

Certified Edge Unlearning for Graph Neural Networks

Kun Wu
kwu14@stevens.edu
Stevens Institute of Technology
Hoboken, New Jersey, USA

Jie Shen
jie.shen@stevens.edu
Stevens Institute of Technology
Hoboken, New Jersey, USA

Yue Ning
yue.ning@stevens.edu
Stevens Institute of Technology
Hoboken, New Jersey, USA

Ting Wang
inbox.ting@gmail.com
Pennsylvania State University
University Park, Pennsylvania, USA

Wendy Hui Wang
hui.wang@stevens.edu
Stevens Institute of Technology
Hoboken, New Jersey, USA

ABSTRACT

The emergence of evolving data privacy policies and regulations has sparked a growing interest in the concept of “machine unlearning”, which involves enabling machine learning models to forget specific data instances. In this paper, we specifically focus on *edge unlearning* in Graph Neural Networks (GNNs), which entails training a new GNN model as if certain specified edges never existed in the original training graph. Unlike conventional unlearning scenarios where data samples are treated as independent entities, edges in graphs exhibit correlation. Failing to carefully account for this data dependency would result in the incomplete removal of the requested data from the model. While retraining the model from scratch by excluding the specific edges can eliminate their influence, this approach incurs a high computational cost. To overcome this challenge, we introduce CEU, a Certified Edge Unlearning framework. CEU expedites the unlearning process by updating the parameters of the pre-trained GNN model in a single step, ensuring that the update removes the influence of the removed edges from the model. We formally prove that CEU offers a rigorous theoretical guarantee under the assumption of convexity on the loss function. Our empirical analysis further demonstrates the effectiveness and efficiency of CEU for both linear and deep GNNs – it achieves significant speedup gains compared to retraining and existing unlearning methods while maintaining comparable model accuracy to retraining from scratch.

CCS CONCEPTS

• **Security and privacy**; • **Computing methodologies** → *Machine learning*;

KEYWORDS

Graph Unlearning; Graph Neural Networks; Machine Learning Security and Privacy

ACM Reference Format:

Kun Wu, Jie Shen, Yue Ning, Ting Wang, and Wendy Hui Wang. 2023. Certified Edge Unlearning for Graph Neural Networks. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23)*, August 6–10, 2023, Long Beach, CA, USA. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3580305.3599271>

1 INTRODUCTION

Legislation such as the General Data Protection Regulation (GDPR) [31], the California Consumer Privacy Act (CCPA) [27], and the Personal Information Protection and Electronic Documents Act (PIPEDA) [28] has introduced requirements for companies to honor user requests for the removal of private data. This has sparked discussions around the concept of the “right to be forgotten” [22], which empowers users to have more control over their data by requesting its deletion from learned models. When a company has already utilized user data to train their machine learning (ML) models, these models must be appropriately manipulated to reflect data deletion requests.

In this paper, we study Graph Neural Networks (GNNs) as the target model and edge removal as the unlearning request. To illustrate this scenario, let us consider an online social network platform where users request the elimination of their sensitive social relations. The platform owner is legally bound to remove the edges associated with these sensitive social relations from any GNN model trained on the graph containing those edges. This ensures that the model no longer “remembers” those sensitive social relations.

Naively erasing edges from a GNN model by fully retraining can be excessively time-consuming, particularly for complex GNN models trained on large graphs. As a result, recent efforts have focused on developing efficient methods for exact unlearning [7, 10] as well as approximate unlearning [9, 26] specifically tailored for GNNs. In this paper, our emphasis is on approximate graph unlearning methods that facilitate the removal of requested edges from the model without retraining from scratch. Our approach is inspired by the concept of *influence function*, which enables the estimation of the impact of individual data samples on learning models [21]. To prove that the resulting model has removed the information related to the deleted edges, our goal is to provide a rigorously certified guarantee [15, 16] of the statistical indistinguishability between the retrained model and the unlearning model.

Despite the plethora of research on machine unlearning for non-graph datasets (e.g., [3–5, 23]), none of these approaches can be directly applied to GNNs due to the presence of data dependency

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '23, August 6–10, 2023, Long Beach, CA, USA.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0103-0/23/08...\$15.00

<https://doi.org/10.1145/3580305.3599271>

within graphs. Failing to carefully account for this data dependency would result in an incomplete removal of the requested data from the model. While recent efforts have been made to develop exact and approximate edge unlearning methods for GNNs [7, 9, 10], exact unlearning methods suffer from potentially significant loss of model accuracy [7, 10]. On the other hand, the existing approximate unlearning methods either lack a certified guarantee [8] or are limited to GNN models with specific structures [9].¹

Our contributions. We design CEU, a Certified Edge Unlearning algorithm that removes requested edges from GNNs without the need for retraining while providing a provable guarantee of the unlearning model. Our contributions are outlined as follows.

► **Unlearning through influence analysis.** We formulate the unlearning problem as finding a closed-form update on the model parameters. To achieve this, we introduce a novel influence function that efficiently computes the necessary update, while also taking into account the neighborhood of the removed edges. We address several theoretical and practical challenges of deriving edge influence by providing an influence estimator that is computationally and memory efficient.

► **Certified unlearning.** We undertake in-depth theoretical analysis and present non-trivial findings. We provide formal proofs demonstrating that CEU can deliver a rigorous (ϵ, δ) -approximation guarantee under the assumption of a strongly convex loss function. Additionally, we derive both worst-case and data-dependent bounds for the statistical distance between the retrained model and the model obtained through unlearning using CEU.

► **Empirical analysis.** Through extensive empirical study, we showcase the efficiency and effectiveness of CEU for both linear and deep GNN models. Specifically, for linear GNNs, we demonstrate that CEU achieves effective unlearning with a remarkable 16.2-fold speedup compared to retraining from scratch. Notably, our method outperforms the exact graph unlearning approach [7] in terms of both model accuracy and unlearning efficiency, exhibiting a 63% improvement in model accuracy and a 3.7-fold speedup. Additionally, it surpasses the existing certified graph unlearning method [9] in terms of unlearning efficiency, achieving a speedup of at least two orders of magnitude. Moving on to deep GNNs, our empirical results highlight the high efficiency of CEU, providing a speedup of up to 5 times compared to retraining while maintaining similar model accuracy. Furthermore, we quantitatively assess the efficacy of unlearning by conducting a link membership inference attack [19] on unlearning models. We demonstrate that the attack accuracy of inferring the removed edges from the unlearning model is comparable to that from the retrained model, indicating the successful removal of the targeted edges.

2 RELATED WORK

Machine unlearning [2, 17, 25] refers to a process that aims to remove the impact of a set of data samples in the training set from a trained model. From the certainty of unlearning, the existing

machine unlearning methods can be divided into two categories: *exact unlearning* and *approximate unlearning*.

