
Certified Graph Unlearning

Eli Chien*

ECE
UIUC

ichien3@illinois.edu

Chao Pan*

ECE
UIUC

chaopan2@illinois.edu

Olgica Milenkovic

ECE
UIUC

milenkov@illinois.edu

Abstract

Graph-structured data is ubiquitous in practice and often processed using graph neural networks (GNNs). With the adoption of recent laws ensuring the “right to be forgotten”, the problem of graph data removal has become of significant importance. To address the problem, we introduce the first known framework for *certified graph unlearning* of GNNs. In contrast to standard machine unlearning, new analytical and heuristic unlearning challenges arise when dealing with complex graph data. First, three different types of unlearning requests need to be considered, including node feature, edge and node unlearning. Second, to establish provable performance guarantees, one needs to address challenges associated with feature mixing during propagation. The underlying analysis is illustrated on the example of simple graph convolutions (SGC) and their generalized PageRank (GPR) extensions, thereby laying the theoretical foundation for certified unlearning of GNNs. Our empirical studies on six benchmark datasets demonstrate excellent performance-complexity trade-offs when compared to complete retraining methods and approaches that do not leverage graph information. For example, when unlearning 20% of the nodes on the Cora dataset, our approach suffers only a 0.1% loss in test accuracy while offering a 4-fold speed-up compared to complete retraining. Our scheme also outperforms unlearning methods that do not leverage graph information with a 12% increase in test accuracy for comparable time complexity. Our implementation is available online.²

1 Introduction

Machine learning algorithms are used in many application domains, including biology, computer vision and natural language processing. Relevant models are often trained either on third-party datasets, internal or customized subsets of publicly available user data. For example, many computer vision models are trained on images of Flickr users [1, 2], while many natural language processing (e.g., sentiment analysis) and recommender systems heavily rely on repositories such as IMDB [3]. Furthermore, numerous ML classifiers in computational biology are trained on data from the UK Biobank [4], which represents a collection of genetic and medical records of roughly half a million patients [5]. With recent demands for increased data privacy, the above referenced and many other data repositories are facing increasing demands for data removal. Certain laws are already in place guaranteeing the rights of certified data removal, including the European Union’s General Data Protection Regulation (GDPR), the California Consumer Privacy Act (CCPA) and the Canadian Consumer Privacy Protection Act (CPPA) [6].

Removing user data from a dataset is insufficient to guarantee the desired level of privacy, since models trained on the original data may still contain information about their patterns and features.

*Equal contribution.

²https://github.com/thupchnsky/sgc_unlearn.




Figure 1: Illustration of three different types of certified graph unlearning problems and a comparison with the case of unlearning without graph information [2]. The colors of the nodes capture properties of node features, and the red frame indicates node embeddings affected by 1-hop propagation. When no graph information is used, the node embeddings are uncorrelated. However, for the case of graph unlearning problems, removing one node or edge can affect the node embeddings of the entire graph for a large enough number of propagation steps.

This consideration gave rise to a new research direction in machine learning, referred to as *machine unlearning* [7], in which the goal is to guarantee that the user data information is also removed from the trained model. Naively, one can retrain the model from scratch to meet the privacy demand, yet retraining comes at a high computational cost and is thus not practical when accommodating frequent removal requests. To avoid complete retraining, various methods for machine unlearning have been proposed, including *exact approaches* [5, 8] as well as *approximate methods* which often come with performance guarantees and are therefore certified unlearners [2, 6].

At the same time, graph-centered machine learning has received significant interest from the learning community due to the ubiquity of graph-structured data. Usually, the data contains two sources of information: Node features and graph topology. Graph Neural Networks (GNN) leverage both types of information simultaneously and achieve state-of-the-art performance in numerous real-world applications, including Google Maps [9], various recommender system [10], self-driving cars [11] and bioinformatics [12]. Clearly, user data is involved in training the underlying GNNs and it may therefore be subject to removal. However, it is still unclear how to perform unlearning of GNNs.

We take the first step towards solving the approximate unlearning problem for graphs by performing a nontrivial theoretical analysis of some simplified GNN architectures. Inspired by the unstructured data *certified removal procedure* [2], we propose the first known approach for *certified graph unlearning*. Our main contributions are as follows. First, we introduce three types of data removal requests for graph unlearning: Node feature unlearning, edge unlearning and node unlearning (see Figure 1). Second, we derive theoretical guarantees for certified graph unlearning mechanisms for all three removal cases on SGCs [13] and their GPR generalizations. In particular, we analyze L_2 -regularized graph models trained with differentiable convex loss functions. The analysis is challenging since propagation on graphs “mixes” node features. Our analysis reveals the intuitive observation that the degree of the unlearned node plays an important role in the unlearning process, while the number of propagation steps may or may not be important for different unlearning scenarios. To the best of our knowledge, the theoretical guarantees established in this work are the first provable certified removal studies for graphs. Furthermore, the proposed analysis also encompasses node classification and node regression problems. Third, our empirical investigation on frequently used datasets for GNN learning shows that our method offers an excellent performance-complexity trade-off. For example, when unlearning 20% nodes on the Cora dataset, our method achieves a 4-fold speedup with only a 0.1% drop in test accuracy compared to complete retraining. We also test our model on datasets for which removal requests are more likely to arise, including the Amazon co-purchase networks. All proofs and detailed discussions of the results are relegated to the Appendix.

2 Related Works

Machine unlearning and certified data removal. The paper [7] introduced the concept of machine unlearning and proposed distributed learners for exact unlearning, while [8] introduced sharding-based methods for unlearning. The authors of [5] described unlearning approaches for k -means clustering. All the aforementioned works focused on *exact unlearning*: The unlearned model is required to perform identically to a completely retrained model. As an alternative, [2] introduced a probabilistic definition of unlearning motivated by differential privacy [14] and proposed the study of “approximate” unlearning. This work was followed by [6], which studied the generalization




Figure 2: Difference between machine unlearning (as defined in [2]) and Differential Privacy (DP).

performance of machine unlearning methods, and [15], which proposed heuristic-based selective forgetting in deep networks. None of these works addressed the machine unlearning problem for graphs. To the best of our knowledge, the only work in this direction is the preprint [16]. However, the strategy proposed therein is sharding, which only works for exact unlearning and is hence completely different from our approximate approach. Also, the approach in [16] relies on partitioning the graph using community detection methods. It therefore implicitly makes the assumption that the graph is homophilic which is not warranted in practice [17, 18]. In contrast, our method works for arbitrary graphs and allows for approximate unlearning while ensuring excellent trade-offs between performance and complexity.

Differential privacy (DP) and DP-GNNs. Machine unlearning, especially the approximation version described in [2], is closely related to differential privacy [14]. In fact, differential privacy is a *sufficient* for machine unlearning. If a model is differentially private, then the adversary cannot distinguish whether the model is trained on the original dataset or on a dataset in which one data point is removed. Hence, even *without* model updating, a DP model will automatically unlearn the removed data point (see also the explanation in [5, 6] and Figure 2). Although DP is sufficient for unlearning, it is not necessary. Also, most of the DP models suffer from a significant degradation in performance even when the privacy constraint is loose [19, 20]. Machine unlearning can therefore be viewed as a means to trade-off between performance and computational cost, with complete retraining and DP on two different ends of the spectrum [2]. Several recent works proposed DP-GNNs [21, 22, 23, 24] – however, even for unlearning *one single node or edge*, these methods require a very high “privacy budget” to learn with sufficient accuracy.

3 Preliminaries

Notation. We reserve bold-font capital letters such as \mathbf{S} for matrices and bold-font lowercase letters such as \mathbf{s} for vectors. We use \mathbf{e}_i to denote the i^{th} standard basis, so that $\mathbf{e}_i^T \mathbf{S}$ and $\mathbf{S} \mathbf{e}_i$ represent the i^{th} row and column vector of \mathbf{S} , respectively. The absolute value $|\cdot|$ is applied component-wise on both matrices and vectors. We also use the symbols $\mathbf{1}$ for the all-one vector and \mathbf{I} for the identity matrix. Furthermore, we let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ stand for an undirected graph with node set $\mathcal{V} = [n]$ of size n and edge set \mathcal{E} . The symbols \mathbf{A} and \mathbf{D} are used to denote the corresponding adjacency and node degree matrix, respectively. The feature matrix is denoted by $\mathbf{X} \in \mathbb{R}^{n \times F}$ and the features have dimension F ; For binary classification, the labels are summarized in $\mathbf{Y} \in \{-1, 1\}^n$, while the nonbinary case is discussed in Section 5. The relevant norms are $\|\cdot\|$, the l_2 norm, and $\|\cdot\|_F$, the Frobenius norm. Note that we use $\|\cdot\|$ for both row and column vectors to simplify the notation. The matrices \mathbf{A} and \mathbf{D} should not be confused with the symbols for an algorithm A and dataset \mathcal{D} .

Certified removal. Let A be a (randomized) learning algorithm that trains on \mathcal{D} , the set of data points before removal, and outputs a model $h \in \mathcal{H}$, where \mathcal{H} represents a chosen space of models. The removal of a subset of points from \mathcal{D} results in \mathcal{D}' . For instance, let $\mathcal{D} = (\mathbf{X}, \mathbf{Y})$. Suppose we want to remove a data point, $(\mathbf{e}_i^T \mathbf{X}, \mathbf{e}_i^T \mathbf{Y})$ from \mathcal{D} , resulting in $\mathcal{D}' = (\mathbf{X}', \mathbf{Y}')$. Here, \mathbf{X}', \mathbf{Y}' are equal to \mathbf{X}, \mathbf{Y} , respectively, except that the row corresponding to the removed data point is deleted. Given $\epsilon, \delta > 0$, an unlearning algorithm M applied to $A(\mathcal{D})$ is said to guarantee an (ϵ, δ) -certified removal for A , where \mathcal{X} denotes the space of possible datasets, if $\forall \mathcal{T} \subseteq \mathcal{H}, \mathcal{D} \subseteq \mathcal{X}, i \in [n]$,

$$\begin{aligned} \mathbb{P}(M(A(\mathcal{D}), \mathcal{D}, \mathcal{D} \setminus \mathcal{D}') \in \mathcal{T}) &\leq e^\epsilon \mathbb{P}(A(\mathcal{D}') \in \mathcal{T}) + \delta, \\ \mathbb{P}(A(\mathcal{D}') \in \mathcal{T}) &\leq e^\epsilon \mathbb{P}(M(A(\mathcal{D}), \mathcal{D}, \mathcal{D} \setminus \mathcal{D}') \in \mathcal{T}) + \delta. \end{aligned} \quad (1)$$

This definition is related to (ϵ, δ) -DP [14] except that we are allowed to update the model based on the removed point (see Figure 2). An (ϵ, δ) -certified removal method guarantees that the updated

model $M(A(\mathcal{D}), \mathcal{D}, \mathcal{D} \setminus \mathcal{D}')$ is “approximately” the same as the model $A(\mathcal{D}')$ obtained by retraining from scratch. Thus, any information about the removed data $\mathcal{D} \setminus \mathcal{D}'$ is “approximately” eliminated from the model. Ideally, we would like to design M such that it satisfies equation (1) and has a complexity that is significantly smaller than that of complete retraining.

4 Certified Graph Unlearning

Unlike standard machine unlearning, certified graph unlearning uses datasets that contain not only node features \mathbf{X} but also the graph topology \mathbf{A} , and therefore require different data removal procedures. We focus on node classification, for which the training dataset equals $\mathcal{D} = (\mathbf{X}, \mathbf{Y}_{T_r}, \mathbf{A})$. Here, \mathbf{Y}_{T_r} is identical to \mathbf{Y} on rows indexed by points of the training set T_r while the remaining rows are all zeros. Wlog, we assume that the training set comprises the first m nodes (i.e. $T_r = [m]$), where $m \leq n$. An unlearning method M achieves (ϵ, δ) -certified graph unlearning with algorithm A if (1) is satisfied for $\mathcal{D} = (\mathbf{X}, \mathbf{Y}_{T_r}, \mathbf{A})$ and \mathcal{D}' , which differ based on the type of graph unlearning: Node feature unlearning, edge unlearning, and node unlearning.

4.1 Unlearning SGC

SGC is a simplification of GCN obtained by removing all nonlinearities from the latter model. This leads to the following update rule: $\mathbf{P}^K \mathbf{X} \mathbf{W} \triangleq \mathbf{Z} \mathbf{W}$, where \mathbf{W} denotes the matrix of learnable weights, $K \geq 0$ equals the number of propagation steps and \mathbf{P} denotes the one-step propagation matrix. The standard choice of the propagation matrix is the symmetric normalized adjacency matrix with self-loops, $\mathbf{P} = \tilde{\mathbf{D}}^{-1/2} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-1/2}$, where $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ and $\tilde{\mathbf{D}}$ equals the degree matrix with respect to $\tilde{\mathbf{A}}$. We will work with the asymmetric normalized version of \mathbf{P} , $\mathbf{P} = \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}}$. This choice is made purely for analytical purposes and our empirical results confirm that this normalization ensures the competitive performance of our unlearning methods.

The resulting node embedding is used for node classification by choosing an appropriate loss (i.e., logistic loss) and minimizing the L_2 -regularized empirical risk. For binary classification, \mathbf{W} can be replaced by a vector \mathbf{w} ; the loss equals $L(\mathbf{w}, \mathcal{D}) = \sum_{i: \mathbf{e}_i^T \mathbf{Y}_{T_r} \neq 0} (\ell(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}, \mathbf{e}_i^T \mathbf{Y}_{T_r}) + \frac{\lambda}{2} \|\mathbf{w}\|^2)$, where $\ell(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}, \mathbf{e}_i^T \mathbf{Y}_{T_r})$ is a convex loss function that is differentiable everywhere. We also write $\mathbf{w}^* = A(\mathcal{D}) = \operatorname{argmin}_{\mathbf{w}} L(\mathbf{w}, \mathcal{D})$, where the optimizer is unique whenever $\lambda > 0$.

Motivated by the unlearning approach from [2] pertaining to unstructured data, we design an unlearning mechanism M for graphs that changes the trained model \mathbf{w}^* to \mathbf{w}^- which represents an approximation of the unique optimizer of $L(\mathbf{w}, \mathcal{D}')$. Denote the Hessian of $L(\cdot, \mathcal{D}')$ at \mathbf{w}^* by $\mathbf{H}_{\mathbf{w}^*} = \nabla^2 L(\mathbf{w}^*, \mathcal{D}')$. The authors of [2] propose the following unlearning mechanism for unlearning the m^{th} training point: $\mathbf{w}^- = M(\mathbf{w}^*, \mathcal{D}, \mathcal{D} \setminus \mathcal{D}') = \mathbf{w}^* + \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta_{guo}$, where $\Delta_{guo} = \lambda \mathbf{w}^* + \nabla \ell(\mathbf{e}_m^T \mathbf{X} \mathbf{w}^*, \mathbf{e}_m^T \mathbf{Y}_{T_r})$. In order to address graph unlearning, we generalize this unlearning mechanism by replacing Δ_{guo} with $\Delta = \nabla L(\mathbf{w}^*, \mathcal{D}) - \nabla L(\mathbf{w}^*, \mathcal{D}')$. More precisely, we have

$$\Delta = \lambda \mathbf{w}^* + \nabla \ell(\mathbf{e}_m^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_m^T \mathbf{Y}_{T_r}) + \sum_{i=1}^{m-1} [\nabla \ell(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r}) - \nabla \ell(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})]. \quad (2)$$

Note that our generalized unlearning mechanism matches the one in [2] when no graph information is involved. This can be seen by setting $K = 0$, which leads to $\mathbf{Z} = \mathbf{X}$ and $\mathbf{e}_i^T \mathbf{Z} = \mathbf{e}_i^T \mathbf{Z}' \forall i \in [m-1]$. Hence, the third term in equation (2) is zero and thus $\Delta = \Delta_{guo}$. Note that the third term in equation (2) is not zero in general when graph information is involved. This also highlights the main difficulty of directly applying the analysis of [2] to graphs, as we need to take care of the third term in equation (2). More details can be found in Appendix 7.2.

When $\nabla L(\mathbf{w}^-, \mathcal{D}') = 0$, \mathbf{w}^- is the unique optimizer of $L(\cdot, \mathcal{D}')$. If $\nabla L(\mathbf{w}^-, \mathcal{D}') \neq 0$, information about the removed data point remains present. One can show that the gradient residual norm $\|\nabla L(\mathbf{w}^-, \mathcal{D}')\|$ determines the error of \mathbf{w}^- when used to approximate the true minimizer of $L(\cdot, \mathcal{D}')$ [2]. Hence, upper bounds on $\|\nabla L(\mathbf{w}^-, \mathcal{D}')\|$ can be used to establish certified removal/unlearning guarantees. More precisely, assume that we have $\|\nabla L(\mathbf{w}^-, \mathcal{D}')\| \leq \epsilon'$ for some $\epsilon' > 0$. Furthermore, consider training with the noisy loss $L_b(\mathbf{w}, \mathcal{D}) = \sum_{i: \mathbf{e}_i^T \mathbf{Y}_{T_r} \neq 0} (\ell(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}, \mathbf{e}_i^T \mathbf{Y}_{T_r}) + \frac{\lambda}{2} \|\mathbf{w}\|^2) + \mathbf{b}^T \mathbf{w}$, where \mathbf{b} is drawn randomly according to some distribution. Then one can leverage the following result.

Theorem 4.1 (Theorem 3 from [2]). *Let A be the learning algorithm that returns the unique optimum of the loss $L_{\mathbf{b}}(\mathbf{w}, \mathcal{D})$. Suppose that $\|\nabla L(\mathbf{w}^-, \mathcal{D}')\| \leq \epsilon'$ for some computable bound $\epsilon' > 0$, independent of \mathbf{b} and achieved by M . If $\mathbf{b} \sim \mathcal{N}(\mathbf{0}, c_0 \epsilon' / \epsilon \cdot \mathbf{I})$ with $c_0 > 0$, then M satisfies (1) with (ϵ, δ) for algorithm A applied to \mathcal{D}' , where $\delta = 1.5e^{-c_0^2/2}$.*

Hence, if we are able to prove that $\|\nabla L(\mathbf{w}^-, \mathcal{D}')\|$ is appropriately bounded for our graph setting as well, then M will ensure (ϵ, δ) -certified graph unlearning. Our main technical contribution is to establish such bounds for all three types of unlearning methods for graphs. For the analysis, we need the loss function ℓ to satisfy the following properties.

