A Differentially Private Guide for Graph Analytics

Felipe T. Brito
Universidade Federal do Ceará
Fortaleza, CE, Brazil
felipe.timbo@lsbd.ufc.br

André L. C. Mendonça
Universidade Federal do Ceará
Fortaleza, CE, Brazil
andre.luis@lsbd.ufc.br

Javam C. Machado
Universidade Federal do Ceará
Fortaleza, CE, Brazil
javam.machado@lsbd.ufc.br

ABSTRACT

In recent years, the application of graph analytics has been shown to provide huge value in many application domains. However, the growing adoption of graph analytics corresponds to an increasing need to protect sensitive information in graph data. In this context, differential privacy has become the de facto standard for privacy-preserving data analysis under strong mathematical guarantees. This tutorial provides a comprehensive overview of differentially private methods and techniques to protect sensitive information while conducting meaningful graph analysis. We explore a variety of definitions, mechanisms, examples, and case studies to demonstrate the application of these methods in various scenarios.

1 MOTIVATION, GOALS, AND OBJECTIVES

Graph analytics refers to techniques and tools used to examine, interpret, and extract insights from complex relationships, structures, and interconnections in graph data structures [38]. Such analyses have applications in various disciplines, including physics, mathematics, computer science, biology, sociology, and economics. Because graphs usually contain sensitive information, releasing this data for analysis without sufficient privacy guarantees may seriously jeopardize the individuals' privacy. Current laws and regulations on data privacy (see, e.g., the *General Data Protection Regulation* [48] in E.U. and the *Federal Communications Commission* regulations [14] in the U.S. around customer proprietary network information) require that individuals are no longer re-identifiable from released information.

Differential privacy (DP) [10] has emerged as the de facto standard notion of privacy for data release. It offers a formal definition of privacy with interesting properties, such as no computational/informational assumptions about attackers, data type-agnosticism, and composability [33]. The main idea behind differential privacy is that a given query is answered by a *randomized algorithm* that queries private information and returns a randomized answer sampled from an *output distribution*. A randomized algorithm is also referred to as a mechanism. A mechanism is differentially private if the probability distribution of the outputs does not change significantly based on the presence or absence of an individual.

In this tutorial, we will explore a set of differentially private methods and techniques applicable to graph analytics, aiming to protect sensitive information and enable meaningful analysis in graph-based data structures. The materials introduced in this tutorial will provide a comprehensive guide to understanding which techniques are more suitable for different graph analyses. We will also aim to identify relevant open problems and research directions for the community in this field of study.

2 TUTORIAL OUTLINE

This tutorial will be presented in 1.5 hours. It will consist of 5 main modules. The first module will focus on the foundations of graph analytics and will last 10 minutes. The next two modules, "Introduction to Differential Privacy" and "Differential Privacy for Graphs", will focus on designing differentially private strategies to preserve the privacy of graph data, each lasting 20 minutes. The subsequent module will focus on applications and use cases of differential privacy for graph analytics and is scheduled for 30 minutes. Finally, the concluding module will focus on identifying open problems and future directions and will last 10 minutes.

2.1 Fundamentals of Graph Analytics

This module presents an overview of graph data structures and common analysis. In contrast to traditional tabular data analytics, graph analytics focuses on analyzing the structure, relationships between vertices and edges, and the flows composing graph data. Common statistics of interest for graph data include node degrees, associated degree distribution, centrality metrics, and other pertinent measures. Additionally, subgraph counts and various distance metrics are also notable cases frequently examined in graph analysis. We will present these metrics and illustrate inherent privacy concerns in graph analytics, motivating the need for formal guarantees to provide sufficient protection against privacy violations.

2.2 Introduction to Differential Privacy

In this module, we introduce the main concepts of differential privacy, highlighting its key principles and core components.