Exact machine unlearning. In exact unlearning, a model is naively retrained from scratch after removing certain data samples from the dataset. This is generally computationally expensive. Several attempts have been made to make unlearning more efficient than retraining from scratch. An earlier study converts ML algorithms to statistical query (SQ) learning so that unlearning only needs to retrain the summation of SQ learning [4]. The SISA (sharded, isolated, sliced, and aggregated) approach [2] trains a set of constituent models on disjoint data *shards*. Only the shards affected by the unlearning requests and their constituent models are retrained. Some recent works [7, 10] extend exact unlearning to the graph setting. In particular, *GraphEraser* [7] adapts the SISA approach to graph unlearning. It splits graphs into disjoint partitions. Upon receiving an unlearning request, only the model on the affected shards is retrained. However, as shown in our empirical studies later (Section 6), *GraphEraser* suffers from a significant loss of model accuracy, as splitting the training graph into disjoint partitions damages the original graph structure. *GraphEditor* [10] designs an exact unlearning solution of linear GNNs. However, it is restricted to the linear structure only. It also cannot deal with efficient batch removal of a large number of edges.

Approximate machine unlearning. Approximate unlearning relaxes the requirement for exact unlearning by requiring that the removed data is statistically unlearned with the guarantee that the unlearning model cannot be distinguished from an exact deletion model [16], where the indistinguishability is defined in a similar manner as differential privacy [12]. Certified unlearning can be realized by adding noise either on the weights [14, 15, 25, 32, 39] or on the loss function [16]. In the context of graph unlearning, Chien *et al.* [9] provide the first certified GNN unlearning solution. However, their approach is restricted to GNN models of certain structures such as Simple Graph Convolution (SGC) and its generalized PageRank (GPR). And their implementation cannot be easily adapted to general GNNs. Furthermore, their approach cannot support batch edge removal. Our empirical results show that CEU is much faster than [9] in batch edge unlearning, with a speed-up of at least two orders of magnitude. Their follow-up work [26] extends to a particular type of nonlinear GNN models based on Graph Scattering Transform (GST). However, [26] considers node unlearning not edge unlearning. On the other hand, the approximate edge unlearning solution proposed by Cheng *et al.* [8] cannot provide any certified guarantee.

3 PROBLEM FORMULATION

Problem setup. Let \mathcal{G} be a set of graphs. In this paper, we only consider undirected graphs. Let Θ be the parameter space of GNN models. A learning algorithm \mathcal{A}_L is a function that maps an instance $G(V, E) \in \mathcal{G}$ to a parameter $\theta \in \Theta$. Let θ_{OR} be the parameters of \mathcal{A}_L trained on G . Any user can submit an edge unlearning request to remove specific edges from G . In practice, unlearning requests are often submitted sequentially. For efficiency, we assume these requests are processed in a batch. Let E_{UL} denote the batch of edges that are requested to be removed. As a response to these requests,

¹Both theoretical analysis and algorithmic techniques of [9] are closely tied to linear GNNs such as simple graph convolutions (SGC) and their generalized PageRank (GPR) extensions.

\mathcal{A}_L has to erase the impacts of E_{UL} on \mathcal{A}_L and produce an unlearning model. A straightforward approach is to retrain the model on $G(V, E \setminus E_{UL})$ from scratch and obtain the model parameters θ_{RE} . However, due to the high computational cost of retraining, an alternative solution is to apply an *unlearning* process \mathcal{A}_{UL} that takes E_{UL} and θ_{OR} as input and outputs an unlearning model.

Certified guarantee. Approximate unlearning requires some format of guarantee that the information related to the deleted data has been removed from the model. Intuitively, if the result of unlearning is likely to be obtained by retraining, then the unlearning algorithm has successfully eliminated the influence of the removed data points from the model. Following this intuition, we adapt the concept of *certified removal* [16, 25] to our setting to measure the difference between the retrained model and one obtained by unlearning. Broadly speaking, certified removal defines the indistinguishability between the retrained model and the unlearning model in a similar manner as (ϵ, δ) -differential privacy [12]. In particular, it defines the notion of (ϵ, δ) -approximate unlearning which is formalized as follows.

DEFINITION 1 ((ϵ, δ) -Approximate Unlearning). *Given a learning algorithm \mathcal{A}_L and two constants $\epsilon, \delta > 0$, an unlearning algorithm \mathcal{A}_{UL} performs (ϵ, δ) -certified unlearning for \mathcal{A}_L if*

$$P(\mathcal{A}_{UL}(D, z, \mathcal{A}_L(D))) \leq e^\epsilon P(\mathcal{A}_L(D \setminus z)) + \delta, \quad (1)$$

and

$$P(\mathcal{A}_L(D \setminus z)) \leq e^\epsilon P(\mathcal{A}_{UL}(D, z, \mathcal{A}_L(D))) + \delta, \quad (2)$$

where z is the sample to be removed.

Intuitively, Def. 1 guarantees that the unlearning model is “approximately” the same as the retrained model, where the difference between the unlearning and retrained model is bounded by the parameters of ϵ and δ . Smaller ϵ and δ indicate that the unlearning model is closer to the retrained model. Trivially, a $(0, 0)$ -approximate unlearning model is equivalent to the retrained model.

We adapt the notion (ϵ, δ) -approximate unlearning to edge removal, and formalize the edge unlearning problem as follows:

DEFINITION 2 ((ϵ, δ) -approximate Edge Unlearning). *Given a graph $G(V, E)$, a set of edges $E_{UL} \subset E$ that are requested to be removed from G , a graph learning algorithm \mathcal{A}_L and its readout function f , then an edge unlearning algorithm \mathcal{A}_{UL} performs (ϵ, δ) -certified unlearning for \mathcal{A}_L if:*

$$P(\mathcal{A}_{UL}(G, E_{UL}, \mathcal{A}_L(G))) \leq e^\epsilon P(\mathcal{A}_L(G_{UL})) + \delta, \quad (3)$$

and

$$P(\mathcal{A}_L(G_{UL})) \leq e^\epsilon P(\mathcal{A}_{UL}(G, E_{UL}, \mathcal{A}_L(G))) + \delta, \quad (4)$$

where $\epsilon, \delta > 0$, and $G_{UL} = G(V, E \setminus E_{UL})$.

While Def. 1 is defined for a single data sample, we extend it to the removal of a set of samples (edges) to handle batch edge removal. Our goal is to seek the unlearning mechanism \mathcal{A}_{UL} that can remove multiple edges at once with (ϵ, δ) -approximate guarantee while its computational complexity is significantly cheaper than retraining.