Assumption 4.2. *For any \mathcal{D} , $i \in [n]$ and $\mathbf{w} \in \mathbb{R}^F$: (1) $\|\nabla \ell(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}, \mathbf{e}_i^T \mathbf{Y})\| \leq c$ (i.e. the norm of $\nabla \ell$ is c -bounded); (2) ℓ' is γ_2 -Lipschitz; (3) $\|\mathbf{e}_i^T \mathbf{X}\| \leq 1$; (4) ℓ' is γ_1 -Lipschitz; (5) ℓ' is c_1 -bounded.*

Assumptions (1)-(3) are also needed for unstructured unlearning of linear classifiers [2]. To account for graph-structured data, we require the additional assumptions (4)-(5) to establish worst-case bounds, which can be avoided when working with data-dependent bounds (Section 5).

In all subsequent derivations, we assume that the unlearned data point corresponds to the m^{th} node for node feature and node unlearning; for edge unlearning, we wish to unlearn the edge $(1, m)$. Generalizations for multiple unlearning requests are discussed in Section 5.

4.2 Node feature unlearning for SGC

We start with the simplest type of unlearning – node feature unlearning – for SGCs. In this case, we remove the node feature and label of one node from \mathcal{D} , resulting in $\mathcal{D}' = (\mathbf{X}', \mathbf{Y}'_{T_r}, \mathbf{A})$. The matrices \mathbf{X}' , \mathbf{Y}'_{T_r} are identical to \mathbf{X} , \mathbf{Y}_{T_r} , respectively, except for the m^{th} row of the former being set to zero. Note that in this case, the graph structure remains unchanged.

Theorem 4.3. *Suppose that Assumption 4.2 holds. For the node feature unlearning scenario, and for $\mathbf{Z} = \mathbf{P}^K \mathbf{X}$ and $\mathbf{P} = \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}}$, we have*

$$\|\nabla L(\mathbf{w}^-, \mathcal{D}')\| = \|(\mathbf{H}_{\mathbf{w}_\eta} - \mathbf{H}_{\mathbf{w}^*}) \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta\| \leq \frac{\gamma_2(2c\lambda + (c\gamma_1 + \lambda c_1) \tilde{\mathbf{D}}_{mm})^2}{\lambda^4(m-1)}, \quad (3)$$

where $\mathbf{H}_{\mathbf{w}_\eta}$ denotes the Hessian of $L(\cdot, \mathcal{D}')$ at $\mathbf{w}_\eta = \mathbf{w}^* + \eta \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta$, for some $\eta \in [0, 1]$. A similar conclusion holds for the case when we wish to unlearn node features of a node that is not in T_r . In this case we just replace $\tilde{\mathbf{D}}_{mm}$ by the degree of the corresponding node. This result shows that the norm bound is large if the unlearned node has a large degree, since a large-degree node will affect the values of many rows in \mathbf{Z} . Our result also demonstrates that the norm bound is *independent* of K , due to the fact that \mathbf{P} is right stochastic. We provide next a sketch of the proof to illustrate the analytical challenges of graph unlearning compared to those of unstructured data unlearning.

Although for node feature unlearning the graph topology does not change, all rows of $\mathbf{Z} = \mathbf{P}^K \mathbf{X}$ may potentially change due to graph information propagation. Thus, the original analysis from [2], which corresponds to the special case $\mathbf{Z} = \mathbf{X}$, cannot be applied directly. There are two particular challenges. The first is to ensure that the norm of each row of \mathbf{Z} is bounded by 1. We provide Lemma 7.1 to guarantee this. It is critical to choose $\mathbf{P} = \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}}$ since all other choices of degree normalization lead to worse bounds (see Appendix 7.9). The second and more difficult challenge is to bound $\|\Delta\|$. When $\mathbf{Z} = \mathbf{X}$, the third term in equation (2) is exactly zero, in accordance with [2]. Due to graph propagation, we have to further bound the norm of the third term, which is highly nontrivial since the upper bound is not allowed to grow with m or n . We first focus on one of the $m-1$ terms in the sum. Using Assumption 4.2, one can bound this term by $\|\mathbf{e}_i^T (\mathbf{Z} - \mathbf{Z}')\|$ (we suppressed the dependency on λ, c, c_1 and γ_1 for simplicity). The key analytical novelty is to explore the sparsity of $\mathbf{Z} - \mathbf{Z}' = \mathbf{P}^K (\mathbf{X} - \mathbf{X}')$. Note that $\mathbf{X} - \mathbf{X}'$ is an all-zero matrix except for its m^{th} row being equal to $\mathbf{e}_m^T \mathbf{X}$. Thus, we have $\|\mathbf{e}_i^T (\mathbf{Z} - \mathbf{Z}')\| = \|\mathbf{e}_i^T \mathbf{P}^K (\mathbf{X} - \mathbf{X}')\| = \|\mathbf{e}_i^T \mathbf{P}^K \mathbf{e}_m \mathbf{e}_m^T \mathbf{X}\| \leq \mathbf{e}_i^T \mathbf{P}^K \mathbf{e}_m$, where the last bound follows from the Cauchy-Schwartz inequality, (3) in Assumption 4.2 and the fact that \mathbf{P}^K is a (component-wise) nonnegative matrix. Thus, summing over $i \in [m-1]$ leads to the upper bound $\mathbf{1}^T \mathbf{P}^K \mathbf{e}_m$, since $m \leq n$. Next, observe that $\mathbf{1}^T \mathbf{P}^K \mathbf{e}_m = \mathbf{1}^T \mathbf{P}^K \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{D}} \mathbf{e}_m = \mathbf{1}^T (\tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}})^K \tilde{\mathbf{D}}^{-1} \mathbf{e}_m \tilde{\mathbf{D}}_{mm} = \mathbf{1}^T \tilde{\mathbf{D}}^{-1} (\tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-1})^K \mathbf{e}_m \tilde{\mathbf{D}}_{mm}$. Since $\tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-1}$ is a left stochastic matrix, $\tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-1} \mathbf{p}$ is a probability vector whenever \mathbf{p} is a probability

vector. Clearly, \mathbf{e}_m is a probability vector. Hence, $(\tilde{\mathbf{A}}\tilde{\mathbf{D}}^{-1})^K \mathbf{e}_m$ is also a probability vector. Since all diagonal entries of $\tilde{\mathbf{D}}^{-1}$ are nonnegative and upper bounded by 1 given the self-loops for all nodes, $\mathbf{1}^T \tilde{\mathbf{D}}^{-1} \mathbf{p} \leq \mathbf{1}^T \mathbf{p} = 1$ for any probability vector \mathbf{p} . Hence, the term above is bounded by $\tilde{\mathbf{D}}_{mm}$. The bound depends on $\tilde{\mathbf{D}}_{mm}$ and does not increase with m or K . Although node feature unlearning is the simplest case of graph unlearning, our sketch of the proof illustrates the difficulties associated with bounding the third term in Δ . Similar, but more complicated approaches are needed for the analysis of edge unlearning and node unlearning.

4.3 Edge and node unlearning for SGC and GPR extension

Edge unlearning for SGC. We describe next the bounds for edge unlearning and highlight the technical issues arising in the analysis of this setting. Here, we remove one edge $(1, m)$ from \mathcal{D} , resulting in $\mathcal{D}' = (\mathbf{X}, \mathbf{Y}_{T_r}, \mathbf{A}')$. The matrix \mathbf{A}' is identical to \mathbf{A} except for $\tilde{\mathbf{A}}'_{1m} = \tilde{\mathbf{A}}'_{m1} = 0$. Furthermore, $\tilde{\mathbf{D}}'$ is the degree matrix corresponding to $\tilde{\mathbf{A}}'$. Note that the node features and labels remain unchanged.

Theorem 4.4. *Suppose that Assumption 4.2 holds. For the edge unlearning scenario, and for $\mathbf{P} = \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}}$ and $\mathbf{Z} = \mathbf{P}^K \mathbf{X}$, we have*

$$\|\nabla L(\mathbf{w}^-, \mathcal{D}')\| = \|(\mathbf{H}_{\mathbf{w}_\eta} - \mathbf{H}_{\mathbf{w}^*}) \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta\| \leq \frac{16\gamma_2 K^2 (c\gamma_1 + c_1\lambda)^2}{\lambda^4 m}. \quad (4)$$

Node unlearning for SGC. We now discuss the most difficult case, node unlearning. In this case, one node is entirely removed from \mathcal{D} , including node features, labels and edges. This results in $\mathcal{D}' = (\mathbf{X}', \mathbf{Y}'_{T_r}, \mathbf{A}')$. The matrices \mathbf{X}' , \mathbf{Y}'_{T_r} are defined similarly to those described for node feature unlearning. The matrix \mathbf{A}' is obtained by replacing the m^{th} row and column in \mathbf{A} by all-zeros (similar changes are introduced in $\tilde{\mathbf{A}}$, with $\tilde{\mathbf{A}}_{mm} = 0$). For simplicity, we let $\tilde{\mathbf{D}}'_{mm} = 1$ as this assumption does not affect the propagation results.

Theorem 4.5. *Suppose that Assumption 4.2 holds. For the node unlearning scenario, and for $\mathbf{Z} = \mathbf{P}^K \mathbf{X}$ and $\mathbf{P} = \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}}$, we have*

$$\|\nabla L(\mathbf{w}^-, \mathcal{D}')\| = \|(\mathbf{H}_{\mathbf{w}_\eta} - \mathbf{H}_{\mathbf{w}^*}) \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta\| \leq \frac{\gamma_2 \left(2c\lambda + K(c\gamma_1 + c_1\lambda) \left(2\tilde{\mathbf{D}}_{mm} - 1\right)\right)^2}{\lambda^4 (m-1)}. \quad (5)$$

The main challenge arising in the proofs of Theorem 4.4 and 4.5 is bounding $\|\Delta\|$ appropriately. Unlike for the node feature unlearning case, now both graph structure and node features can change due to the unlearning request. We establish a series of lemmas to characterize the differences between \mathbf{Z} and \mathbf{Z}' , which play an important role in our proofs (see Appendix 7.6 and 7.7 for complete proofs).

Certified graph unlearning in GPR-based model. Our analysis can be extended to Generalized PageRank (GPR)-based models [25]. The definition of GPR is $\sum_{k=0}^K \theta_k \mathbf{P}^k \mathbf{S}$, where \mathbf{S} denotes a node feature or node embedding. The learnable weights θ_k are called GPR weights and different choices for the weights lead to different propagation rules [26, 27]. GPR-type propagations include SGC and APPNP rules as special cases [17]. If we use linearly transformed features $\mathbf{S} = \mathbf{X}\bar{\mathbf{W}}$, for some weight matrix $\bar{\mathbf{W}}$, the GPR rule can be rewritten as $\mathbf{Z}\mathbf{W} = \frac{1}{K+1} [\mathbf{X}, \mathbf{P}\mathbf{X}, \mathbf{P}^2\mathbf{X}, \dots, \mathbf{P}^K\mathbf{X}] \mathbf{W}$. This constitutes a concatenation of the steps from 0 up to K . The learnable weight matrix $\mathbf{W} \in \mathbb{R}^{(K+1)F \times C}$ combines θ_k and $\bar{\mathbf{W}}$. These represent linearizations of GPR-GNNs [17] and SIGNs [28], simple yet useful models for learning on graphs. For simplicity, we only describe the results for node feature unlearning and delegate the analysis of edge and node unlearning to Appendix 7.8.

Theorem 4.6. *Suppose that Assumption 4.2 holds and considers the node feature unlearning case. For $\mathbf{Z} = \frac{1}{K+1} [\mathbf{X}, \mathbf{P}\mathbf{X}, \mathbf{P}^2\mathbf{X}, \dots, \mathbf{P}^K\mathbf{X}]$ and $\mathbf{P} = \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}}$, we have*

$$\|\nabla L(\mathbf{w}^-, \mathcal{D}')\| = \|(\mathbf{H}_{\mathbf{w}_\eta} - \mathbf{H}_{\mathbf{w}^*}) \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta\| \leq \frac{\gamma_2 (2c\lambda + (c\gamma_1 + \lambda c_1) \tilde{\mathbf{D}}_{mm})^2}{\lambda^4 (m-1)}. \quad (6)$$

Note that the resulting bound is the same as the bound in Theorem 4.3. This is due to the fact that we used the normalization factor $\frac{1}{K+1}$ in \mathbf{Z} . Hence, given the same noise level, the GPR-based models are more sensitive when we trained on the noisy loss $L_{\mathbf{b}}$. Whether the general high-level performance of GPR can overcompensate this drawback depends on the actual datasets considered.

5 Empirical Aspect of Certified Graph Unlearning

Logistic and least-squares regression on graphs. For binary logistic regression, the loss equals $\ell(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}, \mathbf{e}_i^T \mathbf{Y}_{T_r}) = -\log(\sigma(\mathbf{e}_i^T \mathbf{Y}_{T_r} \mathbf{e}_i^T \mathbf{Z} \mathbf{w}))$, where $\sigma(x) = 1/(1+\exp(-x))$ denotes the sigmoid function. As shown in [2], the assumptions (1)-(3) in 4.2 are satisfied with $c = 1$ and $\gamma_2 = 1/4$. We only need to show that (4) and (5) of 4.2 hold as well. By standard analysis, we show that our loss satisfies (4) and (5) in 4.2 with $\gamma_1 = 1/4$ and $c_1 = 1$. For multi-class logistic regression, one can adapt the “one-versus-all other-classes” strategy which leads to the same result. For least-square regression, since the hessian is independent of \mathbf{w} our approach offers $(0, 0)$ -certified graph unlearning even without loss perturbations. See Appendix 7.3 for the complete discussion and derivation.

Sequential unlearning. In practice, multiple users may request unlearning. Hence, it is desirable to have a model that supports sequential unlearning of all types of data points. One can leverage the same proof as in [2] (induction coupled with the triangle inequality) to show that the resulting gradient residual norm bound equals $T\epsilon'$ at the T^{th} unlearning request, where ϵ' is the bound for a single instance of certified graph unlearning. We relegate the complete proof in Appendix 7.17 for completeness.

Data-dependent bounds. The gradient residual norm bounds derived for different types of certified graph unlearning contain a constant factor $1/\lambda^4$, and may be loose in practice. Following [2], we also examined data-dependent bounds.

Corollary 5.1 (Application of Corollary 1 in [2]). *For all three graph unlearning scenarios, we have $\|\nabla L(\mathbf{w}^-, \mathcal{D}')\| \leq \gamma_2 \|\mathbf{Z}'\|_{op} \|\mathbf{H}_{\mathbf{w}^*}^{-1} \Delta\| \|\mathbf{Z}' \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta\|$.*

Hence, there are two ways to accomplish certified graph unlearning. If we do not allow any retraining, we have to leverage the worst-case bound in Section 4 based on the expected number of unlearning requests. Importantly, we will also need to constrain the node degree of nodes to be unlearned (i.e., do not allow for unlearning hub nodes), for both node feature and node unlearning. Otherwise, we can select the noise standard deviation α , ϵ and δ and compute the corresponding “privacy budget” $\alpha\epsilon/\sqrt{2\log(1.5/\delta)}$. Once the accumulated gradient residual norm exceeds this budget, we retrain the model from scratch. Note that this still greatly reduces the time complexity compare to retraining the model for every unlearning request (see Section 6). We provide the pseudo-code for our method that leverages data-dependent bounds for sequential unlearning in Appendix 7.4.

6 Experiment

We test our certified graph unlearning methods by verifying theorems and via comparisons with baseline methods on benchmark datasets.

Settings. We test our methods on benchmark datasets for graph learning, including Cora, Citseer, Pubmed [29, 30, 31] and large-scale dataset ogbn-arxiv [32] and Amazon co-purchase networks Computers and Photo [33, 34]. We either use the public splitting or random splitting based on similar rules as public splitting and focus on node classification. Following [2], we use LBFGS as the optimizer for all methods due to its high efficiency on strongly convex problems. Unless specified otherwise, we fix $K = 2$, $\delta = 10^{-4}$, $\lambda = 10^{-2}$, $\epsilon = 1$, $\alpha = 0.1$ for all experiments, and average the results over 5 independent trails with random initializations. Our baseline methods include complete retraining with graph information after each unlearning request (SGC Retraining), complete retraining without graph information after each unlearning request (No Graph Retraining), and Algorithm 2 in [2]. Additional details can be found in Appendix 7.18.

Dependency on node degrees. While an upper bound does not necessarily capture the dependency of each term correctly, we show in Figure 4 (a) and (b) that our Theorem 4.5 and 4.6 indeed do so. Here, each point corresponds to unlearning one node. We test for all nodes in the training set and fix $\lambda = 10^{-4}$, $\alpha = 0$. Our results show that unlearning a large-degree node is more expensive in terms of the privacy budget (i.e., it induces a larger gradient residual norm). The node degree dependency is unclear if one merely examines Corollary 5.1. For other datasets, refer to Appendix 7.18. The worst-case bounds are looser than the data-dependent bounds, matching the observation from [2].

Performance of certified graph unlearning methods. The performance of our proposed certified graph unlearning methods, including the time complexity of unlearning and test accuracy after unlearning, is shown in Figures 3 and 4. It shows that: (1) Leveraging graph information is necessary

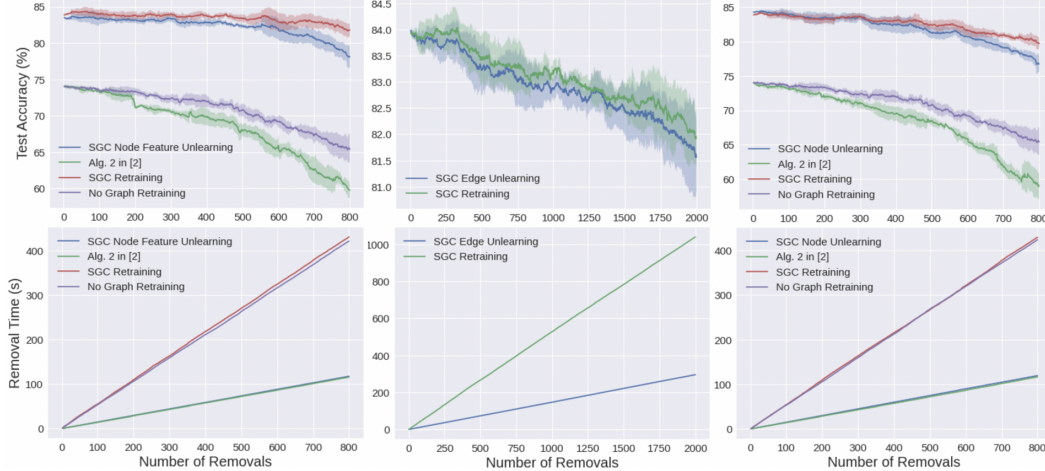


Figure 3: Comparison of proposed SGC node feature unlearning (left column), edge unlearning (middle column) and node unlearning (right column) with baseline methods. The shaded regions in the first row represent the standard deviation of test accuracy. In the second row, we show the accumulated unlearning time as a function of the number of unlearned points. The time needed for each unlearning procedure is given in Appendix 7.18.

