
Privacy Preserving Adjacency Spectral Embedding on Stochastic Blockmodels

Li Chen^{* 1}

Abstract

For graphs generated from stochastic blockmodels, adjacency spectral embedding is asymptotically consistent. Further, adjacency spectral embedding composed with universally consistent classifiers is universally consistent to achieve the Bayes error. However when the graph contains private or sensitive information, treating the data as non-private can potentially leak privacy and incur disclosure risks. In this paper, we propose a differentially private adjacency spectral embedding algorithm for stochastic blockmodels. We demonstrate that our proposed methodology can estimate the latent positions close to, in Frobenius norm, the latent positions by adjacency spectral embedding and achieve comparable accuracy at desired privacy parameters in simulated and real world networks.

1. Introduction

Our framework considers random graphs. A random graph is a graph-valued random variable, with a fixed vertex set and a random edge set generated from a probabilistic distribution. The latent position models (Hoff et al., 2002), random dot product graphs (Young & Scheinerman, 2007) and stochastic blockmodels (Holland et al., 1983b; Wang & Wong, 1987) are a few examples of random graphs. Many algorithms developed on random graphs indicate vertices from the same class share similar connectivity patterns. One important exploitation task on such datasets is vertex classification to identify the class labels of the vertices. For example, we may wish to classify whether a user in a social network holds liberal or conservative political views, or whether an individual in a communication network belongs to a social community.

The adjacency spectral embedding and Laplacian spectral embedding methods are valuable for performing inference on graphs realized from a stochastic blockmodels (Sussman

et al., 2014; Fishkind et al., 2013; Sussman et al., 2012; Tang et al., 2013). Theoretically they are proven to be consistent estimation (Devroye et al., 1996) of the latent positions that generate the stochastic blockmodels. Classification techniques such as k -nearest neighbor composed with spectral embedding are proved to be Bayes optimal, which indicate such composition can theoretically obtain the minimum misclassification error.

On the other hand, graphs that describe social networks, communication networks or purchasing networks can contain sensitive and private information about individuals such as their communication frequencies, patterns, webpage visits, purchasing behaviors and so on. It becomes critical to propose and design inference algorithms on random graphs for not only generating effective and accurate analysis results but also protecting the sensitive nature of the information.

Differential privacy (Dwork, 2011) is a mechanism that measures privacy risk via parameters α or (α, δ) and bounds the likelihood of the algorithm when two data sets differ by one sample. A randomized algorithm is considered differentially private when its outputs cannot expose whether a particular individual's information was used in the computation. Essentially, differential privacy requires the probability distribution on the analysis results, over the outcome space of the underlying algorithm, to be roughly the same, independent of whether or not any individual exists in the data set.

Motivated by the practical challenges faced with inference on random graphs containing sensitive information, in this paper, we propose a differentially private adjacency spectral embedding procedure on stochastic blockmodels to estimate the latent positions and further composite it with universally consistent classifier such as k nearest neighbor for vertex classification. We demonstrate in our experiments the classification efficacy in simulated experiments and two real-world networks.

The rest of the paper is organized as follows. In Section 2, we provide brief background on graph models and differential privacy. In Section 3, we present our proposed algorithm differentially private adjacency spectral embedding. In Section 4, we present results on numerical simulation, a political blog network and a co-purchasing product network.

¹Security and Privacy Research, Intel Labs. Correspondence to: Li Chen <li.chen@intel.com>.

2. Preliminaries

2.1. Graph Models

A random graph is a graph-valued random variable: $\mathcal{G} : \Omega \rightarrow \mathcal{G}_n$, where \mathcal{G}_n represents the collection of all $2^{\binom{n}{2}}$ possible graphs on the vertex set $V = [n]$, and Ω is a probability space. Associated with the adjacency matrix $A \in \mathbb{R}^{n \times n}$, there exists a probability matrix $P \in [0, 1]^{n \times n}$, where each entry P_{ij} denotes the probability of edge between vertex i and vertex j .

In the latent position model (LPM) (Hoff et al., 2002), each vertex i is associated with a latent random variable $X_i \in \mathbb{R}^d$ drawn independently from a specified distribution F on \mathbb{R}^d . These latent variables determine the probabilities of edge existence. The adjacency matrix entries $A_{ij}(X_i, X_j) \sim \text{Bernoulli}(l(X_i, X_j))$ are conditionally independent, where $l : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, 1]$ is defined as the link function.