2.2.1 *Intuitions and definition.* In the original definition of differential privacy, private data is viewed as a collection of records, with each record corresponding to an individual. In essence, differential privacy promises privacy protection by injecting noise into these records, i.e., modifying the original data by introducing randomness [11]. It is important to mention that DP is not a single tool but rather a paradigm that quantifies and manages privacy violation risks. Then, differential privacy can be adopted from simple statistical estimations to machine learning [53]. This section covers the DP definition and the notion of neighboring datasets using examples.

2.2.2 *Privacy budget and sensitivity.* Differential privacy allows a minor change between the output of the real analysis and the output generated when a single register is added or removed from a given dataset. This minor change is controlled by the parameter ϵ, also known as privacy budget. We discuss in this section that it is not simple to define an adequate ϵ for an application [47]. Additionally, we introduce the notion of sensitivity [11], which measures the maximum impact on data analysis by adding or deleting any data record in the dataset. We also cover the smooth sensitivity framework for answering high-sensitivity queries [39]. Attendees will learn how to compute the sensitivity of analysis and how to establish bounds on privacy loss.

Series ISSN: 2367-2005

10.48786/edbt.2024.86

2.2.3 Mechanisms and properties. Mechanisms are ways of achieving differential privacy. In this section, we cover classic differentially private mechanisms, including randomized response [11], Laplace mechanism [11], geometric mechanism [15], and exponential mechanism [32]. In addition, composition theorems, including sequential composition, parallel composition, and post-processing, are all covered to conduct multiple DP analyses.

2.2.4 Local Differential Privacy. The basic (global) setup of differential privacy involves a trusted curator (third party) that has access to the original data and globally adds noise to achieve differential privacy. However, finding a genuinely trusted third party for data collection and processing can be challenging in some practical scenarios, restricting the applicability of global differential privacy approaches. To address this concern, we present local differential privacy (LDP) as an alternative approach that eliminates the need for a trusted data curator [9].

2.3 Differential Privacy for Graphs

The fundamental concept of differential privacy relies on the definition of neighboring datasets. In previous definitions, a neighboring dataset is defined as a dataset obtained by adding or removing a single record. In the context of graphs, which primarily focus on the relationships between individuals, the association between private data and dataset records becomes less clear. This module gives an overview of extended differentially private definitions for graphs [3, 16, 23, 26, 44] as well as DP mechanisms [4, 13, 19] that can be adopted to provide data privacy in graph analytics.

2.3.1 Node DP. This section explores the node-level differential privacy definition [26], which allows limited inference about the existence or absence of a node (entity) in a graph. Node differential privacy aims to protect both the nodes and their adjacent edges. Node DP is much harder to achieve than other differentially private definitions for graphs since it may be infeasible to design node-differentially private algorithms that provide accurate graph analysis.

2.3.2 Edge DP. This section covers another important extended definition of DP for graphs, denoted edge differential privacy [16]. Edge DP was the first adaptation of DP to deal with graphs. It is a method that hides the presence or absence of a single edge in a graph. In certain applications, this level of privacy assurance may be reasonable. However, there are scenarios where it becomes desirable to extend privacy protection beyond individual edges.

2.3.3 Edge-weight DP. The two main alternatives for applying differential privacy on graphs (edge and node DP) are not well suited to weighted graphs. In general, it is impossible to release, e.g., shortest paths, with meaningful utility under edge-DP or node-DP since changing a single edge can significantly change the distances in a graph. Therefore, this section discusses two new notions of differential privacy for weighted graphs. In the first one [44], two graphs are said to be neighbors if they have the same topology and similar weight functions. In contrast to the previous formulation, the authors of this tutorial proposed a new definition for neighboring weighted graphs that consider both the graph topology and the edge weights as private information [2, 3].

2.3.4 Attributed graphs. Another case where neither node-DP nor edge-DP privacy models are adequate to provide desirable graph utility is when graphs have attributes attached to their edges or nodes. Initially, this section covers a proposed new neighboring definition to deal with attributed graphs under DP guarantees, denoted edge-adjacent attributed graphs [23]. Another definition related to attributed graphs will also be presented [30], which considers two attributed graphs to be neighboring if one can be obtained from another by adding/removing one certain attribute along with all related edges.