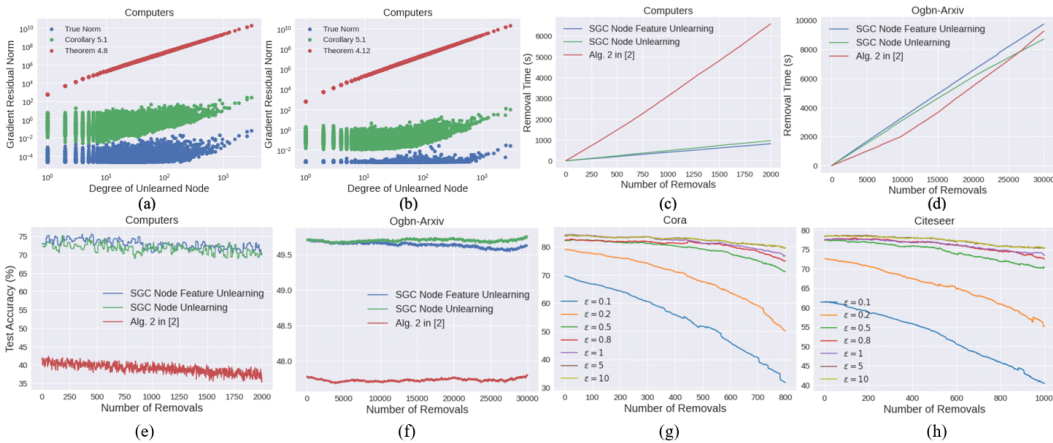


Figure 4: (a), (b) Simulation verification of the result in Theorem 4.5 and 4.6 pertaining to node degrees. (c), (d) Accumulated unlearning time as a function of the number of removed points. The unlearning time of Algorithm 2 from [2] is often higher than that of our proposed certified graph unlearning algorithms, because the number of retraining steps needed may be larger. (e), (f) Performance of certified graph unlearning methods on different datasets. We set $\alpha = 10$, $\lambda = 10^{-4}$ for Computers and $\lambda = 10^{-4}$ for ogbn-arxiv. The number of repeated trails is 3 due to large amount of removed data. (g), (h) Tradeoff between privacy ϵ and performance. To match the number of retraining cycles, we set $\alpha\epsilon = 0.1$.

when designing unlearning methods for node classification tasks. (2) Our method supports unlearning a large proportion of data points with a small loss in test accuracy. (3) Our method is around $4\times$ faster than completely retraining the model after each unlearning request. (4) Our methods have robust performance regardless of the scale of the datasets (see Appendix 7.18).

Trade-off amongst privacy, performance and time complexity. As indicated in Theorem 4.1, there exists a trade-off amongst privacy, performance and time complexity parameters. Compared to exact unlearning (i.e. SGC retraining), approximate unlearning offers $4\times$ speedups in time complexity with competitive performance (see Figure 3). Figure 4 (c)-(d) shows the need for certified graph unlearning compared to unstructured unlearning, where we achieve higher accuracy with comparable or lower unlearning time. Note that it is possible to have a lower unlearning time for SGC unlearning compared to unstructured unlearning, as the latter may require more frequent retraining updates

due to reaching the preset privacy budget. We further examine this trade-off by fixing λ and δ , in which case the trade-off is controlled by ϵ and α . The results are shown in Figure 4 (g) and (h) for Cora and Citeseer respectively, where we set $\alpha\epsilon = 0.1$. The test accuracy increases when we relax our constraints on ϵ , which agrees with our intuition. Remarkably, we can still obtain competitive performance with SGC Retraining when we require ϵ to be as small as 1. In contrast, one needs at least $\epsilon \geq 5$ to unlearn even *one* node or edge by leveraging state-of-the-art DP-GNNs [24, 21] for reasonable performance, albeit our tested datasets are different. This shows the benefit of our certified graph unlearning method as opposed to both retraining from scratch and DP-GNNs. The code for DP-GNNs is not publicly available, which prevents us from testing them on our datasets in a unified treatment.

7 Conclusion

We introduced the first known framework for certified graph unlearning. In this setting, new analytical and heuristic unlearning challenges had to be addressed due to the presence of complex graph feature and topology data. Our analytical contributions pertain to novel proof techniques for certified graph unlearning, while our empirical studies on six benchmark datasets established fundamental performance-complexity trade-offs between unlearning and complete retraining.

Acknowledgments and Disclosure of Funding

This work was funded by NSF grants 1816913 and 1956384.

References

- [1] B. Thomee, D. A. Shamma, G. Friedland, B. Elizalde, K. Ni, D. Poland, D. Borth, and L.-J. Li, “Yfcc100m: The new data in multimedia research,” *Communications of the ACM*, vol. 59, no. 2, pp. 64–73, 2016.
- [2] C. Guo, T. Goldstein, A. Hannun, and L. Van Der Maaten, “Certified data removal from machine learning models,” in *International Conference on Machine Learning*. PMLR, 2020, pp. 3832–3842.
- [3] A. L. Maas, R. E. Daly, P. T. Pham, D. Huang, A. Y. Ng, and C. Potts, “Learning word vectors for sentiment analysis,” in *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*. Portland, Oregon, USA: Association for Computational Linguistics, June 2011, pp. 142–150. [Online]. Available: <http://www.aclweb.org/anthology/P11-1015>
- [4] C. Sudlow, J. Gallacher, N. Allen, V. Beral, P. Burton, J. Danesh, P. Downey, P. Elliott, J. Green, M. Landray *et al.*, “Uk biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age,” *PLoS medicine*, vol. 12, no. 3, p. e1001779, 2015.
- [5] A. Ginart, M. Guan, G. Valiant, and J. Y. Zou, “Making ai forget you: Data deletion in machine learning,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [6] A. Sekhari, J. Acharya, G. Kamath, and A. T. Suresh, “Remember what you want to forget: Algorithms for machine unlearning,” *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [7] Y. Cao and J. Yang, “Towards making systems forget with machine unlearning,” in *2015 IEEE Symposium on Security and Privacy*. IEEE, 2015, pp. 463–480.
- [8] L. Bourtole, V. Chandrasekaran, C. A. Choquette-Choo, H. Jia, A. Travers, B. Zhang, D. Lie, and N. Papernot, “Machine unlearning,” in *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2021, pp. 141–159.
- [9] A. Darrow-Pinion, J. She, D. Wong, O. Lange, T. Hester, L. Perez, M. Nunkesser, S. Lee, X. Guo, B. Wiltshire *et al.*, “Eta prediction with graph neural networks in google maps,” in *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 2021, pp. 3767–3776.

- [10] R. Ying, R. He, K. Chen, P. Eksombatchai, W. L. Hamilton, and J. Leskovec, “Graph convolutional neural networks for web-scale recommender systems,” in *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 2018, pp. 974–983.
- [11] J. Gao, C. Sun, H. Zhao, Y. Shen, D. Anguelov, C. Li, and C. Schmid, “Vectornet: Encoding hd maps and agent dynamics from vectorized representation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 525–11 533.
- [12] X.-M. Zhang, L. Liang, L. Liu, and M.-J. Tang, “Graph neural networks and their current applications in bioinformatics,” *Frontiers in Genetics*, vol. 12, 2021.
- [13] F. Wu, A. Souza, T. Zhang, C. Fifty, T. Yu, and K. Weinberger, “Simplifying graph convolutional networks,” in *International conference on machine learning*. PMLR, 2019, pp. 6861–6871.
- [14] C. Dwork, “Differential privacy. encyclopedia of cryptography and security,” 2011.
- [15] A. Gohatkar, A. Achille, and S. Soatto, “Eternal sunshine of the spotless net: Selective forgetting in deep networks,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 9304–9312.
- [16] M. Chen, Z. Zhang, T. Wang, M. Backes, M. Humbert, and Y. Zhang, “Graph unlearning,” *arXiv preprint arXiv:2103.14991*, 2021.
- [17] E. Chien, J. Peng, P. Li, and O. Milenkovic, “Adaptive universal generalized pagerank graph neural network,” in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=n6jl7fLxrP>
- [18] D. Lim, X. Li, F. Hohne, and S.-N. Lim, “New benchmarks for learning on non-homophilous graphs,” *arXiv preprint arXiv:2104.01404*, 2021.
- [19] K. Chaudhuri, C. Monteleoni, and A. D. Sarwate, “Differentially private empirical risk minimization.” *Journal of Machine Learning Research*, vol. 12, no. 3, 2011.
- [20] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, “Deep learning with differential privacy,” in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016, pp. 308–318.
- [21] A. Daigavane, G. Madan, A. Sinha, A. G. Thakurta, G. Aggarwal, and P. Jain, “Node-level differentially private graph neural networks,” *arXiv preprint arXiv:2111.15521*, 2021.
- [22] I. E. Olatunji, T. Funke, and M. Khosla, “Releasing graph neural networks with differential privacy guarantees,” *arXiv preprint arXiv:2109.08907*, 2021.
- [23] F. Wu, Y. Long, C. Zhang, and B. Li, “Linkteller: Recovering private edges from graph neural networks via influence analysis,” *arXiv preprint arXiv:2108.06504*, 2021.
- [24] S. Sajadmanesh, A. S. Shamsabadi, A. Bellet, and D. Gatica-Perez, “Gap: Differentially private graph neural networks with aggregation perturbation,” *arXiv preprint arXiv:2203.00949*, 2022.
- [25] P. Li, I. Chien, and O. Milenkovic, “Optimizing generalized pagerank methods for seed-expansion community detection,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [26] G. Jeh and J. Widom, “Scaling personalized web search,” in *Proceedings of the 12th international conference on World Wide Web*, 2003, pp. 271–279.
- [27] F. Chung, “The heat kernel as the pagerank of a graph,” *Proceedings of the National Academy of Sciences*, vol. 104, no. 50, pp. 19 735–19 740, 2007.
- [28] F. Frasca, E. Rossi, D. Eynard, B. Chamberlain, M. Bronstein, and F. Monti, “Sign: Scalable inception graph neural networks,” *arXiv preprint arXiv:2004.11198*, 2020.
- [29] P. Sen, G. Namata, M. Bilgic, L. Getoor, B. Galligher, and T. Eliassi-Rad, “Collective classification in network data,” *AI magazine*, vol. 29, no. 3, pp. 93–93, 2008.
- [30] Z. Yang, W. Cohen, and R. Salakhudinov, “Revisiting semi-supervised learning with graph embeddings,” in *International conference on machine learning*. PMLR, 2016, pp. 40–48.
- [31] M. Fey and J. E. Lenssen, “Fast graph representation learning with pytorch geometric,” *arXiv preprint arXiv:1903.02428*, 2019.
- [32] W. Hu, M. Fey, M. Zitnik, Y. Dong, H. Ren, B. Liu, M. Catasta, and J. Leskovec, “Open graph benchmark: Datasets for machine learning on graphs,” *Advances in neural information processing systems*, vol. 33, pp. 22 118–22 133, 2020.

- [33] J. McAuley, C. Targett, Q. Shi, and A. Van Den Hengel, “Image-based recommendations on styles and substitutes,” in *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, 2015, pp. 43–52.
- [34] O. Shchur, M. Mumme, A. Bojchevski, and S. Günnemann, “Pitfalls of graph neural network evaluation,” *Relational Representation Learning Workshop, NeurIPS 2018*, 2018.
- [35] E. Chien, W.-C. Chang, C.-J. Hsieh, H.-F. Yu, J. Zhang, O. Milenkovic, and I. S. Dhillon, “Node feature extraction by self-supervised multi-scale neighborhood prediction,” in *International Conference on Learning Representations*, 2022. [Online]. Available: <https://openreview.net/forum?id=KJggliHbs8>
- [36] F. Gama, A. Ribeiro, and J. Bruna, “Stability of graph scattering transforms,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [37] C. Pan, S. Chen, and A. Ortega, “Spatio-temporal graph scattering transform,” in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=CF-ZluSMXRz>
- [38] J. Ma, B. Chang, X. Zhang, and Q. Mei, “Copulagnn: Towards integrating representational and correlational roles of graphs in graph neural networks,” in *International Conference on Learning Representations*, 2020.
- [39] J. Jia and A. R. Benson, “Residual correlation in graph neural network regression,” in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020, pp. 588–598.

Appendix

7.1 Limitations and future research directions

Batch unlearning. In practice, it is likely that we not only require sequential unlearning, but also batch unlearning: A number of users may request their data to be unlearned within a certain (short) time frame. The approach in [2] can ensure certified removal even in this scenario. The generalization of our approach for batch unlearning is also possible, but will be discussed elsewhere.

Nonlinear models. Also akin to what was described in [2], we can leverage pre-trained (nonlinear) feature extractors or special graph feature transforms to further improve the performance of the overall model. For example, Chien et al. [35] proposed a node feature extraction method termed GIANT-XRT that greatly improves the performance of simple network models such as MLP and SGC. If a public dataset is never subjected to unlearning, one can pre-train GIANT-XRT on that dataset and use it for subsequent certified graph unlearning. If such a public dataset is unavailable, we have to make the node feature extractor DP. In this case, we can either design a DP version of GIANT-XRT or leverage the DP-GNN model described in Section 2. By applying Theorem 5 of [2], the overall model can be shown to guarantee certified graph unlearning, where the parameters ϵ and δ now also depend on the DP guarantees of the node feature extractor. There is also another line of work on Graph Scattering Transforms (GSTs) [36, 37] for use as feature extractors for graph information. Since a GST is a predefined mathematical transform and hence does not require training, it can be easily combined with our approach. The rigorous analysis is delegated to future work.

Societal impacts. The authors believe that for medical and biological sciences, the right to be forgotten may significantly set back potentially life-saving discoveries due to the need to have access to many diverse data samples. But current trends seem to favor privacy over discovery rates and timings. Hence, a compromise between data availability and the right to be forgotten has to be established in the near future.

One current limitation of our work is that the newly proposed proof-techniques do not apply to general graph neural networks where nonlinear activation functions are used. Nevertheless, our work is the first step towards developing certified graph unlearning approaches for general GNNs.

7.2 Intuition behind the model update rule

Our unlearning mechanism proposed in Section 4 is

$$\mathbf{w}^- = \mathbf{w}^* + [\nabla^2 L(\mathbf{w}^*, \mathcal{D}')]^{-1} [\nabla L(\mathbf{w}^*, \mathcal{D}) - \nabla L(\mathbf{w}^*, \mathcal{D}')],$$

and the intuition behind the approach is as follows. Our goal is to have $\nabla L(\mathbf{w}^-, \mathcal{D}') = 0$ for the updated model. Using Taylor series expansion we have

$$\nabla L(\mathbf{w}^-, \mathcal{D}') \approx \nabla L(\mathbf{w}^*, \mathcal{D}') + \nabla^2 L(\mathbf{w}^*, \mathcal{D}')(\mathbf{w}^- - \mathbf{w}^*) = 0.$$

Therefore,

$$\begin{aligned} \mathbf{w}^- - \mathbf{w}^* &= [\nabla^2 L(\mathbf{w}^*, \mathcal{D}')]^{-1} [0 - \nabla L(\mathbf{w}^*, \mathcal{D}')] \\ \mathbf{w}^- &= \mathbf{w}^* + [\nabla^2 L(\mathbf{w}^*, \mathcal{D}')]^{-1} [\nabla L(\mathbf{w}^*, \mathcal{D}) - \nabla L(\mathbf{w}^*, \mathcal{D}')]. \end{aligned}$$

The last equality holds due to the fact that \mathbf{w}^* should be the unique optimizer for the strongly convex loss $L(\mathbf{w}, \mathcal{D})$ over the entire dataset \mathcal{D} .

7.3 Additional discussions

Details regarding Assumption 4.2. Assumptions (2), (4) and (5) in our model and that of [2] require Lipschitz conditions with respect to the first argument of ℓ , but not the second. We also implicitly assume that the second argument (corresponding to labels) does not effect the norm of gradients or Hessians. One example that meets these constraints is the logistic loss: If $\ell(\mathbf{w}^T \mathbf{x}, y) = \ell(y \mathbf{w}^T \mathbf{x})$ then all required assumptions hold.

Least-squares and logistic regression on graphs. Paralleling once again the results of [2], it is clear that our certified graph unlearning mechanism can be used in conjunction with least-squares and logistic regressions. For example, node classification can be performed using a logistic loss. The node regression problem described in [38, 39] is related to least-squares regression. In particular, least-squares regression uses the loss $\ell(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}, \mathbf{e}_i^T \mathbf{Y}_{T_r}) = (\mathbf{e}_i^T \mathbf{Z} \mathbf{w} - \mathbf{e}_i^T \mathbf{Y}_{T_r})^2$. Note that its Hessian is of the form $(\mathbf{e}_i^T \mathbf{Z})^T \mathbf{e}_i^T \mathbf{Z}$, which does not depend on \mathbf{w} . Thus, based on the same arguments presented in [2], our proposed unlearning method M offers $(0, 0)$ -certified graph unlearning even without loss perturbations.

For binary logistic regression, the loss equals $\ell(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}, \mathbf{e}_i^T \mathbf{Y}_{T_r}) = -\log(\sigma(\mathbf{e}_i^T \mathbf{Y}_{T_r} \mathbf{e}_i^T \mathbf{Z} \mathbf{w}))$, where $\sigma(x) = 1/(1 + \exp(-x))$ denotes the sigmoid function. As shown in [2], the assumptions (1)-(3) in 4.2 are satisfied with $c = 1$ and $\gamma_2 = 1/4$. We only need to show that (4) and (5) of 4.2 hold as well. Observe that $\ell'(x, \mathbf{e}_i^T \mathbf{Y}_{T_r}) = (\sigma(\mathbf{e}_i^T \mathbf{Y}_{T_r} x) - 1)$. Since the sigmoid function $\sigma(\cdot)$ is restricted to lie in $[0, 1]$, $|\ell'|$ is bounded by 1, which means that our loss satisfies (5) in 4.2 with $c_1 = 1$. Based on the Mean Value Theorem, one can show that $\sigma(x)$ is $\max_{x \in \mathbb{R}} |\sigma(x)'|$ -Lipschitz. Using some simple algebra, one can also prove that $\sigma(x)' = \sigma(x)(1 - \sigma(x)) \Rightarrow \max_{x \in \mathbb{R}} |\sigma(x)'| = 1/4$. Thus our loss satisfies assumption (4) in 4.2 as well, with $\gamma_1 = 1/4$. For multi-class logistic regression, one can adapt the ‘‘one-versus-all other-classes’’ strategy which leads to the same result.