The stochastic blockmodel (SBM) is a special case of LPM. The latent positions of an SBM are sampled as a mixture of the point masses which are the eigenvectors of B . The SBM (Holland et al., 1983a) is a family of random graph models with a set of n vertices randomly belonging to K blocks. Conditioned on the K -partition, edges between all the pairs of vertices are independent Bernoulli trials with parameters determined by the block memberships of the two vertices.

Definition 1. Stochastic Blockmodel SBM($[n], B, \pi$) Let K be the number of blocks. Let π be a length K vector in the unit simplex Δ^{K-1} specifying the block membership probabilities. The block membership of the vertex i is given by $Y_i \stackrel{iid}{\sim} \text{Multinomial}([K], \pi)$. Let B be a $K \times K$ symmetric block communication probability matrix. Then the graph G is realized from a stochastic blockmodel $G \sim \text{SBM}([n], B, \pi)$ if

$$\mathbb{P}(A|Y_1, \dots, Y_n) = \prod_{i < j} P_{ij}^{A_{ij}} (1 - P_{ij})^{1 - A_{ij}}, \quad (1)$$

$$P_{ij} = \mathbb{P}(A_{ij} = 1 | X_i, X_j) = \mathbb{P}(A_{ij} = 1 | Y_i, Y_j) = B_{Y_i, Y_j}. \quad (2)$$

2.2. Adjacency Spectral Embedding

Adjacency spectral embedding (ASE) (Sussman et al., 2012) has theoretical guarantee to consistently estimate the latent positions of SBM. This technique applies spectral decomposition on A to compute the first d eigen-pairs $(U_A, S_A) \in \mathbb{R}^{n \times d} \times \mathbb{R}^d$, where S_A is diagonal with the top d eigenvalues sorted in absolute values. The resulted embedding $ASE := U_A S_A^{\frac{1}{2}}$ consistently estimates the block memberships of SBM. The consistency property of Laplacian spectral embedding is proved in (Rohe et al., 2011).

Adjacency spectral embedding followed by k -nearest neighbor (k NN) classifiers (Sussman et al., 2014), linear classifiers (Tang et al., 2013) is shown to be universally consistent,

which means such procedure achieves the minimum Bayes error as n goes to infinity.

2.3. Differential Privacy

Denote a randomized algorithm by f taking values in a set S . A and B are datasets that differ by one element.

Definition 2. α -Differential Privacy f provides α -differential privacy if

$$\mathbb{P}(f(A) \in S) \leq \mathbb{P}(f(B) \in S) e^\alpha. \quad (3)$$

Definition 3. (α, δ) -Differential Privacy f provides (α, δ) -differential privacy if

$$\mathbb{P}(f(A) \in S) \leq \mathbb{P}(f(B) \in S) e^\alpha + \delta. \quad (4)$$

The parameter α bounds on the change in probability of any outcome. A low value of α such as 0.1 implies that very little can change in the beliefs about any sample's existence in the database. The (α, δ) -Differential Privacy is a weaker notion compared with the α -Differential Privacy.

3. Differentially Private Adjacency Spectral Embedding

We propose a differentially private mechanism for adjacency spectral embedding (DP-ASE) to estimate the latent positions of stochastic blockmodels (SBM).

Algorithm 1 DP-ASE for SBM

Input: The adjacency matrix A . The embedding dimension d . Desired privacy parameters (α, δ) .

Output: Differentially private approximations of latent positions $\mathcal{X}_{DP} \in \mathbb{R}^{n \times d}$.

Step 1: Set $\beta^2 := \frac{8d^2 \log^2(d/\delta)}{n^2 \alpha^2}$. Sample a matrix E whose entries are i.i.d from normal distribution $\mathcal{N}(0, \beta^2)$. Symmetrize E .

Step 2: Apply spectral embedding on $A_{DP} := A + E$. Compute the first d eigen-pairs of A_{DP} , denoted by $(U_{DP}, S_{DP}) \in \mathbb{R}^{n \times d} \times \mathbb{R}^d$, where S_{DP} has d largest eigenvalues in magnitude sorted in non-increasing order.

Step 3: Denote the d -dimensional coordinate-scaled singular vector matrix of A_{DP} to be $\mathcal{X}_{DP} := U_{DP} S_{DP}^{1/2} \in \mathbb{R}^{n \times d}$.