2.3.5 DP mechanisms for graphs. In this section, we briefly give an overview of mechanisms that can be applied in graph analytics, such as recursive mechanism [4] for subgraph counting, XOR mechanism [19] for general social network analysis, and local dampening mechanism [13] for influential node analysis. The authors of this tutorial proposed the latter.

2.4 Differentially Private Graph Analysis

This module starts with an overview of the research in private graph analysis [22, 29]. We discuss how differentially private techniques are applied to analyze graphs, providing significant utility and guaranteeing data privacy. Details about the proposed methods are presented according to the following categorization of graph analysis and applications: (1) entire graph release and random graphs; (2) degree sequence and subgraph counting; (3) centrality and community detection; (4) shortest paths and distances; and finally (5) graph neural networks (GNNs).

2.4.1 Entire graph release and random graphs. Differentially private release of the entire graph has been studied extensively in recent years. The main advantage of these approaches is that they are agnostic to the analysis in the sense that one can compute any statistics on the released graph. In this scenario, Pygmalion [43] aimed to release the graph topology under edge-DP by extracting a graph's detailed structure into private dK-graph [31] and then generating a synthetic graph. Wang and Wu [49] subsequently proposed an improvement in the utility of the dK-graph model by calibrating noise based on the smooth sensitivity. A different approach [51] that adopts the Hierarchical Random Graph (HRG) model [5] was introduced to release network data. The authors observed that the noise scale enforced by DP could be reduced by estimating the connection probabilities between nodes. Top-m filter (TmF) [36] was proposed to remedy scalability problems in previous work. It adds Laplace noise to each cell in the adjacency matrix and uses an idea similar to High-pass Filter [6] to avoid the materialization of the noisy adjacency matrix. Recently, Iftikhar et al. [18] developed a micro aggregation-based framework for graph anonymization, which perturbs graphs by adding noise to the distributions of the original graphs. PBCN [17] was also developed to release noisy graphs under edge-DP. Due to the difficulty in obtaining high-utility private mechanisms when releasing the entire graph, there is just one recent work [21] related to node differential privacy when compared to edge-DP. The tutorial covers the main approaches to releasing and generating random graphs.

2.4.2 Degree sequence and subgraph counting. Another widely studied graph statistic is the degree sequence of a graph. Hay et al. [16] proposed a constraint inference-based technique to release degree sequences via DP mechanisms. Karwa and Slavković [25] introduced an optimization step after constraint inference. Subgraph counting queries count the number of times a certain subgraph appears in a given graph. Common subgraphs include triangles and stars. In this context, Karwa et al. [24] extended the results of smooth sensitivity to privately release k-stars and k-triangles. In addition, the ladder function [52] was used to achieve high accuracy with efficient time complexities. This approach effectively combines the concept of local sensitivity at a distance t, from the smooth sensitivity framework [39] with the exponential mechanism, also allowing the counting of k-cliques in a graph. Sun et al. [45] presented a technique to privately release some graph statistics, such as triangles, three-hop paths, and k-clique counts under LDP. We highlight the main techniques in this tutorial.

2.4.3 *Centrality and community detection.*

In social network analysis, centrality measures and community detection are essential for understanding complex topologies and relationships between individuals [27, 46]. In this context, Nguyen et al. [37] adopted the Louvain method as the back-end community detection for input perturbation schemes and proposed the LouvainDP method. Ji et al. [20] proposed an algorithm to protect the privacy of both network topology and node attributes in community detection in social networks. Additionally, the XOR mechanism [19] presents promising results in detecting communities given network topology, i.e., the released adjacency matrix. Betweenness centrality of a node when relevant edge information is spread is privately addressed by Roohi et al. [41]. It is essential to mention that some of our works [3, 13] are also suitable for identifying influential nodes in both weighted and non-weighted networks. The tutorial provides a comprehensive overview of these techniques.