7.4 Algorithmic details

The pseudo-codes for training removal-enabled models and the removal procedure for the case of binary classification are presented below. Note that this procedure is the same for all three types of removal requests (node feature unlearning, edge unlearning and node unlearning). During training, we add a random linear term to the training loss by sampling a Gaussian noise vector \mathbf{b} . The choice of standard deviation α is determined by the $\alpha \epsilon / \sqrt{2 \log(1.5/\delta)}$, as described in Section 5.

Algorithm 1 Training procedure

- 1: **Input:** Training data $\mathbf{Z} \in \mathbb{R}^{m \times d}$, training labels $\mathbf{Y} \in \mathbb{R}^m$, loss ℓ , parameters $\alpha, \lambda > 0$.
 - 2: Sample the noise vector $\mathbf{b} \sim \mathcal{N}(\mathbf{0}, \alpha^2 \mathbf{I}_d)$.
 - 3: $\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathbb{R}^d} \sum_{i=1}^m (\ell(\mathbf{z}_i^T \mathbf{w}, y_i) + \frac{\lambda}{2} \|\mathbf{w}\|^2) + \mathbf{b}^T \mathbf{w}$.
 - 4: **return** \mathbf{w}^* .
-

7.5 Proof of Theorem 4.3

Theorem. Under the node feature unlearning scenario, $\mathcal{D} = (\mathbf{X}, \mathbf{Y}_{T_r}, \mathbf{A})$ and $\mathcal{D}' = (\mathbf{X}', \mathbf{Y}'_{T_r}, \mathbf{A})$. Suppose Assumption 4.2 holds. For $\mathbf{Z} = \mathbf{P}^K \mathbf{X}$ and $\mathbf{P} = \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}}$, we have

$$\|\nabla L(\mathbf{w}^-, \mathcal{D}')\| = \|(\mathbf{H}_{\mathbf{w}_\eta} - \mathbf{H}_{\mathbf{w}^*}) \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta\| \leq \frac{\gamma_2 (2c\lambda + (c\gamma_1 + \lambda c_1) \tilde{\mathbf{D}}_{mm})^2}{\lambda^4 (m-1)}. \quad (7)$$

Algorithm 2 Unlearning procedure

- 1: **Input:** Feature matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$, labels $\mathbf{Y} \in \mathbb{R}^n$, one-step propagation matrix \mathbf{P} , loss ℓ , training set indices $T_r = \{i_1, i_2, \dots\}$, sequence of removal requests $R_m = \{j_1, j_2, \dots\}$, parameters $K, \epsilon, \delta, \gamma_2, \alpha, \lambda > 0$.
 - 2: Compute node embedding after propagation $\mathbf{Z} = \mathbf{P}^K \mathbf{X}$.
 - 3: Training set $\mathcal{D} = \{\mathbf{z}_i, y_i\}_{i \in T_r}$.
 - 4: Compute \mathbf{w} using Algorithm 1 ($\mathcal{D}, \ell, \alpha, \lambda$).
 - 5: Accumulated gradient residual norm $\beta = 0$.
 - 6: **for** $j \in R_m$ **do**
 - 7: Update the feature matrix \mathbf{X}' and propagation matrix \mathbf{P}' based on the removed point.
 - 8: Compute new node embedding after propagation $\mathbf{Z}' = \mathbf{P}'^K \mathbf{X}'$.
 - 9: **if** $j \in T_r$ **then**
 - 10: Remove j from the training indices $T_r = T_r \setminus \{j\}$.
 - 11: **end if**
 - 12: Update the training set $\mathcal{D}' = \{\mathbf{z}'_i, y_i\}_{i \in T_r}$.
 - 13: Compute $\Delta = \nabla L(\mathbf{w}, \mathcal{D}) - \nabla L(\mathbf{w}, \mathcal{D}')$.
 - 14: Compute $\mathbf{H} = \nabla^2 L(\mathbf{w}, \mathcal{D}')$.
 - 15: Update the accumulated gradient residual norm $\beta = \beta + \gamma_2 \|\mathbf{Z}'\|_{op} \|\mathbf{H}^{-1} \Delta\| \|\mathbf{Z}' \mathbf{H}^{-1} \Delta\|$.
 - 16: **if** $\beta > \alpha \epsilon / \sqrt{2 \log(1.5/\delta)}$ **then**
 - 17: Recompute \mathbf{w} using Algorithm 1 ($\mathcal{D}', \ell, \alpha, \lambda$), $\beta = 0$.
 - 18: **else**
 - 19: $\mathbf{w} = \mathbf{w} + \mathbf{H}^{-1} \Delta$.
 - 20: **end if**
 - 21: **end for**
 - 22: **return** \mathbf{w} .
-

We need to ensure that the norm of each row of \mathbf{Z} is bounded by 1. We state the following lemma in support of this claim.

Lemma 7.1. *Assume that $\|\mathbf{e}_i^T \mathbf{S}\| \leq 1, \forall i \in [n]$. Then, $\forall i \in [n], K \geq 0, \|\mathbf{e}_i^T \mathbf{P}^K \mathbf{S}\| \leq 1$, where $\mathbf{P} = \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}}$.*

Proof. Our proof is a nontrivial generalization and extension of the proof in [2]. For completeness, we outline every step of the proof. We also emphasize novel approaches used to accommodate our graph certified unlearning scenario.

Let $G(\mathbf{w}) = \nabla L(\mathbf{w}, \mathcal{D}')$. By the Taylor theorem, $\exists \eta \in [0, 1]$ such that

$$\begin{aligned}
 G(\mathbf{w}^-) &= G(\mathbf{w}^* + \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta) = G(\mathbf{w}^*) + \nabla G(\mathbf{w}^* + \eta \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta) \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta \\
 &\stackrel{(a)}{=} G(\mathbf{w}^*) + \mathbf{H}_{\mathbf{w}_\eta} \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta \\
 &= G(\mathbf{w}^*) + \Delta + \mathbf{H}_{\mathbf{w}_\eta} \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta - \Delta \\
 &\stackrel{(b)}{=} 0 + \mathbf{H}_{\mathbf{w}_\eta} \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta - \Delta \\
 &= \mathbf{H}_{\mathbf{w}_\eta} \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta - \mathbf{H}_{\mathbf{w}^*} \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta \\
 &= (\mathbf{H}_{\mathbf{w}_\eta} - \mathbf{H}_{\mathbf{w}^*}) \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta.
 \end{aligned} \tag{8}$$

In (a), we wrote $\mathbf{H}_{\mathbf{w}_\eta} \triangleq \nabla^2 G(\mathbf{w}^* + \eta \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta)$, corresponding to the Hessian at $\mathbf{w}_\eta \triangleq \mathbf{w}^* + \eta \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta$. Equality (b) is due to our choice of $\Delta = \nabla L(\mathbf{w}^*, \mathcal{D}) - \nabla L(\mathbf{w}^*, \mathcal{D}')$ and the fact that \mathbf{w}^* is the minimizer of $L(\cdot, \mathcal{D})$. We would like to point out that our choice of Δ is more general than that [2]: Since unlearning one node may affect the entire node embedding \mathbf{Z} , a generalization of Δ is crucial. When $K = 0$ (i.e., when no graph topology is included), one recovers Δ from [2] as a special case of our model. In the latter part of the proof, we will see how the graph setting makes the analysis more intricate and complex.

By the Cauchy-Schwartz inequality, we have

$$\|G(\mathbf{w}^-)\| \leq \|\mathbf{H}_{\mathbf{w}_\eta} - \mathbf{H}_{\mathbf{w}^*}\| \|\mathbf{H}_{\mathbf{w}^*}^{-1} \Delta\|. \tag{9}$$

Below we bound both norms on the right hand side separately. We start with the term $\|\mathbf{H}_{\mathbf{w}_\eta} - \mathbf{H}_{\mathbf{w}^*}\|$. Note that

$$\begin{aligned}
& \|\nabla^2 \ell(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}_\eta, \mathbf{e}_i^T \mathbf{Y}'_{T_r}) - \nabla^2 \ell(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}'_{T_r})\| \\
&= \|\ell''(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}_\eta, \mathbf{e}_i^T \mathbf{Y}'_{T_r}) - \ell''(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}'_{T_r})\| (\mathbf{e}_i^T \mathbf{Z}')^T \mathbf{e}_i^T \mathbf{Z}' \\
&\stackrel{(a)}{\leq} \gamma_2 \|\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}_\eta - \mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*\| \|\mathbf{e}_i^T \mathbf{Z}'\|^2 \\
&\leq \gamma_2 \|\mathbf{w}_\eta - \mathbf{w}^*\| \|\mathbf{e}_i^T \mathbf{Z}'\|^3 = \gamma_2 \|\eta \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta\| \|\mathbf{e}_i^T \mathbf{Z}'\|^3 \leq \gamma_2 \|\mathbf{H}_{\mathbf{w}^*}^{-1} \Delta\| \|\mathbf{e}_i^T \mathbf{Z}'\|^3. \tag{10}
\end{aligned}$$

Here, (a) follows from the Cauchy-Schwartz inequality and the Lipschitz condition on ℓ'' in Assumption 4.2. Unlike the analysis in [2], we are faced with the problem of bounding the term $\|\mathbf{e}_i^T \mathbf{Z}'\|$. In [2] (where $\mathbf{Z} = \mathbf{X}$), a simple bound equals 1, which may be obtained via (3) in Assumption 4.2. However, in our case, due to graph propagation this norm needs more careful examination and a simple application of the Cauchy-Schwartz inequality does not suffice, as it would lead to a term $\|\mathbf{X}\|_{op}$, where $\|\cdot\|_{op}$ denotes the operator norm. The simple worst case (i.e., when all rows of \mathbf{X} are identical) leads to a meaningless bound $O(n)$.

By leveraging Lemma 7.1, we can further upper bound (10) according to

$$\begin{aligned}
& \|\nabla^2 \ell(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}_\eta, \mathbf{e}_i^T \mathbf{Y}'_{T_r}) - \nabla^2 \ell(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}'_{T_r})\| \leq \gamma_2 \|\mathbf{H}_{\mathbf{w}^*}^{-1} \Delta\| \|\mathbf{e}_i^T \mathbf{Z}'\|^3 \\
&\stackrel{(a)}{\leq} \gamma_2 \|\mathbf{H}_{\mathbf{w}^*}^{-1} \Delta\|, \tag{11}
\end{aligned}$$

where (a) follows from Lemma 7.1.

As a result, we arrive at a bound for $\|\mathbf{H}_{\mathbf{w}_\eta} - \mathbf{H}_{\mathbf{w}^*}\|$ of the form

$$\begin{aligned}
\|\mathbf{H}_{\mathbf{w}_\eta} - \mathbf{H}_{\mathbf{w}^*}\| &\leq \sum_{i=1}^{m-1} \|\nabla^2 \ell(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}_\eta, \mathbf{e}_i^T \mathbf{Y}'_{T_r}) - \nabla^2 \ell(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}'_{T_r})\| \\
&\leq \gamma_2 (m-1) \|\mathbf{H}_{\mathbf{w}^*}^{-1} \Delta\|. \tag{12}
\end{aligned}$$

Next, we bound $\|\mathbf{H}_{\mathbf{w}^*}^{-1} \Delta\|$. Since $L(\cdot, \mathcal{D}')$ is $\lambda(m-1)$ -strongly convex, we have $\|\mathbf{H}_{\mathbf{w}^*}^{-1}\| \leq \frac{1}{\lambda(m-1)}$. For the norm $\|\Delta\|$, we have

$$\begin{aligned}
\Delta &= \nabla L(\mathbf{w}^*, \mathcal{D}) - \nabla L(\mathbf{w}^*, \mathcal{D}') \\
&= \lambda \mathbf{w}^* + \nabla \ell(\mathbf{e}_m^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_m^T \mathbf{Y}_{T_r}) + \sum_{i=1}^{m-1} [\nabla \ell(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r}) - \nabla \ell(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}'_{T_r})]. \tag{13}
\end{aligned}$$

The third term does not appear in [2], since when $K = 0$, $\mathbf{Z} = \mathbf{X}$ and $\mathbf{Z}' = \mathbf{X}'$ are identical except for the m^{th} row. In the graph certified unlearning scenario, even removing one node feature can make the entire node embedding matrix \mathbf{Z} change in every row, which creates new analytical challenges. For example, consider the case $\mathbf{X} = [x_1, x_2, x_3]^T$, where we have a graph with three nodes, each with a 1-dimensional feature. Consider the fully connected graph (i.e., all entries in \mathbf{P} set to $1/3$). Then, unlearning node 1 results in $\mathbf{Z}' = [0, x_2, x_3]^T$ for unstructured unlearning. However, $\mathbf{Z}' = [(x_2 + x_3)/3, (x_2 + x_3)/3, (x_2 + x_3)/3]^T$ for the case of $L = 1$, which is completely different from $\mathbf{Z} = [(x_1 + x_2 + x_3)/3, (x_1 + x_2 + x_3)/3, (x_1 + x_2 + x_3)/3]^T$. Hence, the analysis in [2] cannot be directly applied to graphs, as \mathbf{Z}' changes in more than just one row compared to \mathbf{Z} while unlearning a node feature.

By Minkowski's triangle inequality, we only need to bound the norm of the three individual terms in order to bound the norm of Δ . For $\|\mathbf{w}^*\|$, since \mathbf{w}^* is the global optimum of $L(\cdot, \mathcal{D})$, we have

$$0 = \nabla L(\mathbf{w}^*, \mathcal{D}) = \sum_{i=1}^m \nabla \ell(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r}) + \lambda m \mathbf{w}^*. \tag{14}$$

By (1) in Assumption 4.2, we have

$$\|\mathbf{w}^*\| = \frac{\|\sum_{i=1}^m \nabla \ell(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})\|}{\lambda m} \leq \frac{c}{\lambda}. \tag{15}$$

Once again, by (1) in Assumption 4.2, we have

$$\|\nabla\ell(\mathbf{e}_m^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_m^T \mathbf{Y}_{T_r})\| \leq c. \quad (16)$$

A bound for the last term is established in the last step, as described below.

$$\begin{aligned} & \left\| \sum_{i=1}^{m-1} [\nabla\ell(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r}) - \nabla\ell(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})] \right\| \\ & \leq \sum_{i=1}^{m-1} \|\nabla\ell(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r}) - \nabla\ell(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})\| \\ & = \sum_{i=1}^{m-1} \|\ell'(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})(\mathbf{e}_i^T \mathbf{Z})^T - \ell'(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})(\mathbf{e}_i^T \mathbf{Z}')^T\|. \end{aligned} \quad (17)$$

Observe that

$$\begin{aligned} & \|\ell'(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})(\mathbf{e}_i^T \mathbf{Z})^T - \ell'(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})(\mathbf{e}_i^T \mathbf{Z}')^T\| \\ & \leq \|\ell'(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})(\mathbf{e}_i^T \mathbf{Z})^T - \ell'(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})(\mathbf{e}_i^T \mathbf{Z})^T\| \\ & \quad + \|\ell'(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})(\mathbf{e}_i^T \mathbf{Z})^T - \ell'(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})(\mathbf{e}_i^T \mathbf{Z}')^T\| \end{aligned} \quad (18)$$

The first term can be bounded as

$$\begin{aligned} & \|\ell'(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})(\mathbf{e}_i^T \mathbf{Z})^T - \ell'(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})(\mathbf{e}_i^T \mathbf{Z})^T\| \\ & \leq |\ell'(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r}) - \ell'(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})| \|(\mathbf{e}_i^T \mathbf{Z})^T\| \\ & \stackrel{(a)}{\leq} \gamma_1 \|\mathbf{e}_i^T \mathbf{Z} \mathbf{w}^* - \mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*\| \|(\mathbf{e}_i^T \mathbf{Z})^T\| \\ & \stackrel{(b)}{\leq} \gamma_1 \|(\mathbf{e}_i^T \mathbf{Z} - \mathbf{e}_i^T \mathbf{Z}')^T\| \|\mathbf{w}^*\| \\ & \stackrel{(c)}{\leq} \frac{c\gamma_1}{\lambda} \|(\mathbf{e}_i^T \mathbf{Z} - \mathbf{e}_i^T \mathbf{Z}')^T\|. \end{aligned} \quad (19)$$

Here, (a) is due to (4) in Assumption 4.2, while (b) follows from Lemma 7.1 and the Cauchy-Schwartz inequality. Inequality (c) is a consequence of the bound for $\|\mathbf{w}\|$ that we previously derived.