Theorem 1 (Privacy of DP-ASE). *Algorithm 1 computes an (α, δ) -differentially private approximation to adjacency spectral embedding on stochastic blockmodels.*

Proof Sketch. The proof of Theorem 1 follows directly from the differentially private algorithm SULQ to approximate PCA with privacy guarantee (Blum et al., 2005; Chaudhuri et al., 2012). \square

4. Experiments

An effective differentially private algorithm should not only provide privacy guarantee but also generate reasonably accurate results. In our experiments, we compose a universally consistent classifier k nearest neighbor with DP-ASE, denoted by $k\text{NN}\circ\text{DP-ASE}$ and evaluate its classification error. The exploitation task is to correctly predict the block memberships of the vertices. The comparison baseline is the classification error by $k\text{NN}\circ\text{ASE}$. We set $k = 3$ in the nearest neighbor classifier, use leave-one-out cross validation to estimate the classification error and generate new matrix $E \sim \mathcal{N}(0, \beta^2)$ in each Monte Carlo replicate.

4.1. Simulation

We simulate SBM with $K = 2$ blocks ($Y \in \{1, 2\}$) and parameters $B = \begin{pmatrix} 0.3 & 0.1 \\ 0.1 & 0.2 \end{pmatrix}$, $\pi = [0.4, 0.6]^T$.

Our first experiment is to compare the latent position approximation and classification performance with or without preserving privacy at fixed but low privacy budget, as the number of vertices increases. We apply ASE and DP-ASE at fixed privacy parameters $\alpha = 0.1$ and $\delta = 0.001$. We first vary the number of vertices $n \in \{50, 100, 500, 1000, 1500, 3000, 3100, 3500, 4000\}$ and compare two terms: 1. the Frobenius norm (F-norm) of the latent positions estimated by DP-ASE and ASE respectively, and the F-norm of the latent positions estimated by two ASEs aligned via solving the orthogonal Procrustes fit problem (Gower et al., 2004); 2. the classification errors by $k\text{NN}\circ\text{DP-ASE}$ and $k\text{NN}\circ\text{ASE}$ respectively at each n . As n increases, DP-ASE approximates to ASE in both latent position estimation and classification error, as seen in left two figures in Fig 1.

Next we understand the utility-privacy tradeoff for smaller-sized stochastic blockmodels. We set $n = 300$, vary the privacy parameters $\alpha \in \{0.001, 0.011, 0.021, \dots, 0.05\}$ and $\delta \in \{0.0001, 0.0021, 0.0041, \dots, 0.6\}$ and examine the classification error of $k\text{NN}\circ\text{DP-ASE}$. We inspect the minimum levels of privacy parameters for $k\text{NN}\circ\text{DP-ASE}$ to achieve the Bayes error, which is zero in this simulation. In the heat map of Fig 1, the rows are α and columns are δ . $k\text{NN}\circ\text{DP-ASE}$ achieves the Bayes error at the lower values of the privacy parameters, which provide tighter privacy bounds.

4.2. Real Data Experiments

4.2.1. POLITICAL BLOG SPHERE

The political blog sphere (Adamic & Glance, 2005) contains 1490 blogs during the 2004 presidential election as vertices, and edges if the blogs are linked. Each blog is either liberal or conservative with label distribution at (0.508,

0.492). Hence the chance error is 0.492. Since the number of memberships does not necessarily equal to the correct embedding dimension in practice, we first vary the embedding dimension $d \in \{2, 5, 8, \dots, 100\}$ at fixed $\alpha = 0.1$ and $\delta = 0.01$. As seen in Figure 2 (a), the lowest classification error by $k\text{NN}\circ\text{DP-ASE}$ is 0.25 at $\hat{d} = 2$, but it is higher than the classification error by $k\text{NN}\circ\text{ASE}$ at 0.180. Such degradation is due to the noise added under differential privacy scheme.

Our next experiment is to examine the trade-off between privacy budget and utility by observing, at which values of α , the error rate under privacy scheme is as accurate as the error rate without privacy. In practice, there exists limited guidance on appropriate values of α (Jayaraman & Evans, 2019) and thus we intend to understand such tradeoff via empirical evaluation. We fix $\delta = 0.01$, vary the privacy parameter $\alpha = 0.001 : 0.01 : 10$ and present the classification errors of $k\text{NN}\circ\text{DP-ASE}$ and $k\text{NN}\circ\text{ASE}$ at $\hat{d} = 2$ in Fig 2 (b). As α increases, the error rate by $k\text{NN}\circ\text{DP-ASE}$ decreases and approximates the error rate by $k\text{NN}\circ\text{ASE}$. For $\hat{d} = 2$, the lowest classification error by $k\text{NN}\circ\text{DP-ASE}$ is 0.164 at best privacy parameter $\alpha = 6.451$. However we discover that at a privacy budget α as low as 0.251, $k\text{NN}\circ\text{DP-ASE}$ has error of 0.189 very close to the classification error without privacy at 0.180.