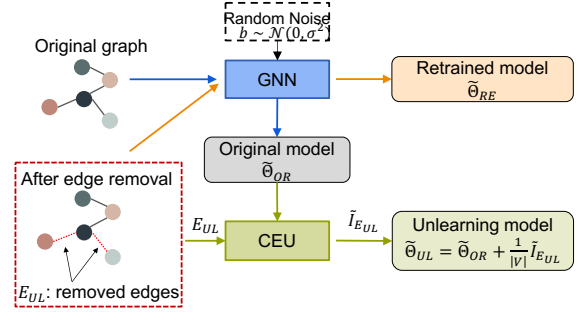


Figure 1: The framework of CEU. Orange lines indicate the process of retraining and green lines indicate unlearning.

4 METHODOLOGY

Given a graph $G(V, E)$ as input, we can find a model represented by θ that fits the data by minimizing an empirical loss. In this paper, we consider cross-entropy loss [11] as our loss function. The original model θ_{OR} is obtained by solving the following program:

$$\theta_{OR} = \arg \min_{\theta} \frac{1}{|V|} \sum_{v \in V} \mathcal{L}(\theta; v, E). \quad (5)$$

Assume a set of edges E_{UL} is deleted from G and let the new graph after the deletion be $G_{UL} = G(V, E \setminus E_{UL})$, retraining the model will obtain a new model parameter θ_{RE} on G_{UL} :

$$\theta_{RE} = \arg \min_{\theta} \frac{1}{|V|} \sum_{v \in V} \mathcal{L}(\theta; v, E \setminus E_{UL}). \quad (6)$$

A major difficulty, as expected, is that obtaining θ_{RE} is prohibitively slow for complex networks and large datasets. To overcome this challenge, we will identify a closed-form update $I_{E_{UL}}$ to θ_{OR} :

$$\theta_{UL} \approx \theta_{OR} - I_{E_{UL}}, \quad (7)$$

where $I_{E_{UL}}$ has the same dimension as the learning model θ_{OR} . Intuitively, θ_{UL} approximates the retraining. Such approximation, however, may not be able to provide any unlearning guarantee, as the direction of the gradient residual of θ_{UL} may still be able to leak information about the removed edges.

Overview of CEU. We design CEU as a two-step process. In Step 1, CEU adds the perturbation to the loss function, aiming to hide the real gradient residual and provide the certified unlearning guarantee. Let $\tilde{\theta}_{OR}$ be the parameters of the model trained with the noisy loss function. In Step 2, CEU estimates the one-shot update on the parameters $\tilde{\theta}_{OR}$ through influence analysis. Figure 1 illustrates an overview of CEU. Next, we describe the details of the two steps.

4.1 Step 1: Adding Perturbation on Loss Function

To enable unlearning with a certified guarantee, we follow the same idea of certified data removal [16] and add a linear noise term to the training loss, aiming to hide the real gradient residual. We use \mathcal{L}_b to denote the loss function with noise formalized as follows:

$$\mathcal{L}_b = L(\theta, E) + \frac{\lambda}{2} \|\theta\|^2 + b^\top \theta, \quad (8)$$

where b is drawn randomly from the Gaussian distribution $\mathcal{N}(0, \sigma^2)$. The randomness in b will mask any potential information leaked by the estimated edge influence. The resulting perturbed learning problem can be solved using standard convex optimization methods.

4.2 Step 2: Unlearning through Influence Analysis

Intuitively, updating model parameters for unlearning can be interpreted from the optimization perspective that the model forgets E_{UL} by “reversing” the influence $\tilde{\theta}_{UL}$ of E_{UL} from the model. The challenge is how to estimate the influence of $\tilde{\theta}_{UL}$ on the model.

Influence functions [21] enable efficient approximation of the effect of some particular training points on a model’s prediction. Intuitively, the influence function computes the parameters after the removal of z by upweighting z on the parameters with some small ζ :

$$\hat{\theta}_{\zeta, z} = \arg \min_{\theta} \frac{1}{m} \sum_{z_i \neq z} \mathcal{L}(\theta; z_i) + \zeta \mathcal{L}(\theta; z), \quad (9)$$

where m is the number of data points in the original dataset, and ζ is a small constant. The influence function is not restricted to a single point. We can define a set of points Z and compute $\hat{\theta}_{\zeta, Z}$.

However, most of the existing influence functions cannot be directly applied to the GNN setting, as removing one edge $e(v_i, v_j)$ from the graph can affect not only the prediction of v_i and v_j but also those of neighboring nodes of v_i and v_j , due to the aggregation function of GNN models. To address this challenge, we design a new influence function for GNNs that take the neighborhood into consideration when estimating the influence of the neighborhood of removing an edge on model parameters.

In general, an ℓ -layer GNN aggregates the information of the ℓ -hop neighborhood of each node. Thus removing an edge $e(v_i, v_j)$ will affect not only v_i and v_j but also all nodes in the ℓ -hop neighborhood of v_i and v_j . To capture such aggregation effect in the derivation of edge influence, first, we define the set of nodes (denoted as V_e) that will be affected by removing an edge $e(v_i, v_j)$ as: $V_e = \mathcal{N}(v_i) \cup \mathcal{N}(v_j) \cup \{v_i, v_j\}$, where $\mathcal{N}(v)$ is the set of nodes connected to v in ℓ hops. Furthermore, we define the set of nodes (denoted as $V_{E_{UL}}$) that will be affected by removing a set of edges E_{UL} as $V_{E_{UL}} = \bigcup_{e \in E_{UL}} V_e$.

To revert the influence of E_{UL} on the target model, we compute the new parameters $\theta_{\zeta, E_{UL}}$ after the removal of E_{UL} as follows:

$$\begin{aligned} \theta_{\zeta, V_{E_{UL}}} = \arg \min_{\theta} \frac{1}{|V|} \sum_{v \in V} \mathcal{L}_b(\theta; v, E) + \zeta \left(\sum_{v \in V_{E_{UL}}} \mathcal{L}_b(\theta; v, E \setminus E_{UL}) \right. \\ \left. - \sum_{v \in V_{E_{UL}}} \mathcal{L}_b(\theta; v, E) \right). \end{aligned} \quad (10)$$

Eqn. (10) contains three terms. While the first term measures the loss of the original model, the second and the third ones together compute the loss of the nodes affected by the removal of E_{UL} . Following this reasoning, Eqn. (10) is equivalent to Eqn. (6) when $\zeta = \frac{1}{|V|}$, where $|V|$ is the total number of nodes in the original graph (the proof is included in our full version [38]). Following this reasoning, instead of solving the problem in Eqn. (10), we formulate the optimization problem as a closed-form update on the original

model $\tilde{\theta}_{OR}$ with the noisy loss function (by Step 1):

$$\tilde{\theta}_{UL} = \tilde{\theta}_{OR} + \frac{1}{|V|} \tilde{I}_{E_{UL}}, \quad (11)$$

where $\tilde{I}_{E_{UL}}$ is the influence of E_{UL} on the target model with noisy loss. By utilizing this formulation, we can describe changes in the training graph structure by edge removal as a one-shot update on model parameters.