The second term can be bounded as

$$\begin{aligned} & \|\ell'(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})(\mathbf{e}_i^T \mathbf{Z})^T - \ell'(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})(\mathbf{e}_i^T \mathbf{Z}')^T\| \\ & \leq |\ell'(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})| \|(\mathbf{e}_i^T \mathbf{Z})^T - (\mathbf{e}_i^T \mathbf{Z}')^T\| \\ & \stackrel{(a)}{\leq} c_1 \|(\mathbf{e}_i^T \mathbf{Z})^T - (\mathbf{e}_i^T \mathbf{Z}')^T\|. \end{aligned} \quad (20)$$

For the inequality in (a), we used (5) from Assumption 4.2. Put together, we have

$$\begin{aligned}
& \left\| \sum_{i=1}^{m-1} [\nabla \ell(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r}) - \nabla \ell(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})] \right\| \\
& \leq \sum_{i=1}^{m-1} \left[\left(\frac{c\gamma_1}{\lambda} + c_1 \right) \|(\mathbf{e}_i^T \mathbf{Z})^T - (\mathbf{e}_i^T \mathbf{Z}')^T\| \right] \\
& = \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \sum_{i=1}^{m-1} \|\mathbf{e}_i^T (\mathbf{Z} - \mathbf{Z}')\| \\
& = \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \sum_{i=1}^{m-1} \|(\mathbf{e}_i^T \mathbf{P}^K (\mathbf{X} - \mathbf{X}')\| \\
& = \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \sum_{i=1}^{m-1} \|(\mathbf{e}_i^T \mathbf{P}^K \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{D}} (\mathbf{X} - \mathbf{X}')\| \\
& \stackrel{(a)}{=} \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \sum_{i=1}^{m-1} \|(\mathbf{e}_i^T \mathbf{P}^K \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{D}} \mathbf{e}_m \mathbf{e}_m^T \mathbf{X})\| \\
& \stackrel{(b)}{\leq} \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \sum_{i=1}^{m-1} \|\mathbf{e}_i^T \mathbf{P}^K \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{D}} \mathbf{e}_m\| \|\mathbf{e}_m^T \mathbf{X}\| \\
& \stackrel{(c)}{\leq} \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \sum_{i=1}^{m-1} \|\mathbf{e}_i^T \mathbf{P}^K \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{D}} \mathbf{e}_m\| \\
& = \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \sum_{i=1}^{m-1} \|\mathbf{e}_i^T \mathbf{P}^K \tilde{\mathbf{D}}^{-1} \mathbf{e}_m \tilde{\mathbf{D}}_{mm}\| \\
& \stackrel{(d)}{=} \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \sum_{i=1}^{m-1} \mathbf{e}_i^T \mathbf{P}^K \tilde{\mathbf{D}}^{-1} \mathbf{e}_m \tilde{\mathbf{D}}_{mm} \\
& \stackrel{(e)}{\leq} \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \mathbf{1}^T \mathbf{P}^K \tilde{\mathbf{D}}^{-1} \mathbf{e}_m \tilde{\mathbf{D}}_{mm} \\
& \stackrel{(f)}{=} \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \mathbf{1}^T \tilde{\mathbf{D}}^{-1} \mathbf{p} \tilde{\mathbf{D}}_{mm} \\
& \stackrel{(g)}{\leq} \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \tilde{\mathbf{D}}_{mm}. \tag{21}
\end{aligned}$$

Inequality (a) follows from the fact that \mathbf{X}' is identical to \mathbf{X} except for the last row and column, which are set to all-zeros. Thus, $\mathbf{X} - \mathbf{X}'$ is a matrix with rows equal to zero-vectors, except for the m^{th} which equals the m^{th} row of \mathbf{X} . Inequality (b) follows from the Cauchy-Schwartz inequality. Inequality (c) is a result of (3) in Assumption 4.2, while (d) is a consequence of the fact that $\mathbf{e}_i^T \mathbf{P}^K \tilde{\mathbf{D}}^{-1} \mathbf{e}_m$ is the value in the i^{th} row and m^{th} column of the matrix $\mathbf{P}^K \tilde{\mathbf{D}}^{-1}$. Also, it is obvious that this matrix is entry-wise nonnegative. Inequality (e) is due to the fact that $\mathbf{P}^K \tilde{\mathbf{D}}^{-1}$ is entry-wise nonnegative. In (f), \mathbf{p} stands for a probability vector and (f) holds since

$$\mathbf{P}^K \tilde{\mathbf{D}}^{-1} = \left(\tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}} \right)^K \tilde{\mathbf{D}}^{-1} = \tilde{\mathbf{D}}^{-1} (\tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-1})^K, \tag{22}$$

and $\tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-1}$ is a left stochastic matrix. Inequality (g) is a consequence of the observation that the maximum entry in $\tilde{\mathbf{D}}^{-1}$ is at most 1 and that the latter is a diagonal matrix. Hence, $\mathbf{1}^T \tilde{\mathbf{D}}^{-1} \mathbf{p} \leq \mathbf{1}^T \mathbf{p}$. Also, $\mathbf{1}^T \mathbf{p} = 1$ by the definition of the probability vector.

Combining the bounds, we obtain

$$\|\Delta\| \leq c + c + \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \tilde{\mathbf{D}}_{mm} = \frac{2c\lambda + (c\gamma_1 + \lambda c_1) \tilde{\mathbf{D}}_{mm}}{\lambda}. \tag{23}$$

Including the bound on $\|\mathbf{H}_{\mathbf{w}^*}^{-1}\|$ and (12), we then obtain

$$\begin{aligned} \|G(\mathbf{w}^-)\| &\leq \gamma_2(m-1)\|\mathbf{H}_{\mathbf{w}^*}^{-1}\Delta\|^2 \leq \gamma_2(m-1)\left(\frac{2c\lambda+(c\gamma_1+\lambda c_1)\tilde{\mathbf{D}}_{mm}}{\lambda(m-1)}\right)^2 \\ &= \frac{\gamma_2(2c\lambda+(c\gamma_1+\lambda c_1)\tilde{\mathbf{D}}_{mm})^2}{\lambda^4(m-1)}. \end{aligned} \quad (24)$$

This completes the proof. \square

7.6 Proof of Theorem 4.4

Theorem. *For the edge unlearning case, we have $\mathcal{D} = (\mathbf{X}, \mathbf{Y}_{T_r}, \mathbf{P})$ and $\mathcal{D}' = (\mathbf{X}, \mathbf{Y}_{T_r}, \mathbf{P}')$. If $\mathbf{P} = \tilde{\mathbf{D}}^{-1}\tilde{\mathbf{A}}$ and $\mathbf{Z} = \mathbf{P}^K\mathbf{X}$, then we have*

$$\|\nabla L(\mathbf{w}^-, \mathcal{D}')\| = \|(\mathbf{H}_{\mathbf{w}_\eta} - \mathbf{H}_{\mathbf{w}^*})\mathbf{H}_{\mathbf{w}^*}^{-1}\Delta\| \leq \frac{16\gamma_2 K^2 (c\gamma_1 + c_1\lambda)^2}{\lambda^4 m}. \quad (25)$$

Similar to what holds for the node feature unlearning case, Theorem 4.4 still holds when neither of the two end nodes of the removed edge belongs to T_r . Since P' is a right stochastic matrix, Lemma 7.1 still applies. Thus, we only need to describe how to bound $\|\Delta\|$. Following an approach similar to the previously described one, we have $\|\Delta\| \leq (\frac{c\gamma_1}{\lambda} + c_1) \sum_{i=1}^m \sum_{j=1}^n \|\mathbf{e}_i^T(\mathbf{P}^K - \mathbf{P}'^K)\mathbf{e}_j\|$. We also need the following technical lemmas.

Lemma 7.2. *For both edge and node unlearning, we have $|\mathbf{e}_i^T[\mathbf{P}^K - (\mathbf{P}')^K]\mathbf{e}_j| \leq \sum_{k=1}^K \mathbf{e}_i^T(\mathbf{P}')^{k-1}|\mathbf{P} - \mathbf{P}'|\mathbf{P}^{K-k}\mathbf{e}_j, \forall i, j \in [n], K \geq 1$.*

Lemma 7.3. *For edge unlearning, we have $\mathbf{1}^T\mathbf{P}'^{k-1}|\mathbf{P} - \mathbf{P}'|\mathbf{P}^{K-k}\mathbf{1} \leq 4, \forall k \in [K]$.*

Combining the two lemmas and after some algebraic manipulation, we arrive at the desired result. It is not hard to see that $|\mathbf{P} - \mathbf{P}'|$ has only two nonzero rows, which correspond to the unlearned edge. One can again construct a left stochastic matrix $\tilde{\mathbf{A}}'\tilde{\mathbf{D}}'^{-1}$ and a right stochastic matrix \mathbf{P} which lead to the result of Lemma 7.3.

Proof. The theorem can be proved as follows. From the previous proof we have

$$\|G(\mathbf{w}^-)\| \leq \|\mathbf{H}_{\mathbf{w}_\eta} - \mathbf{H}_{\mathbf{w}^*}\| \|\mathbf{H}_{\mathbf{w}^*}^{-1}\| \|\Delta\| \leq \gamma_2 \frac{\|\Delta\|^2}{\lambda^2 m}. \quad (26)$$

Since the first term $\|\mathbf{H}_{\mathbf{w}_\eta} - \mathbf{H}_{\mathbf{w}^*}\|$ only involved the updated dataset, the upper bound for this term proved for node feature unlearning still holds. The term $\|\mathbf{H}_{\mathbf{w}^*}^{-1}\|$ can again be bounded using the fact that $L(\cdot, \mathcal{D}')$ is λm -strongly convex. The main difference between node feature and edge unlearning

lies in the bound for Δ . By definition,

$$\begin{aligned}
\Delta &= \nabla L(\mathbf{w}^*, \mathcal{D}) - \nabla L(\mathbf{w}^*, \mathcal{D}') \\
&= \sum_{i=1}^m [\nabla \ell(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r}) - \nabla \ell(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})], \text{ and} \\
\|\Delta\| &\leq \left(\frac{c\gamma_1}{\lambda} + c_1\right) \sum_{i=1}^m \|(\mathbf{Z} - \mathbf{Z}')^T \mathbf{e}_i\| \\
&= \left(\frac{c\gamma_1}{\lambda} + c_1\right) \sum_{i=1}^m \|(\mathbf{P}^K \mathbf{X} - \mathbf{P}'^K \mathbf{X})^T \mathbf{e}_i\| \\
&= \left(\frac{c\gamma_1}{\lambda} + c_1\right) \sum_{i=1}^m \|\mathbf{e}_i^T (\mathbf{P}^K - \mathbf{P}'^K) \mathbf{X}\| \\
&= \left(\frac{c\gamma_1}{\lambda} + c_1\right) \sum_{i=1}^m \|\mathbf{e}_i^T (\mathbf{P}^K - \mathbf{P}'^K) \sum_{j=1}^n \mathbf{e}_j \mathbf{e}_j^T \mathbf{X}\| \\
&\leq \left(\frac{c\gamma_1}{\lambda} + c_1\right) \sum_{i=1}^m \sum_{j=1}^n \|\mathbf{e}_i^T (\mathbf{P}^K - \mathbf{P}'^K) \mathbf{e}_j \mathbf{e}_j^T \mathbf{X}\| \\
&\leq \left(\frac{c\gamma_1}{\lambda} + c_1\right) \sum_{i=1}^m \sum_{j=1}^n \|\mathbf{e}_i^T (\mathbf{P}^K - \mathbf{P}'^K) \mathbf{e}_j\| \|\mathbf{e}_j^T \mathbf{X}\| \\
&\leq \left(\frac{c\gamma_1}{\lambda} + c_1\right) \sum_{i=1}^m \sum_{j=1}^n \|\mathbf{e}_i^T (\mathbf{P}^K - \mathbf{P}'^K) \mathbf{e}_j\| \tag{27}
\end{aligned}$$

By Lemma 7.2 we have

$$\begin{aligned}
&\left(\frac{c\gamma_1}{\lambda} + c_1\right) \sum_{i=1}^m \sum_{j=1}^n \|\mathbf{e}_i^T (\mathbf{P}^K - \mathbf{P}'^K) \mathbf{e}_j\| \\
&\leq \left(\frac{c\gamma_1}{\lambda} + c_1\right) \sum_{i=1}^m \sum_{j=1}^n \sum_{k=1}^K \mathbf{e}_i^T \mathbf{P}'^{k-1} |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{e}_j \\
&\leq \left(\frac{c\gamma_1}{\lambda} + c_1\right) \sum_{k=1}^K \mathbf{1}^T \mathbf{P}'^{k-1} |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1}. \tag{28}
\end{aligned}$$

Using Lemma 7.3 we arrive at $\|\Delta\| \leq \left(\frac{c\gamma_1}{\lambda} + c_1\right) 4K$. Plugging this expression into (26) completes the proof. \square

7.7 Proof of Theorem 4.5

Theorem. *Under the node unlearning scenario, we have $\mathcal{D} = (\mathbf{X}, \mathbf{Y}_{T_r}, \mathbf{P})$ and $\mathcal{D}' = (\mathbf{X}', \mathbf{Y}'_{T_r}, \mathbf{P}')$. Suppose also that Assumption 4.2 holds. For $\mathbf{Z} = \mathbf{P}^K \mathbf{X}$ and $\mathbf{P} = \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}}$, we have*

$$\|\nabla L(\mathbf{w}^-, \mathcal{D}')\| = \|(\mathbf{H}_{\mathbf{w}_\eta} - \mathbf{H}_{\mathbf{w}^*}) \mathbf{H}_{\mathbf{w}^*}^{-1} \Delta\| \leq \frac{\gamma_2 \left(2c\lambda + K(c\gamma_1 + c_1\lambda) \left(2\tilde{\mathbf{D}}_{mm} - 1\right)\right)^2}{\lambda^4(m-1)}. \tag{29}$$

Again, the main challenge is to bound Δ . First we observe that $(\mathbf{P}')^K \mathbf{X}' = (\mathbf{P}')^K \mathbf{X}$. This holds because node m is removed from the graph in \mathcal{D}' , and thus its corresponding node features do not affect \mathbf{Z}' . Similarly to the proof Lemma 7.2, we first derive the bound $\sum_{k=1}^K \mathbf{1}^T (\mathbf{P}')^{k-1} |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1}$. For each term, $\mathbf{1}^T (\mathbf{P}')^{k-1} |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} = \sum_{l=1}^n \mathbf{1}^T (\mathbf{P}')^{k-1} \mathbf{e}_l \mathbf{e}_l^T |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1}$. To proceed, we need the following two lemmas.

Lemma 7.4. *For node unlearning and $\forall k \in [K]$ and $\forall l \in [n]$, $\mathbf{1}^T (\mathbf{P}')^{k-1} (\tilde{\mathbf{D}}')^{-1} \mathbf{e}_l \leq 1$.*

Lemma 7.5. For node unlearning and $\forall k \in [K]$, $\sum_{l=1}^n \mathbf{e}_l^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} \leq 2\tilde{\mathbf{D}}_{mm} - 1$.

These two lemmas give rise to the term $K(2\tilde{\mathbf{D}}_{mm} - 1)$ in the bound of Theorem 4.5 and the rest of the analysis is similar to that of the previous cases. Lemma 7.5 is rather technical, and relies on the following proposition that exploits the structure of $|\mathbf{P} - \mathbf{P}'|$.

Proposition 7.6. For node unlearning and $\forall i, j \neq m$, $\mathbf{e}_i^T |\mathbf{P} - \mathbf{P}'| \mathbf{e}_j = \mathbf{e}_i^T (\mathbf{P}' - \mathbf{P}) \mathbf{e}_j$. For $i = m$ or $j = m$, $\mathbf{e}_i^T |\mathbf{P} - \mathbf{P}'| \mathbf{e}_j = \mathbf{e}_i^T \mathbf{P} \mathbf{e}_j$.

Proof. The proof is similar to the proof of Theorem 4.3, although several parts need modifications. First, the result of Lemma 7.1 needs to be replaced by the following claim.

Lemma 7.7. Assume that $\|\mathbf{e}_i^T \mathbf{S}\| \leq 1$, $\forall i \neq m$ and that $\mathbf{e}_m^T \mathbf{S} = \mathbf{0}^T$. Then $\forall i \in [n]$, $K \geq 0$, we have $\|\mathbf{e}_i^T (\mathbf{P}')^K \mathbf{S}\| \leq 1$, where $\mathbf{P} = \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}}$ and $\mathbf{P}' = (\tilde{\mathbf{D}}')^{-1} \tilde{\mathbf{A}}'$.

Next we have to modify the proof regarding the bound of $\|\Delta\|$. Following a proof similar to that of Theorem 4.3, we have

$$\|\Delta\| \leq 2c + \left(\frac{c\gamma_1}{\lambda} + c_1\right) \sum_{i=1}^{m-1} \|(\mathbf{Z} - \mathbf{Z}')^T \mathbf{e}_i\|. \quad (30)$$

Plugging in the expressions for \mathbf{Z} and \mathbf{Z}' leads to

$$\begin{aligned} \sum_{i=1}^{m-1} \|(\mathbf{Z} - \mathbf{Z}')^T \mathbf{e}_i\| &= \sum_{i=1}^{m-1} \|(\mathbf{P}^K \mathbf{X} - (\mathbf{P}')^K \mathbf{X}')^T \mathbf{e}_i\| \\ &\stackrel{(a)}{=} \sum_{i=1}^{m-1} \|(\mathbf{P}^K \mathbf{X} - (\mathbf{P}')^K \mathbf{X})^T \mathbf{e}_i\| = \sum_{i=1}^{m-1} \|([\mathbf{P}^K - (\mathbf{P}')^K] \mathbf{X})^T \mathbf{e}_i\| \\ &\stackrel{(b)}{=} \sum_{i=1}^{m-1} \|\mathbf{e}_i^T [\mathbf{P}^K - (\mathbf{P}')^K] \sum_{j=1}^n \mathbf{e}_j \mathbf{e}_j^T \mathbf{X}\| \\ &\stackrel{(c)}{\leq} \sum_{i=1}^{m-1} \sum_{j=1}^n \|\mathbf{e}_i^T [\mathbf{P}^K - (\mathbf{P}')^K] \mathbf{e}_j \mathbf{e}_j^T \mathbf{X}\| \\ &\stackrel{(d)}{\leq} \sum_{i=1}^{m-1} \sum_{j=1}^n \|\mathbf{e}_i^T [\mathbf{P}^K - (\mathbf{P}')^K] \mathbf{e}_j\| \|\mathbf{e}_j^T \mathbf{X}\| \\ &\stackrel{(e)}{\leq} \sum_{i=1}^{m-1} \sum_{j=1}^n \|\mathbf{e}_i^T [\mathbf{P}^K - (\mathbf{P}')^K] \mathbf{e}_j\|. \end{aligned} \quad (31)$$

The equality (a) is due to the fact that $(\mathbf{P}')^K \mathbf{X}' = (\mathbf{P}')^K \mathbf{X}$, as the m^{th} row and column of $(\mathbf{P}')^K$ are all-zeros. Thus, changing the last row of \mathbf{X}' makes no difference of $(\mathbf{P}')^K \mathbf{X}'$. Equation (b) is a consequence of the fact that $\mathbf{I} = \sum_{j=1}^n \mathbf{e}_j \mathbf{e}_j^T$. Inequality (c) follows from Minkowski's inequality, while (d) follows from the Cauchy-Schwartz inequality. Inequality (e) holds based on (3) in Assumption 4.2.

By Lemma 7.2, we can proceed with our analysis as follows:

$$\begin{aligned}
& \sum_{i=1}^{m-1} \|(\mathbf{Z} - \mathbf{Z}')^T \mathbf{e}_i\| \\
& \leq \sum_{i=1}^{m-1} \sum_{j=1}^n \|\mathbf{e}_i^T [\mathbf{P}^K - (\mathbf{P}')^K] \mathbf{e}_j\| \\
& \stackrel{(a)}{\leq} \sum_{i=1}^{m-1} \sum_{j=1}^n \sum_{k=1}^K \mathbf{e}_i^T (\mathbf{P}')^{k-1} |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{e}_j \\
& \leq \sum_{k=1}^K \mathbf{1}^T (\mathbf{P}')^{k-1} |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1}, \tag{32}
\end{aligned}$$

where (a) is due to Lemma 7.2 and the fact that $\mathbf{e}_i^T \mathbf{P} \mathbf{e}_j$ is a scalar, equal to the i^{th} row j^{th} column of the matrix \mathbf{P} .