2.4.4 *Shortest paths and distances.*

Sealfon [44] aimed to release weighted shortest paths between pairs of nodes and approximate distances between all pairs of nodes without revealing sensitive information. Fan and Li [12] revisited the problem of privately releasing approximate distances between all pairs of nodes and recently improved Sealfon's results. Li et al. [28] presented a merging barrels and consistency inference (MBCI) approach to releasing weighted graphs under DP guarantees. Recently, our work [2, 3] proposed both global and local DP approaches to release weighted graphs, also privately analyzing the weighted shortest paths and distances. In the tutorial, we present those methods.

2.4.5 *Graph neural networks.*

Recently, graph neural networks (GNNs) have gained significant attention due to their ability to capture complex relationships and dependencies between nodes in a graph [50]. In this context, some differentially private techniques have been proposed. Olatunji et al. [40] introduced PrivGnn, a framework that protects sensitive data while releasing the trained GNN model. Daigavane et al. [8] proposed a graph neighborhood sampling scheme while preserving node-level privacy. The works mentioned above enforce privacy only during training and/or model release. This fact puts user information at risk if the data curator is malicious. Contrarily, Bhaila et al. [1] proposed RGNN, a reconstruction-based GNN learning framework that can guarantee node privacy while adopting local DP. The tutorial provides an overview of these and other techniques [7, 35, 42].

2.5 Open Problems and Research Directions

Despite significant improvements in addressing the privacy-preserving problem while analyzing graph data, several issues still need to be tackled. This module explores these remaining issues, providing strategies and potential solutions to advance privacy-preserving techniques in graph analysis further.

3 INTENDED AUDIENCE

This tutorial is aimed at individuals with a foundational understanding of probability (including distributions, means, and variances), data analysis, and graphs/networks. The target audience includes data scientists and analysts, researchers, machine learning practitioners, privacy and security professionals, software engineers, and application developers who deal with sensitive data. The tutorial does not assume prior knowledge of differential privacy.

4 EARLIER VERSION OF THE TUTORIAL

A version of this tutorial was presented at the Brazilian Symposium on Databases - SBBD 2023 [34] and focused on a general review of node-DP and edge-DP models applied only to social network analysis. The current EDBT tutorial proposal provides a comprehensive and deep review of the mentioned models for graphs in general and other neighborhood definitions to offer privacy protections for more graph analytics. We also cover applications of differential privacy in machine learning and data mining. In addition, recent open problems and research directions are presented in this tutorial.

5 BIOGRAPHY

Felipe T. Brito has a Ph.D. in Computer Science from Universidade Federal do Ceará (UFC), Brazil, with a sandwich year at AT&T Labs, New York, USA. He obtained an M.Sc. and a B.Sc. in Computer Science, also from UFC. He is an applied research scientist at the Laboratório de Sistemas e Bancos de Dados (LSBD/UFC), currently working with differentially private mechanisms for weighted graph data structures. His topics of interest include data privacy, machine learning, and data science.

André L. C. Mendonça is a Ph.D. candidate at the Universidade Federal do Ceará (UFC), with a 3-months sandwich period at Laboratoire d'Informatique de Grenoble (LIGLab), Grenoble, France. He obtained an M.Sc. and B.Sc. in Computer Science also from UFC. He is a research assistant at the Laboratório de Sistemas e Bancos de Dados (LSBD/UFC), currently working with differential privacy mechanisms for attributed graphs. His topics of interest include data privacy and social networks.