In this paper, we take a *second-order* update strategy that utilizes second-order derivatives to calculate the closed-form update $\tilde{I}_{E_{UL}}$. Our second-order update result is present in the following theorem.

THEOREM 3. *Given the parameters θ_{OR} obtained by \mathcal{A}_{UL} on a graph G , and the loss function \mathcal{L} , assume that \mathcal{L} is twice-differentiable and convex in θ , then the influence of a set of edges E_{UL} is:*

$$\tilde{I}_{E_{UL}} = -H_{\tilde{\theta}_{OR}}^{-1} \left(\nabla_{\theta} \sum_{v \in V_{E_{UL}}} \mathcal{L}_b(\theta_{OR}; v, E \setminus E_{UL}) - \nabla_{\theta} \sum_{v \in V_{E_{UL}}} \mathcal{L}_b(\tilde{\theta}_{OR}; v, E) \right), \quad (12)$$

where $H_{\tilde{\theta}_{OR}} := \nabla^2 \frac{1}{|V|} \sum_{v \in V} \mathcal{L}_b(\tilde{\theta}_{OR}, v, E)$, and $H_{\tilde{\theta}_{OR}}^{-1}$ is the inverse Hessian of the loss at $\tilde{\theta}_{OR}$.

The proof of Theorem 3 can be found in our full version [38]. Theorem 3 assumes the loss function is convex. Given the non-convexity nature of GNN models, the Hessian matrix can be non-invertible and thus there may not have a solution for the influence estimation. To address this issue, we follow [21] and add a damping term λ_1 to $H_{\tilde{\theta}_{OR}}$ (i.e., $H_{\tilde{\theta}_{OR}} + \lambda_1 I$) if $H_{\tilde{\theta}_{OR}}$ has negative eigenvalues, where λ_1 is the same as the regularization rate λ in Eqn. (8). Our empirical analysis (Sec. 6) will show this solution enables effective unlearning in practice.

There are several practical and theoretical challenges in calculating the influence (Eqn. (12)). First, for large graphs, even storing a Hessian matrix in memory is expensive: in our experiments, we will show that Hessian matrices are huge, e.g. the Hessian matrix on the CS dataset has a size of around $10^5 \times 10^5$ which would cost 50 GB memory. Second, even under the promise that the linear system is feasible, computing the inverse of a matrix of huge size is prohibitive. To address these two challenges, we design an algorithm that approximates the inverse Hessian. Note that the existing certified graph unlearning method [9] did not use any influence estimator. Instead, it computes the exact inverse Hessian.

The starting point of our algorithm is a novel perspective that solving the linear system can be thought of as finding a *stationary point* of the quadratic function f : $f(x) = \arg \min_x \frac{1}{2} x^T B x - k^T x$, where $B = \tilde{H}_{\tilde{\theta}_{OR}}$, and

$$k = \nabla_{\theta} \sum_{v \in V_{E_{UL}}} \mathcal{L}_b(\tilde{\theta}_{OR}; v, E \setminus E_{UL}) - \nabla_{\theta} \sum_{v \in V_{E_{UL}}} \mathcal{L}_b(\tilde{\theta}_{OR}; v, E). \quad (13)$$

The random noise b by Step 1 does not appear in B due to the second-order derivation. It does not appear in k either because it was contained in both terms in Eqn. (13) and thus was canceled. By leveraging the stationary point, a convergence guarantee can be established using gradient-descent-type algorithms [1].

We employ the implementation [24] that combines Hessian-vector product (HVP) [29] and the conjugate gradient (CG) [35] to approximate the inverse Hessian. CG exhibits promising computational efficiency for minimizing quadratic functions [30]. It

is well-known that, as long as some regularity conditions (e.g., the objective function is Lipschitz and bounded) are met, the CG asymptotically converges to a stationary point. This stationary point corresponds to a solution of $\tilde{I}_{E_{UL}}$ (Eqn. (12)). Hence, we have the following convergence guarantee of influence estimation.

LEMMA 4 (THEOREM 2.1 OF [30]). *The CG method generates a sequence of iterates $\{x_t\}_{t \geq 1}$ such that $\lim_{t \rightarrow +\infty} f(x_t) = 0$. In addition, the per-iteration time complexity is $O(|x|)$ where $|x|$ denotes the dimension of x .*

We note, however, that an appealing feature of Eqn. (12) is that it does not need to find a solution with an exact-zero gradient. This enables us to terminate CG early by monitoring the magnitude of the gradients. Our empirical study also shows that CG can get a good approximation in a small number of iterations.

Besides time efficiency, we have the following lemma showing that the CG method is memory-efficient.

LEMMA 5. *The CG method can be implemented using $O(|\theta|)$ memory.*

The proof of Lemma 5 can be found in our full version [38].

5 CERTIFIED UNLEARNING GUARANTEE

As the unlearning model by CEU approximates the retrained model, ideally it should provide the theoretical guarantee that the unlearning model is statistically indistinguishable from the retrained one. Next, we derive conditions under which the second-order update by CEU can provide the (ϵ, δ) -approximate unlearning guarantee. To construct theoretical guarantees for our approach, we make the following assumptions on the GNN models.

ASSUMPTION 6. *For the given GNN model and its loss function L :* (1) L is a strictly convex loss function that is twice differentiable; (2) $\|\nabla L\|_2 \leq c_1$; (3) $\nabla^2 L$ is γ_1 -Lipschitz; (4) ∇L is γ_2 -Lipschitz; and (5) the node features x_v is bounded: $\|x_v\|_2 \leq 1, \forall v \in V$. Here c_1, γ_1, γ_2 are positive constants.

These assumptions can be satisfied by a wide range of GNNs such as Simple Graph Convolution (SGC) [6, 37] and Graph Linear Network (GLN) [36] which can achieve the comparable performance compared with deep GNNs [13, 36, 37]. It is important to note that, although our theoretical analysis relies on the assumption of strictly convex loss function, our algorithmic techniques are generic and can be applied to various GNN models, including non-convex ones. We will show that CEU can achieve notable empirical performance on both linear and deep GNNs (Section 6).