Next, we bound each term $\mathbf{1}^T (\mathbf{P}')^{k-1} |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1}$ separately. For $k \in [K]$, we have

$$\begin{aligned}
& \mathbf{1}^T (\mathbf{P}')^{k-1} |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} \\
& = \mathbf{1}^T (\mathbf{P}')^{k-1} (\tilde{\mathbf{D}}')^{-1} \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} \\
& = \mathbf{1}^T (\mathbf{P}')^{k-1} (\tilde{\mathbf{D}}')^{-1} \sum_{l=1}^n \mathbf{e}_l \mathbf{e}_l^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} \\
& = \sum_{l=1}^n \left(\mathbf{1}^T (\mathbf{P}')^{k-1} (\tilde{\mathbf{D}}')^{-1} \mathbf{e}_l \right) \left(\mathbf{e}_l^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} \right). \tag{33}
\end{aligned}$$

Note that for each index l , the corresponding term in the sum is just a product of two scalars. Let first analyze $\mathbf{1}^T (\mathbf{P}')^{k-1} (\tilde{\mathbf{D}}')^{-1} \mathbf{e}_l$. This term can be bounded as

$$\begin{aligned}
& \mathbf{1}^T (\mathbf{P}')^{k-1} |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} \\
& = \sum_{j=1}^n \left(\mathbf{1}^T (\mathbf{P}')^{k-1} (\tilde{\mathbf{D}}')^{-1} \mathbf{e}_j \right) \left(\mathbf{e}_j^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} \right) \\
& \stackrel{(a)}{\leq} \sum_{j=1}^n \mathbf{e}_j^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1}. \tag{34}
\end{aligned}$$

where (a) follows from Lemma 7.4.

We now turn our attention to the term $\mathbf{e}_l^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1}$, which can be bounded by Lemma 7.5 as

$$\begin{aligned}
& \mathbf{e}_l^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} \\
& \leq \sum_{l=1}^n \mathbf{e}_l^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} \\
& \leq 2\tilde{\mathbf{D}}_{mm} - 1. \tag{35}
\end{aligned}$$

Using these two bounds in (32) gives

$$\begin{aligned}
& \sum_{i=1}^{m-1} \|(\mathbf{Z} - \mathbf{Z}')^T \mathbf{e}_i\| \\
& \leq \sum_{i=1}^{m-1} \sum_{j=1}^n \|\mathbf{e}_i^T [\mathbf{P}^K - (\mathbf{P}')^K] \mathbf{e}_j\| \\
& \leq \sum_{k=1}^K \mathbf{1}^T (\mathbf{P}')^{k-1} |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} \\
& \leq \sum_{k=1}^K (2\tilde{\mathbf{D}}_{mm} + 1) = K(2\tilde{\mathbf{D}}_{mm} - 1). \tag{36}
\end{aligned}$$

Using this bound in the expression for $\|\Delta\|$ we obtain

$$\begin{aligned}
\|\Delta\| & \leq 2c + \left(\frac{c\gamma_1}{\lambda} + c_1\right) \sum_{i=1}^{m-1} \|(\mathbf{Z} - \mathbf{Z}')^T \mathbf{e}_i\| \\
& \leq 2c + \left(\frac{c\gamma_1}{\lambda} + c_1\right) K (2\tilde{\mathbf{D}}_{mm} - 1) \\
& \Rightarrow \|G(\mathbf{w}^-)\| \leq \|\mathbf{H}_{\mathbf{w}_\eta} - \mathbf{H}_{\mathbf{w}^*}\| \|\mathbf{H}_{\mathbf{w}^*}^{-1} \Delta\| \\
& \leq \gamma_2(m-1) \|\mathbf{H}_{\mathbf{w}^*}^{-1} \Delta\|^2 \\
& \leq \gamma_2(m-1) \left(\frac{2c + \left(\frac{c\gamma_1}{\lambda} + c_1\right) K (2\tilde{\mathbf{D}}_{mm} - 1)}{\lambda(m-1)} \right)^2 \\
& = \frac{\gamma_2 \left(2c\lambda + K(c\gamma_1 + c_1\lambda) (2\tilde{\mathbf{D}}_{mm} - 1)\right)^2}{\lambda^4(m-1)}. \tag{37}
\end{aligned}$$

This completes the proof. \square

7.8 Proof of Theorem 4.6

Theorem. *In the node feature unlearning scenario, we are given $\mathcal{D} = (\mathbf{X}, \mathbf{Y}_{T_r}, \mathbf{A})$ and $\mathcal{D}' = (\mathbf{X}', \mathbf{Y}'_{T_r}, \mathbf{A})$. Suppose that Assumption 4.2 holds. For $\mathbf{Z} = \frac{1}{K+1} [\mathbf{X}, \mathbf{P}\mathbf{X}, \mathbf{P}^2\mathbf{X}, \dots, \mathbf{P}^K\mathbf{X}]$ and $\mathbf{P} = \tilde{\mathbf{D}}^{-1}\tilde{\mathbf{A}}$, we have*

$$\|\nabla L(\mathbf{w}^-, \mathcal{D}')\| = \|(\mathbf{H}_{\mathbf{w}_\eta} - \mathbf{H}_{\mathbf{w}^*})\mathbf{H}_{\mathbf{w}^*}^{-1}\Delta\| \leq \frac{\gamma_2(2c\lambda + (c\gamma_1 + \lambda c_1)\tilde{\mathbf{D}}_{mm})^2}{\lambda^4(m-1)}. \tag{38}$$

Proof. The proof is almost identical to the proof of Theorem 4.3. We only need to bound the norms of the terms in \mathbf{Z} . We start by modifying Lemma 7.1 for the GPR case.

Lemma 7.8. *Assume that $\|\mathbf{e}_i^T \mathbf{S}\| \leq 1, \forall i \in [n]$. Then $\forall i \in [n], K \geq 0$, we have $\|\frac{1}{\sqrt{K+1}}\mathbf{e}_i^T [\mathbf{S}, \mathbf{P}\mathbf{S}, \mathbf{P}^2\mathbf{S}, \dots, \mathbf{P}^K\mathbf{S}]\| \leq 1$, where $\mathbf{P} = \tilde{\mathbf{D}}^{-1}\tilde{\mathbf{A}}$.*

Another part of the proof that needs to be changed is to establish a bound on

$$\sum_{i=1}^{m-1} [\nabla \ell(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r}) - \nabla \ell(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}'_{T_r})]. \tag{39}$$

Following a proof similar to that of Theorem 4.3, we have

$$\begin{aligned}
& \left\| \sum_{i=1}^{m-1} [\nabla \ell(\mathbf{e}_i^T \mathbf{Z} \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r}) - \nabla \ell(\mathbf{e}_i^T \mathbf{Z}' \mathbf{w}^*, \mathbf{e}_i^T \mathbf{Y}_{T_r})] \right\| \\
& \leq \sum_{i=1}^{m-1} \left[\left(\frac{c\gamma_1}{\lambda} + c_1 \right) \|(\mathbf{e}_i^T \mathbf{Z})^T - (\mathbf{e}_i^T \mathbf{Z}')^T\| \right] \\
& = \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \sum_{i=1}^{m-1} \|(\mathbf{Z} - \mathbf{Z}')^T \mathbf{e}_i\| \\
& = \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \sum_{i=1}^{m-1} \left\| \left(\frac{1}{K+1} [\mathbf{X} - \mathbf{X}', \mathbf{P}(\mathbf{X} - \mathbf{X}'), \dots, \mathbf{P}^K(\mathbf{X} - \mathbf{X}')] \right)^T \mathbf{e}_i \right\| \\
& = \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \sum_{i=1}^{m-1} \left\| \left(\frac{1}{K+1} [\mathbf{e}_m \mathbf{e}_m^T \mathbf{X}, \mathbf{P} \mathbf{e}_m \mathbf{e}_m^T \mathbf{X}, \dots, \mathbf{P}^K \mathbf{e}_m \mathbf{e}_m^T \mathbf{X}] \right)^T \mathbf{e}_i \right\| \\
& = \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \sum_{i=1}^{m-1} \left\| \frac{1}{K+1} [\mathbf{e}_i^T \mathbf{e}_m \mathbf{e}_m^T \mathbf{X}, \mathbf{e}_i^T \mathbf{P} \mathbf{e}_m \mathbf{e}_m^T \mathbf{X}, \dots, \mathbf{e}_i^T \mathbf{P}^K \mathbf{e}_m \mathbf{e}_m^T \mathbf{X}]^T \right\| \\
& \leq \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \sum_{i=1}^{m-1} \left\| \frac{1}{K+1} [\mathbf{e}_i^T \mathbf{e}_m, \mathbf{e}_i^T \mathbf{P} \mathbf{e}_m, \dots, \mathbf{e}_i^T \mathbf{P}^K \mathbf{e}_m]^T \right\| \|(\mathbf{e}_m^T \mathbf{X})^T\| \\
& \leq \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \sum_{i=1}^{m-1} \left\| \frac{1}{K+1} [\mathbf{e}_i^T \mathbf{e}_m, \mathbf{e}_i^T \mathbf{P} \mathbf{e}_m, \dots, \mathbf{e}_i^T \mathbf{P}^K \mathbf{e}_m]^T \right\| \\
& \stackrel{(a)}{\leq} \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \sum_{i=1}^{m-1} \frac{1}{K+1} \sum_{k=1}^K \mathbf{e}_i^T \mathbf{P}^k \mathbf{e}_m \\
& \leq \frac{c\gamma_1}{\lambda} + c_1 \sum_{k=1}^K \mathbf{1}^T \mathbf{P}^k \mathbf{e}_m = \frac{c\gamma_1}{\lambda} + c_1 \sum_{k=1}^K \mathbf{1}^T \mathbf{P}^k \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{D}} \mathbf{e}_m = \frac{c\gamma_1}{\lambda} + c_1 \sum_{k=1}^K \mathbf{1}^T \mathbf{P}^k \tilde{\mathbf{D}}^{-1} \mathbf{e}_m \tilde{\mathbf{D}}_{mm} \\
& \stackrel{(b)}{=} \frac{c\gamma_1}{\lambda} + c_1 \sum_{k=1}^K \mathbf{1}^T \tilde{\mathbf{D}}^{-1} \mathbf{p}^{(k)} \tilde{\mathbf{D}}_{mm} \\
& \leq \frac{c\gamma_1}{\lambda} + c_1 \sum_{k=1}^K \mathbf{1}^T \mathbf{p}^{(k)} \tilde{\mathbf{D}}_{mm} = \frac{c\gamma_1}{\lambda} + c_1 \times K \tilde{\mathbf{D}}_{mm} \\
& \leq \left(\frac{c\gamma_1}{\lambda} + c_1 \right) \tilde{\mathbf{D}}_{mm}, \tag{40}
\end{aligned}$$

where (a) is due to the fact that the ℓ_1 norm is an upper bound for the ℓ_2 norm. Also note that $\mathbf{e}_i^T \mathbf{e}_m = 0, \forall i \neq m$. In (b), $\forall k \in [K]$, $\mathbf{p}^{(k)}$ are probability vectors. This completes the proof. \square

Remark. Note that the GPR extension for the edge and node unlearning cases can be derived through a similar analysis. One can also see that the key step is inequality (a), which still holds for the edge and node unlearning cases. The results are similar to Theorem 4.4 and Theorem 4.5, except that the definition of \mathbf{Z} is replaced by one corresponding to the GPR case, as in Theorem 4.6.

7.9 Proof of Lemma 7.1

Lemma. Assume that $\|\mathbf{e}_i^T \mathbf{S}\| \leq 1, \forall i \in [n]$. Then $\forall i \in [n], K \geq 0$, we have $\|\mathbf{e}_i^T \mathbf{P}^K \mathbf{S}\| \leq 1$, where $\mathbf{P} = \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}}$.

Proof. We prove this lemma by induction. Let $\mathbf{Z}^{(k)} = \mathbf{P}^k \mathbf{S}$. For the base case $k = 0$ it is true by assumption that $\|\mathbf{e}_i^T \mathbf{S}\| \leq 1 \forall i \in [n]$. Assume next that the claim is true for the case $k = K - 1$.

Then we have

$$\begin{aligned} \|\mathbf{e}_i^T \mathbf{P}^K \mathbf{S}\| &= \|\mathbf{e}_i^T \mathbf{P} \mathbf{Z}^{(K-1)}\| = \left\| \frac{1}{\tilde{\mathbf{D}}_{ii}} \sum_{j: \tilde{\mathbf{A}}_{ij}=1} \mathbf{e}_j^T \mathbf{Z}^{(K-1)} \right\| \leq \frac{1}{\tilde{\mathbf{D}}_{ii}} \sum_{j: \tilde{\mathbf{A}}_{ij}=1} \|\mathbf{e}_j^T \mathbf{Z}^{(K-1)}\| \\ &\stackrel{(a)}{\leq} \frac{1}{\tilde{\mathbf{D}}_{ii}} \sum_{j: \tilde{\mathbf{A}}_{ij}=1} 1 = \frac{1}{\tilde{\mathbf{D}}_{ii}} \times \tilde{\mathbf{D}}_{ii} = 1, \end{aligned} \quad (41)$$

where (a) is based on the induction hypothesis for $k = K - 1$. \square

Remark: Note that if we choose another propagation matrix \mathbf{P} compared to the one used in the SGC analysis, the above expression for $K = 1$ becomes

$$\begin{aligned} \|\mathbf{e}_i^T \mathbf{P} \mathbf{S}\| &= \left\| \frac{1}{\sqrt{\tilde{\mathbf{D}}_{ii}}} \sum_{j: \tilde{\mathbf{A}}_{ij}=1} \frac{\mathbf{e}_j^T \mathbf{S}}{\sqrt{\tilde{\mathbf{D}}_{jj}}} \right\| \leq \frac{1}{\sqrt{\tilde{\mathbf{D}}_{ii}}} \sum_{j: \tilde{\mathbf{A}}_{ij}=1} \frac{\|\mathbf{e}_j^T \mathbf{S}\|}{\sqrt{\tilde{\mathbf{D}}_{jj}}} \\ &\leq \frac{1}{\sqrt{\tilde{\mathbf{D}}_{ii}}} \sum_{j: \tilde{\mathbf{A}}_{ij}=1} \frac{1}{\sqrt{\tilde{\mathbf{D}}_{jj}}}. \end{aligned} \quad (42)$$

We cannot easily simplify the sum $\sum_{j: \tilde{\mathbf{A}}_{ij}=1} \frac{1}{\sqrt{\tilde{\mathbf{D}}_{jj}}}$. One way to approach the problem is to simply use the fact that the degree of a node is at least 1 and can thus be further upper bounded by $\tilde{\mathbf{D}}_{ii}$. This leads to the bound

$$\|\mathbf{e}_i^T \mathbf{P} \mathbf{S}\| \leq \sqrt{\tilde{\mathbf{D}}_{ii}}. \quad (43)$$

Obviously, this bound is worse than the one in Lemma 7.1 even when $K = 1$. For general K , there will be an additional exponent $K/2$ for the maximal degree, which is undesirable. Nevertheless, our bound is tight since for the worst case of a star graph with a center at node i , so that $\mathbf{D}_{jj} = 2$ for all $j \neq i$. The same argument applies for other degree normalizations. Thus it is critical to choose $\mathbf{P} = \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}}$ to obtain the desired bound in Lemma 7.1.

7.10 Proof of Lemma 7.2

Lemma. For either the edge or node unlearning case, and $\forall i, j \in [n]$, $K \geq 1$, we have

$$|\mathbf{e}_i^T [\mathbf{P}^K - (\mathbf{P}')^K] \mathbf{e}_j| \leq \sum_{k=1}^K \mathbf{e}_i^T (\mathbf{P}')^{k-1} |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{e}_j. \quad (44)$$

Proof. The proof consist of two parts. We first show that

$$\mathbf{P}^K - (\mathbf{P}')^K = \sum_{k=1}^K (\mathbf{P}')^{k-1} (\mathbf{P} - \mathbf{P}') \mathbf{P}^{K-k}.$$

Then we proceed to analyze the absolute values of all terms in the sum.