Javam C. Machado is a full professor at the Computer Science Department of the Universidade Federal do Ceará (UFC), Brazil, and an SBC and ACM member. He obtained a Ph.D. degree in Computer Science from the Université de Grenoble, France. He served as the coordinator of the SBC Special Database Commission (2017) and was a visiting researcher at TelecomSudParis – FR (2001) and AT&T Labs-Research – USA (2018; 2020; 2023). Javam has published more than 170 scientific papers and has advised 40 M.Sc. and 5 Ph.D. students. His topics of interest include data privacy and algorithmic discrimination.

ACKNOWLEDGMENTS

This research was partially funded by Lenovo, as part of its R&D investment under Brazilian Informatics Law. Additional funding was provided by CAPES/Brazil under grant number 189723/2018-01, by CNPq/Brazil under grant number 316729/2021-3, and also by LSBD/UFC.

REFERENCES

[1] Karuna Bhaila, Wen Huang, Yongkai Wu, and Xintao Wu. 2023. Local Differential Privacy in Graph Neural Networks: a Reconstruction Approach. *arXiv preprint arXiv:2309.08569* (2023).

[2] Felipe T Brito. 2023. *Differentially private release of count-weighted graphs*. Ph.D. Dissertation. Universidade Federal do Ceara.

[3] Felipe T. Brito, Victor A. E. de Farias, Cheryl J. Flynn, Subhabrata Majumdar, Javam C. Machado, and Divesh Srivastava. 2023. Global and Local Differentially Private Release of Count-Weighted Graphs. *Proc. ACM Manag. Data* 1, 2 (2023), 154:1–154:25.

[4] Shixi Chen and Shuigeng Zhou. 2013. Recursive mechanism: towards node differential privacy and unrestricted joins. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*. 653–664.

[5] Aaron Clauset, Cristopher Moore, and Mark EJ Newman. 2006. Structural inference of hierarchies in networks. In *ICML Workshop on Statistical Network Analysis*. Springer, 1–13.

[6] Graham Cormode, Cecilia Procopiuc, Divesh Srivastava, and Thanh TL Tran. 2012. Differentially private summaries for sparse data. In *Proceedings of the 15th International Conference on Database Theory*. 299–311.

[7] Enyan Dai, Tianxiang Zhao, Huaisheng Zhu, Junjie Xu, Zhimeng Guo, Hui Liu, Jiliang Tang, and Suhang Wang. 2022. A comprehensive survey on trustworthy graph neural networks: Privacy, robustness, fairness, and explainability. *arXiv preprint arXiv:2204.08570* (2022).

[8] Ameya Daigavane, Gagan Madan, Aditya Sinha, Abhradeep Guha Thakurta, Gaurav Aggarwal, and Prateek Jain. 2021. Node-level differentially private graph neural networks. *arXiv preprint arXiv:2111.15521* (2021).

[9] John C Duchi, Michael I Jordan, and Martin J Wainwright. 2013. Local privacy and statistical minimax rates. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. IEEE, 429–438.

[10] Cynthia Dwork. 2006. Differential privacy. In *International Colloquium on Automata, Languages, and Programming*. Springer, 1–12.

[11] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*. Springer, 265–284.

[12] Chenglin Fan and Ping Li. 2022. Distances release with differential privacy in tree and grid graph. In *2022 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2190–2195.

[13] Victor AE Farias, Felipe T Brito, Cheryl Flynn, Javam C Machado, Subhabrata Majumdar, and Divesh Srivastava. 2023. Local dampening: Differential privacy for non-numeric queries via local sensitivity. *The VLDB Journal* (2023), 1–24.

[14] Federal Communications Commission. 2018. Customer privacy. https://www.fcc.gov/general/customer-privacy. Online; accessed 13 October 2022.

[15] Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. 2012. Universally utility-maximizing privacy mechanisms. *SIAM J. Comput.* 41, 6 (2012), 1673–1693.

[16] Michael Hay, Chao Li, Gerome Miklau, and David Jensen. 2009. Accurate estimation of the degree distribution of private networks. In *2009 Ninth IEEE International Conference on Data Mining*. IEEE, 169–178.