Following the state-of-the-art certified removal work [16], we utilize the gradient residual $\|\nabla \mathcal{L}\|_2$ for the proof of certified guarantee. Intuitively, for strongly convex loss functions, the gradient residual is zero as the optimum is unique. Hence, the norm of the gradient residual $\|\nabla \mathcal{L}\|_2$ can reflect the distance between the retrained and the unlearning models. Based on this, CEU can establish the (ϵ, δ) -approximation guarantee by following Theorem 7.

THEOREM 7 (THEOREM 3 FROM [16]). *Let \mathcal{A}_L be the learning algorithm that returns the unique optimum of the loss \mathcal{L}_b and let \mathcal{A}_{UL} be the unlearning mechanism. Suppose that $\|\nabla \mathcal{L}_b\|_2 \leq \epsilon'$ for some computable bound $\epsilon' > 0$. If $b \sim \mathcal{N}(0, c\epsilon'/\epsilon)^d$ with some constants*

$c, \epsilon > 0$, where d is the parameter dimension, then \mathcal{A}_{UL} provides (ϵ, δ) -approximation guarantee for \mathcal{A}_L , where $\delta = 1.5e^{-c^2/2}$.

Intuitively, Theorem 7 requires the *gradient residual norm* $\|\nabla \mathcal{L}_b\|_2$ to be bounded appropriately in order to provide the approximation guarantee. Thus, our theoretical analysis mainly focuses on finding the bound of $\|\nabla \mathcal{L}_b\|_2$. First, we present the worst-case bound of $\|\nabla \mathcal{L}_b\|_2$ in Theorem 8.

THEOREM 8 (Worst-case Bound). *Assume Assumption 6 holds. Then we have the following worst-case bound of $\|\nabla \mathcal{L}_b\|_2$:*

$$\|\nabla \mathcal{L}_b(\tilde{\theta}_{UL}, E \setminus E_{UL})\|_2 \leq \frac{\gamma_1 \gamma_2^2 c_1^2}{\lambda^4 |V|} \left(\sum_{v \in E_{UL}} n_v \right)^2, \quad (14)$$

where n_v is the number of neighbors of node v , λ is the regularization rate (Eqn. (8)), and $|V|$ is the number of nodes in the training graph.

The proof of Theorem 8 is provided in our full version [38]. As $\frac{1}{\lambda^4}$ in Theorem 8 can be large, the worst-case bound can be impractically loose. Therefore, next, we derive the *data-dependent* bound on $\|\nabla \mathcal{L}_b\|_2$ in Theorem 9.

THEOREM 9 (Data-dependent Bound). *Suppose Assumption 6 holds. Then we have the following data-dependent bound of $\|\nabla \mathcal{L}_b\|_2$:*

$$\|\nabla \mathcal{L}_b(\tilde{\theta}_{UL}, E \setminus E_{UL})\|_2 \leq \gamma_1 \frac{1}{|V|^2} \|\tilde{H}_{\theta_{OR}}^{-1} \Delta\|_2^2, \quad (15)$$

where

$$\Delta = \nabla_{\theta} \sum_{v \in V_{E_{UL}}} \mathcal{L}(\tilde{\theta}_{OR}; v, E) - \nabla_{\theta} \sum_{v \in V_{E_{UL}}} \mathcal{L}(\tilde{\theta}_{OR}; v, E \setminus E_{UL}).$$

The proof of Theorem 9 can be found in our full version [38]. The data-dependent bound can be computed efficiently by using the influence estimator (Sec. 4.2). We will show that the data-dependent bound is much tighter than the worst-case bound in Section 6.

6 EXPERIMENTS

In this section, we empirically verify the efficiency and effectiveness of CEU. The code and datasets are publicly available ².

6.1 Experimental Setup

All experiments are executed on a GPU server with NVIDIA A100 (40G). All the algorithms are implemented in Python with PyTorch. Each experiment is repeated 10 times and the average is reported.

Datasets. We use three datasets, namely **Cora** [33], **Citeseer** [41], and **CS** [34] datasets, that are popularly used for performance evaluation of GNNs [34, 42]. The statistical information of these datasets can be found in Appendix A.

GNN models. We consider two types of GNN models: (1) **Linear models:** We consider a simplified GCN model that contains only one layer and a softmax function (without normalization). (2) **Deep models:** We consider three representative GNN models, namely **GCN** [20], **GraphSAGE** [18], and **GIN** [40]. For these GNN models, we consider various network complexity (up to four hidden layers) in the experiments, with the same number of neurons as 32 at each layer respectively. All GNN models are trained for 1,000 epochs with

²https://github.com/kunwu522/certified_edge_unlearning

an early stop condition that the validation loss does not decrease for 20 epochs. We randomly split each graph into a training set (70%), a validation set (10%), and a test set (20%). More details of the setup of model parameters can be found in Appendix A.

Edges for removal. We randomly pick $k = \{200, 400, 600, 800, 1,000\}$ edges from Cora and CiteSeer datasets, and $k = \{2,000, 4,000, 6,000, 8,000, 10,000\}$ edges from CS dataset for removal. We pick more edges from the CS dataset as its number of edges is orders of magnitude higher than the other two datasets (Table 4).

Metrics. We evaluate the performance of CEU in terms of *efficiency*, *efficacy*, and *model accuracy*: (1) **Unlearning efficiency** is measured as the running time of CEU; (2) **Target model accuracy** is measured as the *accuracy* of node classification, i.e., the percentage of nodes that are correctly classified by the model; (3) **Unlearning efficacy**: We utilize *StealLink* [19], a SOTA edge membership inference attack, to empirically evaluate the extent to which the model has forgotten the removed edges.³ *StealLink* predicts whether particular edges exist in the training graph. We measure the unlearning efficacy as AUC of *StealLink*'s inference of whether the removed edges were present in the original graph. Intuitively, a higher AUC indicates lower unlearning efficacy. AUC close to 0.5 indicates that the model has removed the requested edges.

Noise setup. We follow the same setting of [9] and set $\lambda = 0.01$ and $\sigma = 0.1$ (Eqn. (8)). We use the same ϵ ($\epsilon = 0.1 - 10$) as in [9].

Baselines. We consider three baselines of exact and approximate GNN unlearning for comparison with CEU.

- **Exact unlearning:** We consider GraphEraser [2], the SOTA exact edge unlearning method. GraphEraser has two partitioning strategies denoted as *balanced LPA (BLPA)* and *balanced embedding k-means (BEKM)*. We consider both BLPA and BEKM as the baseline methods as these two methods exhibit varying performance in different settings.
- **Uncertified unlearning (UEU):** We estimate the influence of the removed edges on the original model (i.e., no noise on the loss function), and apply similar influence analysis (Section 4.2) to derive the one-shot update on model parameters. More details of UEU can be found in our full version [38].
- **Certified unlearning:** We consider Certified Graph Unlearning (CGU) [9] as a baseline.⁴

Two retraining settings. As CEU adds noise to the loss function of the target model, we consider two different retraining settings denoted as "Retrain" and "R+N" respectively.