4.2.2. PRODUCT CO-PURCHASING NETWORK

The Amazon product co-purchasing network (Yang & Leskovec, 2015) contains products as nodes with edge existence such that if a product i is frequently co-purchased with product j , then an undirected edge exists between i and j . The labels are product categories provided by Amazon. The original network contains 0.33M vertices with 0.92M edges. We select the six largest categories out of 5000 categories from the network. Since some products belong to multiple categories, we choose only one category to associate with each product and generate a subgraph of 794 nodes with label distribution of (0.413, 0.393, 0.194). The classification chance error is 0.587. Since the number of memberships does not necessarily equal to the correct embedding dimension in practice, we first vary the embedding dimension $d \in \{2, 3, \dots, 60\}$ at fixed $\alpha = 0.1$ and $\delta = 0.01$ and use 5-nearest neighbor classifier. We find the best classification error under differential privacy scheme at 0.584 is only slightly lower than chance error. The contribution of noise under differential privacy is likely to cause the degradation in classification performance of $k\text{NN}\circ\text{DP-ASE}$. The classification error without differential privacy is lowest at 0.489. We expect to see an improvement in classification efficacy, if better disjoint categories are selected to construct the subgraph or multi-label classification is performed. From the minimum error rates by $k\text{NN}\circ\text{DP-ASE}$ and $k\text{NN}\circ\text{ASE}$ respectively, we select two best embedding dimensions to be




Figure 1. (1) F -norm of the estimated latent positions. As n increases, the distance of the estimated latent positions via ASE and DP-ASE become closer. (2) Classification error. When n is small, $k\text{NN}\circ\text{DP-ASE}$ has higher classification error than $k\text{NN}\circ\text{ASE}$. As n increases, classification error by $k\text{NN}\circ$ goes to the Bayes error at zero. (3) 3D plot of F -norm difference as α and δ vary. (4) Heat map of classification error as α and δ vary. $k\text{NN}\circ\text{DP-ASE}$ achieves Bayes error at privacy parameters with lower values.




Figure 2. Experiment on political blog sphere. (a) Classification error by $k\text{NN}\circ\text{DP-ASE}$ and $k\text{NN}\circ\text{ASE}$ at embedding dimensions $d \in \{2, 5, 8, \dots, 100\}$ with fixed $\alpha = 0.1$ and $\delta = 0.01$. (b) Privacy-utility tradeoff analysis. Utility-privacy tradeoff evaluation. As α increases, the error rate by $k\text{NN}\circ\text{DP-ASE}$ decreases and approximates the error rate by $k\text{NN}\circ\text{ASE}$. For $\hat{d} = 2$, the lowest classification error by $k\text{NN}\circ\text{DP-ASE}$ is 0.164 at best privacy parameter $\alpha = 6.451$. However we discover that at a privacy budget α as low as 0.251, $k\text{NN}\circ\text{DP-ASE}$ has error of 0.189 very close to the classification error without privacy at 0.180.

$$\hat{d}_1 = 35 \text{ and } \hat{d}_2 = 60.$$

Our next experiment is to examine the trade-off between privacy budget and utility by observing, at which values of α , the error rate under privacy scheme is as accurate as the error rate without privacy. In practice, there exists limited guidance on appropriate values of α (Jayaraman & Evans, 2019) and thus we intend to understand such tradeoff via empirical evaluation. We fix $\delta = 0.01$, vary the privacy parameter $\alpha = 0.001 : 0.01 : 10$ and present the classification errors of $k\text{NN}\circ\text{DP-ASE}$ and $k\text{NN}\circ\text{ASE}$ at $\hat{d} \in \{\hat{d}_1, \hat{d}_2\}$ respectively in Fig 3 (b) and (c). As α increases, the error rate by $k\text{NN}\circ\text{DP-ASE}$ decreases and approximates the error rate by $k\text{NN}\circ\text{ASE}$. For $\hat{d} = 35$, the lowest classification error by $k\text{NN}\circ\text{DP-ASE}$ is 0.52 at best privacy parameter $\alpha = 3.42$. For $\hat{d} = 60$, the lowest classification error by $k\text{NN}\circ\text{DP-ASE}$ is 0.50 at best privacy parameter $\alpha = 9.25$.