The proof of the first part follows from a telescoping property for the sum,

$$\begin{aligned} \sum_{k=1}^K (\mathbf{P}')^{k-1} (\mathbf{P} - \mathbf{P}') \mathbf{P}^{K-k} &= \sum_{k=1}^K (\mathbf{P}')^{k-1} \mathbf{P}^{K-k+1} - (\mathbf{P}')^k \mathbf{P}^{K-k} \\ &= (\mathbf{P}')^0 \mathbf{P}^K - (\mathbf{P}')^1 \mathbf{P}^{K-1} + (\mathbf{P}')^1 \mathbf{P}^{K-1} - (\mathbf{P}')^2 \mathbf{P}^{K-2} + \dots + (\mathbf{P}')^{K-1} \mathbf{P}^1 - (\mathbf{P}')^K \mathbf{P}^0 \\ &= \mathbf{P}^K - (\mathbf{P}')^K. \end{aligned} \quad (45)$$

Next, note that both \mathbf{P}' and \mathbf{P} are nonnegative matrices, and the same is true of their k^{th} powers, $k \geq 2$. Thus,

$$\begin{aligned} |\mathbf{e}_i^T [\mathbf{P}^K - (\mathbf{P}')^K] \mathbf{e}_j| &= \left| \sum_{k=1}^K \mathbf{e}_i^T (\mathbf{P}')^{k-1} (\mathbf{P} - \mathbf{P}') \mathbf{P}^{K-k} \mathbf{e}_j \right| \\ &\leq \sum_{k=1}^K \mathbf{e}_i^T (\mathbf{P}')^{k-1} |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{e}_j. \end{aligned} \quad (46)$$

This completes the proof. \square

7.11 Proof of Lemma 7.3

Lemma. For the edge unlearning scenario, and $\forall k \in [K]$, we have

$$\mathbf{1}^T \mathbf{P}'^{k-1} |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} \leq 4. \quad (47)$$

Proof. Let us start by analyzing the matrix $|\mathbf{P} - \mathbf{P}'| = |\tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}} - \tilde{\mathbf{D}}'^{-1} \tilde{\mathbf{A}}'|$. Note that all its rows are zeros except for the 1st and m^{th} row. The first row of the matrix equals

$$\begin{aligned} & \mathbf{e}_1^T |\mathbf{P} - \mathbf{P}'| \\ &= \left[\left(\frac{1}{d_1 - 1} - \frac{1}{d_1} \right) \tilde{\mathbf{A}}_{11}, \dots, \left(\frac{1}{d_1 - 1} - \frac{1}{d_1} \right) \tilde{\mathbf{A}}_{1(m-1)}, \frac{1}{d_1}, \left(\frac{1}{d_1 - 1} - \frac{1}{d_1} \right) \tilde{\mathbf{A}}_{1(m+1)}, \dots \right] \\ &= \left[\left(\frac{1}{d_1(d_1 - 1)} \right) \tilde{\mathbf{A}}_{11}, \dots, \left(\frac{1}{d_1(d_1 - 1)} \right) \tilde{\mathbf{A}}_{1(m-1)}, \frac{1}{d_1}, \left(\frac{1}{d_1(d_1 - 1)} \right) \tilde{\mathbf{A}}_{1(m+1)}, \dots \right] \\ &= \left[\left(\frac{1}{d_1(d_1 - 1)} \right) \tilde{\mathbf{A}}_{11}, \dots, \left(\frac{1}{d_1(d_1 - 1)} \right) \tilde{\mathbf{A}}_{1(m-1)}, \frac{1}{d_1(d_1 - 1)}, \left(\frac{1}{d_1(d_1 - 1)} \right) \tilde{\mathbf{A}}_{1(m+1)}, \dots \right] \\ &+ \frac{d_1 - 2}{d_1(d_1 - 1)} \mathbf{e}_m^T = \mathbf{e}_1^T \tilde{\mathbf{D}}'^{-1} \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}} + \frac{d_1 - 2}{d_1(d_1 - 1)} \mathbf{e}_m^T, \end{aligned} \quad (48)$$

where the last equality holds since $\tilde{\mathbf{A}}_{1m} = 1$. Similar arguments apply for the m^{th} row, for which we have

$$\mathbf{e}_m^T |\mathbf{P} - \mathbf{P}'| = \mathbf{e}_m^T \tilde{\mathbf{D}}'^{-1} \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}} + \frac{d_m - 2}{d_m(d_m - 1)} \mathbf{e}_1^T. \quad (49)$$

For a fixed $k \in [K]$,

$$\begin{aligned} \mathbf{1}^T \mathbf{P}'^{k-1} |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} &= \mathbf{1}^T \mathbf{P}'^{k-1} \mathbf{e}_1 \mathbf{e}_1^T \tilde{\mathbf{D}}'^{-1} \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}} \mathbf{P}^{K-k} \mathbf{1} \\ &+ \mathbf{1}^T \mathbf{P}'^{k-1} \frac{d_1 - 2}{d_1(d_1 - 1)} \mathbf{e}_1 \mathbf{e}_m^T \mathbf{P}^{K-k} \mathbf{1} \\ &+ \mathbf{1}^T \mathbf{P}'^{k-1} \mathbf{e}_m \mathbf{e}_m^T \tilde{\mathbf{D}}'^{-1} \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}} \mathbf{P}^{K-k} \mathbf{1} \\ &+ \mathbf{1}^T \mathbf{P}'^{k-1} \frac{d_m - 2}{d_m(d_m - 1)} \mathbf{e}_m \mathbf{e}_1^T \mathbf{P}^{K-k} \mathbf{1}. \end{aligned} \quad (50)$$

We analyze these four terms separately. For the first term, we have

$$\begin{aligned} & \mathbf{1}^T \mathbf{P}'^{k-1} \mathbf{e}_1 \mathbf{e}_1^T \tilde{\mathbf{D}}'^{-1} \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}} \mathbf{P}^{K-k} \mathbf{1} \\ &= \mathbf{1}^T \mathbf{P}'^{k-1} \tilde{\mathbf{D}}'^{-1} \mathbf{e}_1 \mathbf{e}_1^T \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}} \mathbf{P}^{K-k} \mathbf{1} \\ &= \mathbf{1}^T \mathbf{P}'^{k-1} \tilde{\mathbf{D}}'^{-1} \mathbf{e}_1 \mathbf{e}_1^T \mathbf{P}^{K-k+1} \mathbf{1} \end{aligned} \quad (51)$$

By the same argument as used in the proof for node feature unlearning, $\mathbf{1}^T \mathbf{P}'^{k-1} \tilde{\mathbf{D}}'^{-1} \mathbf{e}_1 = \mathbf{1}^T \tilde{\mathbf{D}}'^{-1} \mathbf{p} \leq 1$, for some probability vector \mathbf{p} . Also, $\mathbf{e}_1^T \mathbf{P}^{K-k+1} \mathbf{1} \leq 1$, which holds due to the fact that \mathbf{P} is a right-stochastic matrix. We have hence shown that the first term in (51) is bounded by 1. For the second term, note that $\frac{d_1 - 2}{d_1(d_1 - 1)} \leq \frac{1}{(d_1 - 1)}$. Hence,

$$\begin{aligned} & \mathbf{1}^T \mathbf{P}'^{k-1} \frac{d_1 - 2}{d_1(d_1 - 1)} \mathbf{e}_1 \mathbf{e}_m^T \mathbf{P}^{K-k} \mathbf{1} \\ &\leq \mathbf{1}^T \mathbf{P}'^{k-1} \frac{1}{d_1 - 1} \mathbf{e}_1 \mathbf{e}_m^T \mathbf{P}^{K-k} \mathbf{1} \\ &= \mathbf{1}^T \mathbf{P}'^{k-1} \tilde{\mathbf{D}}'^{-1} \mathbf{e}_1 \mathbf{e}_m^T \mathbf{P}^{K-k} \mathbf{1} \leq 1, \end{aligned} \quad (52)$$

where the final inequality follows the same argument as the one used for bounding the first term. For the third and fourth term, the analysis is similar to these two cases and both terms can be shown to be bounded by 1. Hence, we have

$$\mathbf{1}^T \mathbf{P}'^{k-1} |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} \leq 4. \quad (53)$$

This completes the proof. \square

7.12 Proof of Lemma 7.4

Lemma. For all $k \in [K]$ and $l \in [n]$,

$$\mathbf{1}^T (\mathbf{P}')^{k-1} (\tilde{\mathbf{D}}')^{-1} \mathbf{e}_l \leq 1. \quad (54)$$

Proof. For $k = 1$, the claim is obviously true for all $l \in [n]$, as the largest entry in $\tilde{\mathbf{D}}^{-1}$ is upper bounded by 1. For $k \geq 2$ and $l \neq m$ we have

$$\begin{aligned} \mathbf{1}^T (\mathbf{P}')^{k-1} (\tilde{\mathbf{D}}')^{-1} \mathbf{e}_l &= \mathbf{1}^T ((\tilde{\mathbf{D}}')^{-1} \tilde{\mathbf{A}}')^{k-1} (\tilde{\mathbf{D}}')^{-1} \mathbf{e}_l \\ &= \mathbf{1}^T (\tilde{\mathbf{D}}')^{-1} (\tilde{\mathbf{A}}' (\tilde{\mathbf{D}}')^{-1})^{k-1} \mathbf{e}_l \\ &\stackrel{(a)}{=} \mathbf{1}^T (\tilde{\mathbf{D}}')^{-1} \mathbf{p} \leq 1. \end{aligned} \quad (55)$$

In (a), \mathbf{p} stands for a probability vector and the result follows since $\tilde{\mathbf{A}}' (\tilde{\mathbf{D}}')^{-1}$ is a left-stochastic matrix if one ignores the node m . For $l = m$, it is easy to see that $\tilde{\mathbf{A}}' \tilde{\mathbf{D}}' \mathbf{e}_m = 0$ by the fact that the m^{th} row and column of $\tilde{\mathbf{A}}'$ are all-zeros. This completes the proof. \square

7.13 Proof of Lemma 7.5

Lemma. For node unlearning, and $\forall k \in [K]$, $\sum_{l=1}^n \mathbf{e}_l^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} \leq 2\tilde{\mathbf{D}}_{mm} - 1$.

Proof. First, note that

$$\sum_{l=1}^n \mathbf{e}_l^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| = \sum_{l=1}^n \mathbf{e}_l^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \sum_{r=1}^n \mathbf{e}_r \mathbf{e}_r^T = \sum_{r=1}^n \sum_{l=1}^n \mathbf{e}_l^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{e}_r \mathbf{e}_r^T. \quad (56)$$

Then for $i, j \neq m$, by Proposition 7.6, we have

$$\begin{aligned} &\mathbf{e}_i^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{e}_r \mathbf{e}_r^T \\ &\stackrel{(a)}{=} \mathbf{e}_i^T \tilde{\mathbf{D}}' (\mathbf{P}' - \mathbf{P}) \mathbf{e}_r \mathbf{e}_r^T \\ &= \mathbf{e}_i^T \left(\tilde{\mathbf{A}}' - \tilde{\mathbf{D}}' \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}} \right) \mathbf{e}_r \mathbf{e}_r^T \\ &= \left(\tilde{\mathbf{A}}'_{lr} - \frac{\tilde{\mathbf{D}}'_{il}}{\tilde{\mathbf{D}}_{il}} \tilde{\mathbf{A}}_{lr} \right) \mathbf{e}_r^T \\ &\stackrel{(b)}{=} \left(\tilde{\mathbf{A}}_{lr} - \frac{\tilde{\mathbf{D}}'_{il}}{\tilde{\mathbf{D}}_{il}} \tilde{\mathbf{A}}_{lr} \right) \mathbf{e}_r^T \end{aligned} \quad (57)$$

We used Proposition 7.6 in (a) since $\mathbf{e}_i^T \tilde{\mathbf{D}}' = \tilde{\mathbf{D}}'_{il} \mathbf{e}_i^T$. The equality (b) is due to the fact that for $i, j \neq m$, $\tilde{\mathbf{A}}'_{lr} = \tilde{\mathbf{A}}_{lr}$. Recall that $\tilde{\mathbf{A}}'$ and $\tilde{\mathbf{A}}$ only differ in the m^{th} row and column.

We consider next the only two possible scenarios, (1) l is a neighbor of m ; (2) l is not a neighbor of m . For (1), we know that $\tilde{\mathbf{D}}'_{il} = \tilde{\mathbf{D}}_{il} - 1 \geq 1$. This leads to

$$\tilde{\mathbf{A}}_{lr} - \frac{\tilde{\mathbf{D}}'_{il}}{\tilde{\mathbf{D}}_{il}} \tilde{\mathbf{A}}_{lr} = \tilde{\mathbf{A}}_{lr} \left(1 - \frac{\tilde{\mathbf{D}}'_{il}}{\tilde{\mathbf{D}}_{il}} \right) = \tilde{\mathbf{A}}_{lr} \left(1 - \frac{\tilde{\mathbf{D}}_{il} - 1}{\tilde{\mathbf{D}}_{il}} \right) = \frac{\tilde{\mathbf{A}}_{lr}}{\tilde{\mathbf{D}}_{il}}. \quad (58)$$

For (2), we know that $\tilde{\mathbf{D}}'_{rr} = \tilde{\mathbf{D}}_{rr}$. Thus, $\tilde{\mathbf{A}}_{lr} - \frac{\tilde{\mathbf{D}}'_{il}}{\tilde{\mathbf{D}}_{il}} \tilde{\mathbf{A}}_{lr} = 0$.

Next, we consider the case $l \neq m, r = m$. Again, by Proposition 7.6, we have

$$\mathbf{e}_i^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{e}_m \mathbf{e}_m^T = \mathbf{e}_i^T \tilde{\mathbf{D}}' \mathbf{P} \mathbf{e}_m \mathbf{e}_m^T = \frac{\tilde{\mathbf{D}}'_{il}}{\tilde{\mathbf{D}}_{il}} \tilde{\mathbf{A}}_{lm} \mathbf{e}_m^T. \quad (59)$$

Now, for case (1), we have $\tilde{\mathbf{A}}_{lm} = 1$ and $\tilde{\mathbf{D}}'_{il} = \tilde{\mathbf{D}}_{il} - 1 \geq 1$. This leads to

$$\mathbf{e}_i^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{e}_m \mathbf{e}_m^T = \frac{\tilde{\mathbf{D}}'_{il}}{\tilde{\mathbf{D}}_{il}} \tilde{\mathbf{A}}_{lm} \mathbf{e}_m^T = \frac{\tilde{\mathbf{D}}_{il} - 1}{\tilde{\mathbf{D}}_{il}} \mathbf{e}_m^T. \quad (60)$$

For case (2), we clearly have $\mathbf{e}_l^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{e}_m \mathbf{e}_m^T = 0$ as $\tilde{\mathbf{A}}_{lm} = 0$.

Hence, for each $j \neq m$ and under the setting in case (1), $\mathbf{e}_j^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'|$ equals the row vector

$$\left[\frac{\tilde{\mathbf{A}}_{l1}}{\tilde{\mathbf{D}}_{ll}}, \frac{\tilde{\mathbf{A}}_{l2}}{\tilde{\mathbf{D}}_{ll}}, \dots, \frac{\tilde{\mathbf{A}}_{l(m-1)}}{\tilde{\mathbf{D}}_{ll}}, 0, \frac{\tilde{\mathbf{A}}_{l(m+1)}}{\tilde{\mathbf{D}}_{ll}}, \dots \right] + \left[0, \dots, 0, \frac{\tilde{\mathbf{D}}_{ll} - 1}{\tilde{\mathbf{D}}_{ll}}, 0, \dots \right], \quad (61)$$

where the m^{th} entry of the first row vector equals 0 and the second row vector is all-zeros except for the m^{th} entry. Note that the first row vector times $\frac{\tilde{\mathbf{D}}_{ll}}{\tilde{\mathbf{D}}_{ll} - 1} > 1$ is a probability vector. Hence, by the property of \mathbf{P}^{K-k} being a right-stochastic matrix, we have

$$\left[\frac{\tilde{\mathbf{A}}_{l1}}{\tilde{\mathbf{D}}_{ll}}, \frac{\tilde{\mathbf{A}}_{l2}}{\tilde{\mathbf{D}}_{ll}}, \dots, \frac{\tilde{\mathbf{A}}_{l(m-1)}}{\tilde{\mathbf{D}}_{ll}}, 0, \frac{\tilde{\mathbf{A}}_{l(m+1)}}{\tilde{\mathbf{D}}_{ll}}, \dots \right] \mathbf{P}^{K-k} \mathbf{1} \leq 1. \quad (62)$$

Since $\frac{\tilde{\mathbf{D}}_{ll} - 1}{\tilde{\mathbf{D}}_{ll}} < 1$, we also have

$$\left[0, \dots, 0, \frac{\tilde{\mathbf{D}}_{jj} - 1}{\tilde{\mathbf{D}}_{jj}}, 0, \dots \right] \mathbf{P}^{K-k} \mathbf{1} \leq 1. \quad (63)$$

Together, this shows that for each $j \neq m$ and for the case (1), one has

$$\mathbf{e}_j^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} \leq 2. \quad (64)$$

For case (2), note that $\mathbf{e}_l^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'|$ is an all-zero row vector. Note also that, excluding self-loops, there are at most $\tilde{\mathbf{D}}_{mm} - 1$ neighbors l of m (case (1)). Thus,

$$\sum_{l \neq m} \mathbf{e}_l^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} \leq 2\tilde{\mathbf{D}}_{mm} - 2. \quad (65)$$

To conclude the proof, we analyze the term $l = m$. For any $i \in [n]$, by Proposition 7.6 we have

$$\begin{aligned} & \mathbf{e}_m^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{e}_r \mathbf{e}_r^T \\ &= \mathbf{e}_m^T \tilde{\mathbf{D}}' \mathbf{P} \mathbf{e}_r \mathbf{e}_r^T = \frac{\tilde{\mathbf{D}}'_{mm}}{\tilde{\mathbf{D}}_{mm}} \tilde{\mathbf{A}}_{mr} \mathbf{e}_r^T \stackrel{(a)}{=} \frac{\tilde{\mathbf{A}}_{mr}}{\tilde{\mathbf{D}}_{mm}} \mathbf{e}_r^T = \mathbf{e}_m^T \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}} \mathbf{e}_r \mathbf{e}_r^T = \mathbf{e}_m^T \mathbf{P} \mathbf{e}_r \mathbf{e}_r^T \end{aligned} \quad (66)$$

where (a) holds by definition, and since $\tilde{\mathbf{D}}'_{mm} = 1$. Thus,

$$\mathbf{e}_m^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} = \mathbf{e}_m^T \mathbf{P} \mathbf{P}^{K-k} \mathbf{1} = \mathbf{p}^T \mathbf{1} = 1, \quad (67)$$

for some probability vector \mathbf{p} . We have hence shown that for any $k \in [K]$,

$$\sum_{j=1}^n \mathbf{e}_j^T \tilde{\mathbf{D}}' |\mathbf{P} - \mathbf{P}'| \mathbf{P}^{K-k} \mathbf{1} \leq 2\tilde{\mathbf{D}}_{mm} - 2 + 1 = 2\tilde{\mathbf{D}}_{mm} - 1. \quad (68)$$

This completes the proof. \square

7.14 Proof of Lemma 7.7

Lemma. Assume that $\|\mathbf{e}_i^T \mathbf{S}\| \leq 1$, $\forall i \neq m$ and that $\mathbf{e}_m^T \mathbf{S} = \mathbf{0}^T$. Then $\forall i \in [n]$, $K \geq 0$, we have $\|\mathbf{e}_i^T (\mathbf{P}')^K \mathbf{S}\| \leq 1$, where $\mathbf{P} = \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{A}}$ and $\mathbf{P}' = (\tilde{\mathbf{D}}')^{-1} \tilde{\mathbf{A}}'$.