6.2 Tightness of Bounds

The tightness of both worst-case and data-dependent bounds of the gradient residual norm determines the strictness of the certified guarantee. To evaluate the tightness of both bounds, we consider the 1-layer GCN model and measure the real gradient residual norm values (as the ground truth) as well as the two bounds. Figure 2 reports the value of the two bounds as a function of the number of removed edges. We have two main observations. First, as expected, the worst-case bound is much looser than the data-dependent bound. It can be

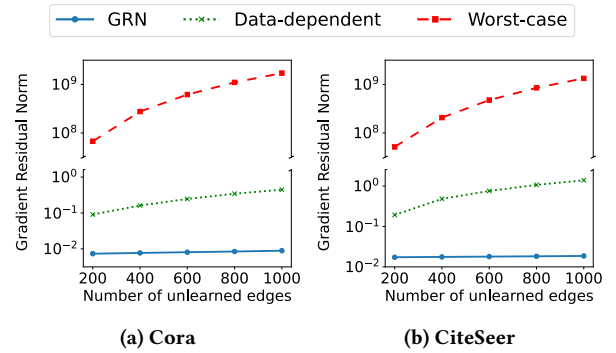


Figure 2: Tightness of bounds (GRN: Gradient residual norm).

Table 1: Model accuracy of CEU, retraining (Retrain and R+N), and baselines (BLPA, BEKM, UEU) (Linear GCN, CS dataset).

| Type | Method | Number of removed edges | | | | | |
|---------|---------|-------------------------|------|------|------|------|------|
| | | 0 | 2K | 4K | 6K | 8K | 10K |
| Retrain | Retrain | 0.93 | 0.93 | 0.93 | 0.93 | 0.93 | 0.93 |
| | R+N | 0.91 | 0.91 | 0.91 | 0.91 | 0.90 | 0.90 |
| Unlearn | BLPA | 0.84 | 0.69 | 0.80 | 0.84 | 0.84 | 0.68 |
| | BEKM | 0.64 | 0.80 | 0.56 | 0.83 | 0.77 | 0.67 |
| | UEU | 0.93 | 0.93 | 0.93 | 0.93 | 0.93 | 0.93 |
| | CEU | 0.91 | 0.91 | 0.91 | 0.91 | 0.90 | 0.90 |

several orders of magnitude larger than the data-dependent bound. The looseness in the bound comes from $\frac{1}{\lambda^4}$ in the bound. Second, the data-dependent bound is close to the ground-truth gradient residual norm, regardless of the growth in the number of removed edges. Given the tightness of the data-dependent bounds, CEU is expected to handle batch removal of a large number of edges.

6.3 Performance of Linear GCN Models

In this section, we only consider linear GCN models (i.e., 1-layer GCN model), and evaluate the performance of CEU for this model on three graph datasets in terms of model accuracy, unlearning efficiency, and unlearning efficacy. The results of UEU show the impact of noise on model performance compared with CEU.

Besides these results, we have additional results of the following studies: (1) the impacts of types of removed edges on unlearning performance; (2) the performance of sequential unlearning. These results can be found in our full version [38].

Model accuracy. Table 1 reports the results of GCN model accuracy on the CS dataset. The results on Cora and Citeseer datasets are similar and can be found in Appendix C.1. We have the following observations. First, the model accuracy obtained by CEU stays very close to that of the retrained model, regardless of the number of removed edges. The difference in model accuracy between the retrained and unlearning models remains negligible (in the range of [0.01%, 0.11%]). Second, in terms of comparison with both exact unlearning baselines (BEKM, BLPA), the model accuracy by CEU is significantly higher than these two baselines in all the settings. For example, when removing 4,000 edges, both BEKM and BLPA only

³Implementation of StealLink: https://github.com/xinleihe/link_stealing_attack

⁴Implementation of CGU [9]: https://github.com/thupchnsky/sgc_unlearn




Figure 3: Time performance of CEU, retraining (Retrain and R+U), and baselines (BLPA, BEKM, UEU) for linear GCN model.

Table 2: Unlearning efficacy of CEU, retraining, and UEU (Linear GCN, Cora dataset).

| $ E_{UL} $ | Original | Retrain | UEU | R+N | CEU |
|------------|----------|---------|-------|-------|-------|
| 200 | 0.930 | 0.577 | 0.572 | 0.533 | 0.535 |
| 400 | 0.936 | 0.582 | 0.580 | 0.541 | 0.543 |
| 600 | 0.935 | 0.582 | 0.580 | 0.547 | 0.547 |
| 800 | 0.936 | 0.589 | 0.585 | 0.549 | 0.552 |
| 1,000 | 0.935 | 0.586 | 0.592 | 0.559 | 0.553 |

can deliver model accuracy of around 0.56 and 0.80, while CEU can deliver a model accuracy of around 0.91 (63% and 14% improvement). This demonstrates the weakness of the exact unlearning through graph partitioning - breaking the graph structure can bring non-negligible model accuracy loss. Third, regarding the comparison with the approximate unlearning baseline (UEU), CEU has very similar model accuracy, although UEU does not add perturbation to the model loss function. This demonstrates that CEU addresses the trade-off between privacy and model accuracy—it can deliver a provable unlearning guarantee while requiring negligible sacrifice on model accuracy.

Unlearning efficiency. We report the time performance results of CEU in Figure 3. Our observations are followings. First, CEU is significantly faster than retraining from scratch. It speeds up by 11.4 \times , 6.4 \times , and 16.2 \times for Cora, CiteSeer, and CS datasets, respectively. Second, CEU is much faster than both BEKM and BLPA baselines, especially when training large graphs. For example, CEU is 3.7 \times faster than both BLPA and BEKM on the CS dataset when 4,000 edges and 10,000 edges were removed respectively (Figure 3 (c)). This demonstrates the advantage of the approximate unlearning methods. Third, for both approximate unlearning methods, CEU has comparable time performance as UEU although UEU is slightly faster than CEU.

Unlearning efficacy. Table 2 reports the attack performance of attack accuracy of the removed edges E_{UL} against the original model, retraining model (with and without noise), UEU, and CEU on the Cora dataset. We observe the following phenomena. First,

StealLink is highly effective to predict the existence of E_{UL} in the original graph (“Original” column), as the AUC of the attack against the original model is higher than 0.9 (much higher than 0.5). Second, the AUC of the attack is noticeably reduced to close to 0.5 for both retrained and unlearning models (“R+N” and “CEU” columns). This demonstrates that CEU has a similar ability to make the model forget the removed edges as retraining. Third, the AUC of both retraining and learning with noise (“R+N” and “CEU” columns) is lower than that without noise (“Retrain” and “UEU” columns). This demonstrates that the perturbation added to the loss function helps to reduce the privacy vulnerability of the removed edges.