[17] Haiping Huang, Dongjun Zhang, Fu Xiao, Kai Wang, Jiateng Gu, and Ruchuan Wang. 2020. Privacy-preserving approach PBCN in social network with differential privacy. *IEEE Transactions on Network and Service Management* 17, 2 (2020), 931–945.

[18] Masooma Iftikhar, Qing Wang, and Yu Lin. 2020. dk-microaggregation: Anonymizing graphs with differential privacy guarantees. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 191–203.

[19] Tianxi Ji, Pan Li, Emre Yilmaz, Erman Ayday, Yanfang Ye, and Jinyuan Sun. 2021. Differentially private binary-and matrix-valued data query: an XOR mechanism. *Proceedings of the VLDB Endowment* 14, 5 (2021), 849–862.

[20] Tianxi Ji, Changqing Luo, Yifan Guo, Jinlong Ji, Weixian Liao, and Pan Li. 2019. Differentially private community detection in attributed social networks. In *Asian Conference on Machine Learning*. PMLR, 16–31.

[21] Xun Jian, Yue Wang, and Lei Chen. 2021. Publishing graphs under node differential privacy. *IEEE Transactions on Knowledge and Data Engineering* (2021).

[22] Honglu Jiang, Jian Pei, Dongxiao Yu, Jiguo Yu, Bei Gong, and Xiuzhen Cheng. 2021. Applications of differential privacy in social network analysis: A survey. *IEEE Transactions on Knowledge and Data Engineering* 35, 1 (2021), 108–127.

[23] Zach Jorgensen, Ting Yu, and Graham Cormode. 2016. Publishing attributed social graphs with formal privacy guarantees. In *Proceedings of the 2016 international conference on management of data*. 107–122.

[24] Vishesh Karwa, Sofya Raskhodnikova, Adam Smith, and Grigory Yaroslavtsev. 2011. Private analysis of graph structure. *PVLDB* 4, 11 (2011), 1146–1157.

[25] Vishesh Karwa and Aleksandra B Slavković. 2012. Differentially private graphical degree sequences and synthetic graphs. In *International Conference on Privacy in Statistical Databases*. Springer, 273–285.

[26] Shiva Prasad Kasiviswanathan, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. 2013. Analyzing graphs with node differential privacy. In *Theory of Cryptography Conference*. Springer, 457–476.

[27] Jesse Laeuchli, Yunior Ramírez-Cruz, and Rolando Trujillo-Rasua. 2022. Analysis of centrality measures under differential privacy models. *Appl. Math. Comput.* 412 (2022), 126546.

[28] Xiaoye Li, Jing Yang, Zhenlong Sun, and Jianpei Zhang. 2017. Differential privacy for edge weights in social networks. *Security and Communication Networks* 2017 (2017).

[29] Yang Li, Michael Purcell, Thierry Rakotoarivelo, David Smith, Thilina Ranbaduge, and Kee Siong Ng. 2023. Private graph data release: A survey. *Comput. Surveys* 55, 11 (2023), 1–39.

[30] Zichun Liu, Liusheng Huang, Hongli Xu, Wei Yang, and Shaowei Wang. 2020. PrivAG: Analyzing attributed graph data with local differential privacy. In *2020 IEEE 26th International Conference on Parallel and Distributed Systems (ICPADS)*. IEEE, 422–429.

[31] Priya Mahadevan, Dmitri Krioukov, Kevin Fall, and Amin Vahdat. 2006. Systematic topology analysis and generation using degree correlations. *ACM SIGCOMM Computer Communication Review* 36, 4 (2006), 135–146.

[32] Frank McSherry and Kunal Talwar. 2007. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*. IEEE, 94–103.

[33] Frank D McSherry. 2009. Privacy integrated queries: an extensible platform for privacy-preserving data analysis. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of data*. 19–30.