Proof. The proof is similar to the proof of Lemma 7.1, and based on induction. The base case $k = 0$ is obviously true by assumption. Now, assume that the claim is true for $k = K - 1$ and let

$\mathbf{Z}^{(K-1)} = (\mathbf{P}')^{(K-1)}\mathbf{S}$. Then, $\forall i \neq m$,

$$\begin{aligned}
\|\mathbf{e}_i^T (\mathbf{P}')^K \mathbf{S}\| &= \|\mathbf{e}_i^T \mathbf{P}' \mathbf{Z}^{(K-1)}\| = \left\| \frac{1}{\tilde{\mathbf{D}}_{ii}} \sum_{j: \tilde{\mathbf{A}}'_{ij}=1} \mathbf{e}_j^T \mathbf{Z}^{(K-1)} \right\| \\
&\leq \frac{1}{\tilde{\mathbf{D}}_{ii}} \sum_{j: \tilde{\mathbf{A}}'_{ij}=1} \|\mathbf{e}_j^T \mathbf{Z}^{(K-1)}\| \\
&\stackrel{(a)}{\leq} \frac{1}{\tilde{\mathbf{D}}_{ii}} \sum_{j: \tilde{\mathbf{A}}'_{ij}=1} 1 \\
&\leq \frac{1}{\tilde{\mathbf{D}}_{ii}} \sum_{j: \tilde{\mathbf{A}}'_{ij}=1} 1 = \frac{1}{\tilde{\mathbf{D}}_{ii}} \tilde{\mathbf{D}}_{ii}' = 1.
\end{aligned} \tag{69}$$

Here, (a) is due to our hypothesis for $k = K - 1$. For $i = m$, note that $\tilde{\mathbf{A}}'_{mj} = 0$, $\forall j \in [n]$. Thus, $\|\mathbf{e}_m^T (\mathbf{P}')^K \mathbf{S}\| = 0 \leq 1$. This completes the proof. \square

7.15 Proof of Lemma 7.8

Lemma. Assume that $\|\mathbf{e}_i^T \mathbf{S}\| \leq 1$, $\forall i \in [n]$. Then, $\forall i \in [n]$, $K \geq 0$, we have $\left\| \frac{1}{\sqrt{K+1}} \mathbf{e}_i^T [\mathbf{S}, \mathbf{P}\mathbf{S}, \mathbf{P}^2\mathbf{S}, \dots, \mathbf{P}^K\mathbf{S}] \right\| \leq 1$, where $\mathbf{P} = \tilde{\mathbf{D}}^{-1}\tilde{\mathbf{A}}$.

Proof. By Lemma 7.1, we have $\|\mathbf{e}_i^T \mathbf{P}^k \mathbf{S}\| \leq 1$, $\forall k \in [K]$. Thus,

$$\left\| \frac{1}{\sqrt{K+1}} \mathbf{e}_i^T [\mathbf{S}, \mathbf{P}\mathbf{S}, \mathbf{P}^2\mathbf{S}, \dots, \mathbf{P}^K\mathbf{S}] \right\|^2 = \frac{1}{K+1} \left(\sum_{k=0}^K \|\mathbf{e}_i^T \mathbf{P}^k \mathbf{S}\|^2 \right) \leq 1, \tag{70}$$

which complete the proof. \square

Remark. Using the normalization $\frac{1}{\sqrt{K+1}}$ also leads to a norm bounded by 1. Hence, the norm of each row of \mathbf{Z} is bounded by 1. We need the normalization $\frac{1}{K+1}$ instead of $\frac{1}{\sqrt{K+1}}$ to accommodate another claim in the proof.

7.16 Proof of Proposition 7.6

Proposition. We have $\mathbf{e}_i^T |\mathbf{P} - \mathbf{P}'| \mathbf{e}_j = \mathbf{e}_i^T (\mathbf{P}' - \mathbf{P}) \mathbf{e}_j$, $\forall i, j \neq m$. For $i = m$ or $j = m$, $\mathbf{e}_i^T |\mathbf{P} - \mathbf{P}'| \mathbf{e}_j = \mathbf{e}_i^T \mathbf{P} \mathbf{e}_j$.

Proof. For the first case when $\forall i, j \neq m$,

$$\mathbf{e}_i^T (\mathbf{P} - \mathbf{P}') \mathbf{e}_j = \frac{\tilde{\mathbf{A}}_{ij}}{\tilde{\mathbf{D}}_{ii}} - \frac{\tilde{\mathbf{A}}'_{ij}}{\tilde{\mathbf{D}}'_{ii}}. \tag{71}$$

Recall that by definition, in this case we have $\tilde{\mathbf{A}}_{ij} = \tilde{\mathbf{A}}'_{ij}$. Now, there are two cases to consider: (1) i is a neighbor of m ; (2) i is not a neighbor of m . For (1), we know that $\tilde{\mathbf{D}}'_{ii} = \tilde{\mathbf{D}}_{ii} - 1 \geq 1$. As a result,

$$\frac{\tilde{\mathbf{A}}_{ij}}{\tilde{\mathbf{D}}_{ii}} - \frac{\tilde{\mathbf{A}}'_{ij}}{\tilde{\mathbf{D}}'_{ii}} = \frac{\tilde{\mathbf{A}}_{ij}}{\tilde{\mathbf{D}}_{ii}} - \frac{\tilde{\mathbf{A}}_{ij}}{\tilde{\mathbf{D}}_{ii} - 1} < 0. \tag{72}$$

This directly implies $\mathbf{e}_i^T |\mathbf{P} - \mathbf{P}'| \mathbf{e}_j = \mathbf{e}_i^T (\mathbf{P}' - \mathbf{P}) \mathbf{e}_j$. For (2), we know that $\tilde{\mathbf{D}}'_{ii} = \tilde{\mathbf{D}}_{ii}$ and thus $\mathbf{e}_i^T |\mathbf{P} - \mathbf{P}'| \mathbf{e}_j = 0 = \mathbf{e}_i^T (\mathbf{P}' - \mathbf{P}) \mathbf{e}_j$. These claims complete the proof for the first part. For the case that $i = m$ or $j = m$, note that since both the m^{th} row and column are all-zeros for \mathbf{P}' , we simply have $\mathbf{e}_i^T |\mathbf{P} - \mathbf{P}'| \mathbf{e}_j = \mathbf{e}_i^T \mathbf{P} \mathbf{e}_j$. Note that in establishing the claim we also used the fact that \mathbf{P} is nonnegative. This completes the proof. \square

7.17 Proof of sequential unlearning

The proof is nearly identical to those presented in [2]. We include it here for completeness. The proof is basically induction with triangle inequality.

Let the original training dataset be $\mathcal{D}^{(0)}$ and let $\mathcal{D}^{(t)}$ be the training dataset after t^{th} unlearning request of any kinds. Suppose the weight after $T \geq 1$ unlearning requests is $\mathbf{w}^{(T)}$ and the gradient residual is $\mathbf{u}_T = \nabla L(\mathbf{w}^{(T)}, \mathcal{D}^{(T)})$. Our goal is to prove that

$$\|\mathbf{u}_T\| \leq T\epsilon', \quad (73)$$

where ϵ' is the gradient residual norm bound for a single unlearning request.

First we note that our update rule of $\mathbf{w}^{(t)}$ is

$$\begin{aligned} \mathbf{w}^{(t+1)} &= \mathbf{w}^{(t)} + \left[\mathbf{H}_{\mathbf{w}^{(t)}}^{(t)} \right]^{-1} \Delta^{(t)}, \\ \text{where } \Delta^{(t)} &= \nabla L(\mathbf{w}^{(t)}, \mathcal{D}^{(t)}) - \nabla L(\mathbf{w}^{(t)}, \mathcal{D}^{(t+1)}) \text{ and } \mathbf{H}_{\mathbf{w}^{(t)}}^{(t)} = \nabla^2 L(\mathbf{w}^{(t)}, \mathcal{D}^{(t+1)}). \end{aligned}$$

Also note that $\mathbf{w}^{(0)} = \operatorname{argmin} L(\mathbf{w}, \mathcal{D}^{(0)})$, which is the unique minimizer of $L(\mathbf{w}, \mathcal{D}^{(0)})$ due to strong convexity.

For base case $T = 1$ it is trivially true based on our definition of ϵ' .

Assume that (73) is true for step T . Consider the modified loss function

$$L^{(T)}(\mathbf{w}, \mathcal{D}) = L(\mathbf{w}, \mathcal{D}) - \mathbf{u}_T^T \mathbf{w}. \quad (74)$$

Then $\mathbf{w}^{(T)}$ is the exact solution of $L^{(T)}(\mathbf{w}, \mathcal{D}^{(T)})$. Now, consider the update of

$$\bar{\mathbf{w}}^{(T+1)} = \mathbf{w}^{(T)} + \left[\nabla^2 L^{(T)}(\mathbf{w}^{(T)}, \mathcal{D}^{(T+1)}) \right]^{-1} \left(\nabla L^{(T)}(\mathbf{w}^{(T)}, \mathcal{D}^{(T)}) - \nabla L^{(T)}(\mathbf{w}^{(T)}, \mathcal{D}^{(T+1)}) \right)$$

Since $\mathbf{w}^{(T)}$ is the exact solution of $L^{(T)}(\mathbf{w}, \mathcal{D}^{(T)})$, the same analysis of gradient residual norm bound we have applies for $\nabla L^{(T)}(\bar{\mathbf{w}}^{(T+1)}, \mathcal{D}^{(T+1)})$, which means

$$\|\nabla L^{(T)}(\bar{\mathbf{w}}^{(T+1)}, \mathcal{D}^{(T+1)})\| \leq \epsilon'. \quad (75)$$

Next, we show that $\bar{\mathbf{w}}^{(T+1)} = \mathbf{w}^{(T+1)}$. By simple algebra, we have

$$\nabla^2 L^{(T)}(\mathbf{w}^{(T)}, \mathcal{D}^{(T+1)}) = \nabla^2 L(\mathbf{w}, \mathcal{D}^{(T+1)}) - \nabla^2 \mathbf{u}_T^T \mathbf{w} \Big|_{\mathbf{w}=\mathbf{w}^{(T)}} = \nabla^2 L(\mathbf{w}^{(T)}, \mathcal{D}^{(T+1)}) = \mathbf{H}_{\mathbf{w}^{(T)}}^{(T)}.$$

On the other hand, it is trivial to see that

$$\nabla L^{(T)}(\mathbf{w}^{(T)}, \mathcal{D}^{(T)}) - \nabla L^{(T)}(\mathbf{w}^{(T)}, \mathcal{D}^{(T+1)}) = \nabla L(\mathbf{w}^{(T)}, \mathcal{D}^{(T)}) - \nabla L(\mathbf{w}^{(T)}, \mathcal{D}^{(T+1)}) = \Delta^{(T)}.$$

Thus we have

$$\begin{aligned} \bar{\mathbf{w}}^{(T+1)} &= \mathbf{w}^{(T)} + \left[\nabla^2 L^{(T)}(\mathbf{w}^{(T)}, \mathcal{D}^{(T+1)}) \right]^{-1} \left(\nabla L^{(T)}(\mathbf{w}^{(T)}, \mathcal{D}^{(T)}) - \nabla L^{(T)}(\mathbf{w}^{(T)}, \mathcal{D}^{(T+1)}) \right) \\ &= \mathbf{w}^{(T)} + \left[\mathbf{H}_{\mathbf{w}^{(T)}}^{(T)} \right]^{-1} \Delta^{(T)} = \mathbf{w}^{(T+1)}. \end{aligned} \quad (76)$$

Then (75) becomes

$$\|\nabla L^{(T)}(\mathbf{w}^{(T+1)}, \mathcal{D}^{(T+1)})\| \leq \epsilon'. \quad (77)$$

Finally, by the definition of $L^{(T)}$ we have

$$\begin{aligned} \nabla L^{(T)}(\mathbf{w}^{(T+1)}, \mathcal{D}^{(T+1)}) &= \nabla \left(L(\mathbf{w}, \mathcal{D}^{(T+1)}) - \mathbf{u}_T^T \mathbf{w} \right) \Big|_{\mathbf{w}=\mathbf{w}^{(T+1)}} \\ &= \nabla L(\mathbf{w}^{(T+1)}, \mathcal{D}^{(T+1)}) - \mathbf{u}_T \\ \Rightarrow \mathbf{u}_{T+1} &= \nabla L(\mathbf{w}^{(T+1)}, \mathcal{D}^{(T+1)}) = \nabla L^{(T)}(\mathbf{w}^{(T+1)}, \mathcal{D}^{(T+1)}) + \mathbf{u}_T. \\ \Rightarrow \|\mathbf{u}_{T+1}\| &\leq \|\nabla L^{(T)}(\mathbf{w}^{(T+1)}, \mathcal{D}^{(T+1)})\| + \|\mathbf{u}_T\| \leq \epsilon' + T\epsilon' = (T+1)\epsilon'. \end{aligned}$$

By induction, we have proved that indeed (73) is true and our sequential unlearning is valid.

Remark. Note that the analysis above holds for both unstructured unlearning and graph unlearning scenarios. Once we have the one step results as in Theorem 4.5, 4.3, 4.4 or 4.6, we have the explicit form of ϵ' in (73). It is also worth pointing out that our proposed update rule works for both cases where the underlying loss is $L^{(T)}$ or L . This is due to the fact that our update rule is invariant to the linear shift of $\mathbf{a}^T \mathbf{w}$ for any vector \mathbf{a} , since the update rule only involves Hessian and gradient difference Δ , both of which are apparently invariant to the linear shift of the form $\mathbf{a}^T \mathbf{w}$. One can also notice that the entire analysis still holds for the noisy training loss $L_{\mathbf{b}}$, as it also just introduces the linear shift $\mathbf{b}^T \mathbf{w}$.

7.18 Additional experimental details

Table 1: Statistics of benchmark datasets.

Name	#nodes	#edges	#features	#classes	train/val/test
Cora	2,708	10,556	1,433	7	1,208/500/1,000
Citeseer	3,327	9,104	3,703	6	1,827/500/1,000
PubMed	19,717	88,648	500	3	18,217/500/1,000
Computers	13,752	491,722	767	10	12,252/500/1,000
Photo	7,650	238,162	745	8	6,150/500/1,000
ogbn-arxiv	169,343	1,166,243	128	40	90,941/29,799/48,603

All our experiments were executed on a Linux machine with 48 cores, 376GB of system memory, and two NVIDIA Tesla P100 GPUs with 12GB of GPU memory each. Information about all datasets can be found in Table 1. The data split is public and obtained from PyTorch Geometric [31]. We used the “full” split option for Cora, Citeseer and Pubmed. Since there is no public split for Computers and Photo, we adopted a similar setting as for the citation networks via random splits (i.e., 500 nodes in the validation set and 1,000 nodes in the test set). The data split for ogbn-arxiv is the public split provided by the Open Graph Benchmark [32].

Dependency on the node degree. We verified our Theorem 4.5 and Theorem 4.6 for node degree dependencies on Photo, Cora, Citeseer and Pubmed. The results are presented in Figure 5.

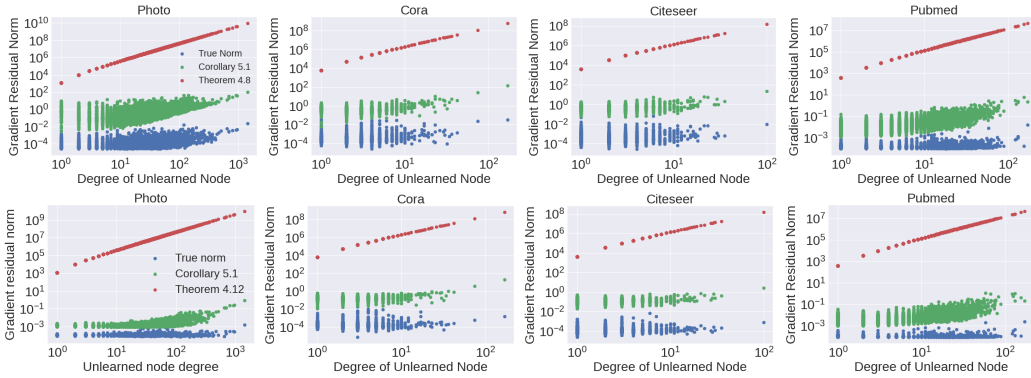


Figure 5: Additional examination of the degree dependency result from Theorem 4.5 (top) and Theorem 4.6 (bottom).

Nonaccumulative time for each of the unlearning procedures. Figure 6 shows the average time complexity for each unlearning step on the Cora dataset. The spikes for certified graph unlearning methods and Algorithm 2 in [2] correspond to retraining after a removal.

Additional experiments. The performance of our proposed certified graph unlearning methods on three datasets, including Citeseer, Pubmed and Amazon Photo, is shown in Figure 7. It is worth pointing out that our bound on the gradient residual norm in Section 4 does not guarantee the generalization ability of the updated model. Therefore, It could happen that the test accuracy




Figure 6: Nonaccumulative time for each removal step on the Cora dataset. The setting is the same as in Figure 3.

increases as we remove information from the training set, as shown in the second row of Figure 7, or that the performance is not very stable, as seen in the third row of Figure 7.

We also performed additional experiments on the Cora dataset, with results shown in Figure 8. The first row shows the average performance over 10 repeated trails with random splitting, and the conclusion is the same as the one stated in Section 6. The second row shows the performance on GPR-based models. Note that when the number of removal requests becomes large, the performance of GPR-based models degrades much faster than that of SGC-based models. This observation is consistent with our discussion of GPR-based models. None of the retraining methods involves noise.




Figure 7: Performance of certified graph unlearning methods on different datasets. First Row: We removed up to 55% of the training data in Citeseer. Second Row: We removed up to 50% of the training data in Pubmed. Third Row: We set $\alpha = 10$, $\lambda = 10^{-4}$, and removed up to 30% of the training data in Amazon Photo.

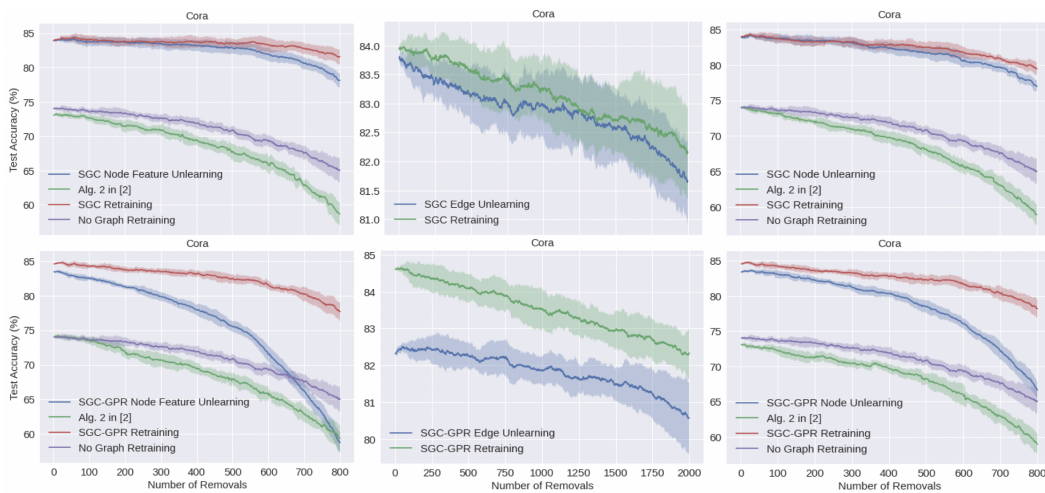


Figure 8: Performance of certified graph unlearning methods on Cora. First Row: The reported results are based on averaging over 10 repeated trails with random splitting. Second Row: GPR-based models are used to obtain node embeddings. All other settings are the same as in Figure 3.