Effects of ϵ on model accuracy. We study the effect of various ϵ values (for (ϵ, δ) -unlearning) on unlearning performance. The noise b is determined by using the data-dependency bound (Theorem 9) as ϵ' and ϵ together. Figure 4 reports the model accuracy with various ϵ values. We observe that, unsurprisingly, the model accuracy degrades when ϵ grows (i.e., more noise is added). For instance, when ϵ changes from 0.1 to 10, we witness the model accuracy drops from 0.925 to 0.9 when removing 2,000 edges from the CS dataset (Figure 4 (c)). Such model accuracy drop is more significant on Cora and CiteSeer datasets. The drop in model accuracy meets our expectation as higher ϵ allows a larger statistical distance between the retrained model and the unlearning model, and thus lowers the accuracy of the unlearning model.

Comparison with CGU [9]. As the approach presented in [9] is specifically designed for Simple Graph Convolutional (SGC) models, we apply both CEU and CGU to the SGC model to ensure a fair comparison of their performance. Figure 5 (a) reports the model accuracy of the retrained model and both CGU and CEU on the Cora dataset. The results on CiteSeer and CS datasets are similar and can be found in the full version [38]. We observe that the model accuracy of CEU stays close to CGU in all the settings. On the other hand, as shown in Figure 5 (b), CEU is much faster than CGU, with a speed-up by at least two orders of magnitude. Indeed, the speed-up is more compelling when more edges are removed. This shows the advantage of CEU for batch edge removal to CGU.

Besides model accuracy and unlearning efficiency, we also evaluated the unlearning efficacy of both CEU and CGU, and observed




Figure 4: Effect of unlearning parameter ϵ on model accuracy by CEU (Linear GCN model).

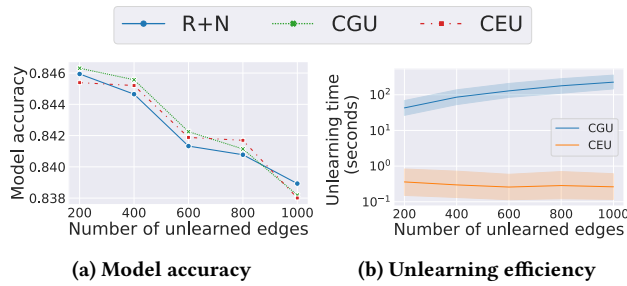


Figure 5: CGU [9] vs. CEU (SGC model, Cora dataset).

that CGU and CEU have comparable unlearning efficacy. Due to the limited space, we include the results in the full version [38].

6.4 Performance of Deep GNN Models

So far, we only consider the linear GCN model that meets Assumption 6. Next, we evaluate the performance of CEU on deep GNN models that do not meet Assumption 6. We consider GCN, GraphSAGE, and GIN models of various complexity (2-layer, 3-layer, and 4-layer) with ReLU as the activation function. We do not compare with the existing certified edge unlearning method [9] as it cannot be used on non-linear GNN models. Hence, we only compare CEU with the two baselines of exact edge unlearning (BLPA and BEKM).

Model accuracy. Figure 6 reports the model accuracy of the retrained model and CEU for the GCN model of various complexity. The results of GraphSAGE and GIN as well as the other two datasets are similar; they can be found in Appendix C.2. We observe two phenomena. First, although the model accuracy degrades for both retrained and unlearning GNN models of higher complexity, the model accuracy of the unlearning model remains close to that of the retrained model. The largest difference between model accuracy is only around 1.4% (Figure 6 (c)). Second, CEU outperforms two baselines (BLPA and BEKM) in terms of model accuracy for all the settings. For example, the model accuracy of CEU on the 4-layer GCN is 30% higher than BEKM when removing 10,000 edges (Figure 6 (c)). This demonstrates the advantage of CEU to the exact graph unlearning. We also observe that the model accuracy of both retrained and unlearning models is insensitive to the number of

edges. This is because the removed edges only takes a small portion (no more than 6%) of the original data.

Unlearning efficiency. Figure 7 shows the running time of retraining and CEU on GCN models with CS dataset. The time performance results of the other two datasets are included in Appendix C.2. We observe that, although the running time for both retraining and CEU grows with the increase in the complexity of GNN models, CEU is always significantly faster than retraining in all the settings, with the speedup factor as large as 5.2 \times . Furthermore, CEU is dominantly faster than the two baselines of exact unlearning (BLPA and BEKM), with a speedup as large as one-order magnitude.

Unlearning efficacy. Table 3 presents the attack performance of StealLink [19] of inferring the removed edges E_{UL} against the original model, the retrained model, the unlearning model by CEU, as well as by two baselines of exact unlearning (BLPA and BEKM) for GCN model on CS dataset. The results of the other settings can be found in Appendix C.2. We observe that, while StealLink is highly effective in predicting the presence of E_{UL} from the original model (“Orig.” column), its attack accuracy is significantly reduced to close to 0.5 when being launched against all the retraining/unlearning models. This indicates that CEU exhibits a similar capability as either retraining or exact unlearning to make deep GNN models forget the removed edges.

7 CONCLUSION

In this paper, we design CEU an efficient edge unlearning method that handles batch edge removal from GNNs. We prove that CEU can provide the theoretical guarantee of unlearning for GNN models under certain assumptions of convexity of the model’s loss function. Our extensive set of experiments demonstrates that CEU can achieve significant speedup gains over retraining while delivering similar model accuracy for both linear and deep GNN models.