Figure 3. Experiment on product co-purchasing network. Utility-privacy tradeoff evaluation. As α increases, the error rate by $k\text{NN}\circ\text{DP-ASE}$ decreases and approximates the error rate by $k\text{NN}\circ\text{ASE}$. For $\hat{d} = 35$, the lowest classification error by $k\text{NN}\circ\text{DP-ASE}$ is 0.52 at best privacy parameter $\alpha = 3.42$. For $\hat{d} = 60$, the lowest classification error by $k\text{NN}\circ\text{DP-ASE}$ is 0.50 at best privacy parameter $\alpha = 9.25$.

5. Conclusion

In this paper we design a differentially private approximation to adjacency spectral embedding (DP-ASE) on stochastic blockmodels and investigate the theoretical and empirical performance of the proposed algorithm. In our numerical experiments, we demonstrate that as n increases, for simulated stochastic blockmodels, a universally consistent classifier composed with DP-ASE can achieve the best error as ASE. From varying the privacy parameters in the simulation experiments, we observe that such mechanism can achieve the minimum possible error at desired privacy parameter levels, which can provide a tighter privacy bound. In practice, there has not been a principled guidance on the best privacy parameters to select while preserving inference accuracy. In our empirical analysis, we find that at reasonable privacy budget DP-ASE can achieve classification error as close as without preserving privacy. We continue our research with theoretical proofs on the consistency of DP-ASE and the universal consistency of a universally consistent classifier composed with DP-ASE.

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*, pp. 36–43. ACM, 2005.
- Blum, A., Dwork, C., McSherry, F., and Nissim, K. Practical privacy: the sulq framework. In *Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pp. 128–138. ACM, 2005.
- Chaudhuri, K., Sarwate, A., and Sinha, K. Near-optimal differentially private principal components. In *Advances in Neural Information Processing Systems*, pp. 989–997, 2012.
- Devroye, L., Györfi, L., and Lugosi, G. *A probabilistic theory of pattern recognition*, volume 31. New York: Springer, 1996.
- Dwork, C. Differential privacy. *Encyclopedia of Cryptography and Security*, pp. 338–340, 2011.
- Fishkind, D. E., Sussman, D. L., Tang, M., Vogelstein, J. T., and Priebe, C. E. Consistent adjacency-spectral partitioning for the stochastic block model when the model parameters are unknown. *SIAM Journal on Matrix Analysis and Applications*, 34(1):23–39, 2013.
- Gower, J. C., Dijkstra, G. B., et al. *Procrustes problems*, volume 30. Oxford University Press on Demand, 2004.
- Hoff, P. D., Raftery, A. E., and Handcock, M. S. Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97(460):1090–1098, 2002.
- Holland, P. W., Laskey, K. B., and Leinhardt, S. Stochastic blockmodels: First steps. *Social networks*, 5(2):109–137, 1983a.
- Holland, P. W., Laskey, K. B., and Leinhardt, S. Stochastic blockmodels: First steps. *Social networks*, 5(2):109–137, 1983b.
- Jayaraman, B. and Evans, D. When relaxations go bad:” differentially-private” machine learning. *arXiv preprint arXiv:1902.08874*, 2019.
- Rohe, K., Chatterjee, S., and Yu, B. Spectral clustering and the high-dimensional stochastic blockmodel. *The Annals of Statistics*, 39(4):1878–1915, 2011.
- Sussman, D. L., Tang, M., Fishkind, D. E., and Priebe, C. E. A consistent adjacency spectral embedding for stochastic blockmodel graphs. *Journal of the American Statistical Association*, 107(499):1119–1128, 2012.
- Sussman, D. L., Tang, M., and Priebe, C. E. Consistent latent position estimation and vertex classification for random dot product graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(1):48–57, 2014.
- Tang, M., Sussman, D. L., and Priebe, C. E. Universally consistent vertex classification for latent positions graphs. *The Annals of Statistics*, 41(3):1406–1430, 2013.
- Wang, Y. J. and Wong, G. Y. Stochastic blockmodels for directed graphs. *Journal of the American Statistical Association*, 82(397):8–19, 1987.
- Yang, J. and Leskovec, J. Defining and evaluating network communities based on ground-truth. *Knowledge and Information Systems*, 42(1):181–213, 2015.
- Young, S. J. and Scheinerman, E. R. Random dot product graph models for social networks. In *Algorithms and models for the web-graph*, pp. 138–149. Springer, 2007.