[34] André LC Mendonça, Felipe T Brito, and Javam C Machado. 2023. Privacy-Preserving Techniques for Social Network Analysis. In *Anais Estendidos do XXXVIII Simpósio Brasileiro de Bancos de Dados*. SBC, 174–178.

[35] Tamara T Mueller, Dmitrii Usynin, Johannes C Paetzold, Rickmer Braren, Daniel Rueckert, and Georgios Kaissis. 2024. Differentially Private Guarantees for Analytics and Machine Learning on Graphs: A Survey of Results. *Journal of Privacy and Confidentiality* 14, 1 (2024).

[36] Hiep H Nguyen, Abdessamad Imine, and Michaël Rusinowitch. 2015. Differentially private publication of social graphs at linear cost. In *2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. IEEE, 596–599.

[37] Hiep H Nguyen, Abdessamad Imine, and Michaël Rusinowitch. 2016. Detecting communities under differential privacy. In *Proceedings of the 2016 ACM on Workshop on Privacy in the Electronic Society*. 83–93.

[38] M Usman Nisar, Arash Fard, and John A Miller. 2013. Techniques for graph analytics on big data. In *2013 IEEE International Congress on Big Data*. IEEE, 255–262.

[39] Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. 2007. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*. 75–84.

[40] Iyiola E Olatunji, Thorben Funke, and Megha Khosla. 2021. Releasing graph neural networks with differential privacy guarantees. *arXiv preprint arXiv:2109.08907* (2021).

[41] Leyla Roohi, Benjamin IP Rubinstein, and Vanessa Teague. 2019. Differentially-private two-party egocentric betweenness centrality. In *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, 2233–2241.

[42] Sina Sajadmanesh and Daniel Gatica-Perez. 2021. Locally private graph neural networks. In *Proceedings of the 2021 ACM SIGSAC conference on computer and communications security*. 2130–2145.

[43] Alessandra Sala, Xiaohan Zhao, Christo Wilson, Haitao Zheng, and Ben Y Zhao. 2011. Sharing graphs using differentially private graph models. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*. 81–98.

[44] Adam Sealfon. 2016. Shortest paths and distances with differential privacy. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*. 29–41.

[45] Haipei Sun, Xiaokui Xiao, Issa Khalil, Yin Yang, Zhan Qin, Hui Wang, and Ting Yu. 2019. Analyzing subgraph statistics from extended local views with decentralized differential privacy. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. 703–717.

[46] Christine Task and Chris Clifton. 2012. A guide to differential privacy theory in social network analysis. In *2012 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. IEEE, 411–417.

[47] United States Census Bureau. 2021. Census Bureau Sets Key Parameters to Protect Privacy in 2020 Census Results. https://www.census.gov/newsroom/press-releases/2021/2020-census-key-parameters.html. Online; accessed 10 January 2023.

[48] Paul Voigt and Axel Von dem Bussche. 2017. The eu general data protection regulation (gdpr). *A Practical Guide, 1st Ed., Cham: Springer International Publishing* 10, 3152676 (2017), 10–5555.

[49] Yue Wang and Xintao Wu. 2013. Preserving differential privacy in degree-correlation based graph generation. *Transactions on data privacy* 6, 2 (2013), 127.

[50] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* 32, 1 (2020), 4–24.

[51] Qian Xiao, Rui Chen, and Kian-Lee Tan. 2014. Differentially private network data release via structural inference. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 911–920.

[52] Jun Zhang, Graham Cormode, Cecilia M Procopiuc, Divesh Srivastava, and Xiaokui Xiao. 2015. Private release of graph statistics using ladder functions. In *Proceedings of the 2015 ACM SIGMOD international conference on management of data*. 731–745.

[53] Tianqing Zhu, Dayong Ye, Wei Wang, Wanlei Zhou, and S Yu Philip. 2020. More than privacy: Applying differential privacy in key areas of artificial intelligence. *IEEE Transactions on Knowledge and Data Engineering* 34, 6 (2020), 2824–2843.