducted reserarch with human subjects.

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [No]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [No]