There are several research directions for future work. An interesting direction will be extending to handle the removal of nodes from graphs. It is seemingly straightforward that node unlearning can be easily adapted from edge unlearning, as removing a node v from a graph is equivalent to removing all the edges that connect with v in the graph. However, node unlearning indeed is more challenging than edge unlearning, as removing a node entirely from the model

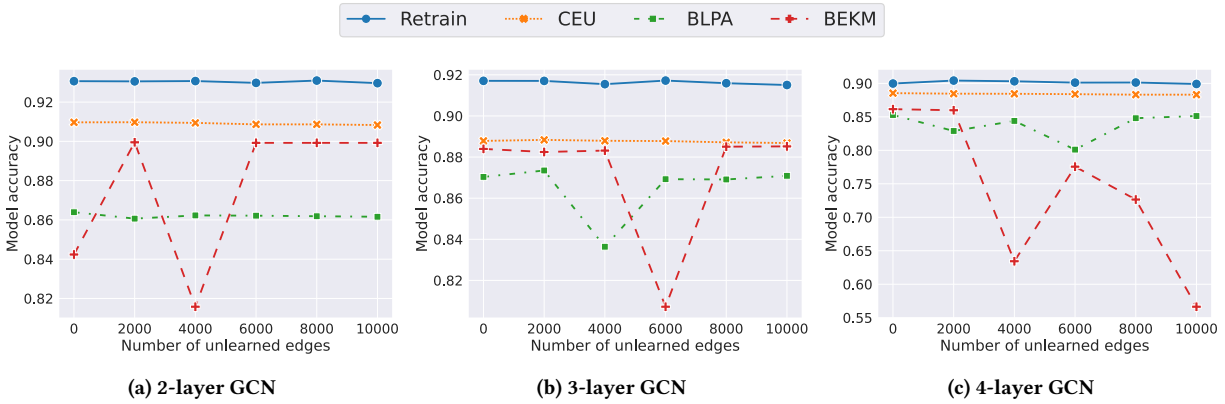


Figure 6: Model accuracy of CEU on deep GNN models (GCN, CS dataset).

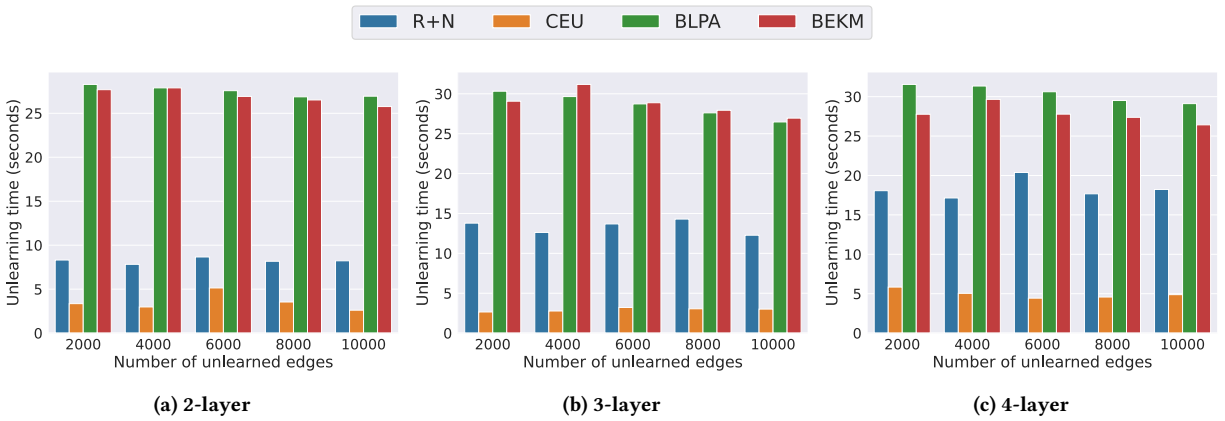


Figure 7: Time performance of retraining and CEU on deep GNN models (GCN, CS dataset).

Table 3: Unlearning efficacy of CEU and baselines on deep GNN models (GCN, CS dataset).

| $ E_{UL} $ | 2-layer | | | | | 3-layer | | | | | 4-layer | | | | |
|------------|---------|---------|-------|-------|-------|---------|---------|-------|-------|-------|---------|---------|-------|-------|-------|
| | Orig. | Retrain | CEU | BLPA | BEKM | Orig. | Retrain | CEU | BLPA | BEKM | Orig. | Retrain | CEU | BLPA | BEKM |
| 2K | 0.960 | 0.547 | 0.547 | 0.502 | 0.503 | 0.957 | 0.543 | 0.547 | 0.495 | 0.486 | 0.955 | 0.543 | 0.551 | 0.503 | 0.510 |
| 4K | 0.960 | 0.545 | 0.552 | 0.503 | 0.499 | 0.956 | 0.545 | 0.549 | 0.495 | 0.506 | 0.956 | 0.547 | 0.553 | 0.501 | 0.501 |
| 6K | 0.959 | 0.550 | 0.555 | 0.498 | 0.504 | 0.957 | 0.544 | 0.552 | 0.499 | 0.502 | 0.955 | 0.547 | 0.550 | 0.503 | 0.492 |
| 8K | 0.959 | 0.553 | 0.554 | 0.502 | 0.497 | 0.956 | 0.549 | 0.550 | 0.501 | 0.500 | 0.956 | 0.547 | 0.553 | 0.502 | 0.507 |
| 10K | 0.960 | 0.550 | 0.554 | 0.500 | 0.500 | 0.956 | 0.551 | 0.554 | 0.500 | 0.505 | 0.956 | 0.549 | 0.554 | 0.500 | 0.502 |

requires removing not only the edges connected with the node but also its features and labels. We will explore how to design efficient and certified node learning methods for the future work. Another interesting direction is to add additional constraints on unlearning. A possible constraint is the *unlearning capacity*, i.e., the maximum number of edges that can be deleted while still ensuring good model accuracy.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their feedback. This project is supported by the National Science Foundation (NSF) under Grant No. 2029038, 2135988, 1948432, 2047843, 2212323, 2119331, 1951729, and 1953893. Any opinions, findings, conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agency.

REFERENCES

- [1] D. P. Bertsekas. *Nonlinear Programming*. Massachusetts: Athena Scientific, 1999.
- [2] L. Bourtole, V. Chandrasekaran, C. A. Choquette-Choo, H. Jia, A. Travers, B. Zhang, D. Lie, and N. Papernot. Machine unlearning. In *Symposium on Security and Privacy (SP)*, pages 141–159, 2021.
- [3] J. Brophy and D. Lowd. Machine unlearning for random forests. In *International Conference on Machine Learning*, pages 1092–1104, 2021.
- [4] Y. Cao and J. Yang. Towards making systems forget with machine unlearning. In *Symposium on Security and Privacy*, pages 463–480, 2015.
- [5] G. Cauwenberghs and T. Poggio. Incremental and decremental support vector machine learning. *Advances in neural information processing systems*, 13, 2000.
- [6] M. Chen, Z. Wei, Z. Huang, B. Ding, and Y. Li. Simple and deep graph convolutional networks. In *International conference on machine learning (ICML)*, pages 1725–1735. PMLR, 2020.
- [7] M. Chen, Z. Zhang, T. Wang, M. Backes, M. Humbert, and Y. Zhang. Graph unlearning. In *ACM Conference on Computer and Communications Security (CCS)*, 2022.
- [8] J. Cheng, G. Dasoulas, H. He, C. Agarwal, and M. Zitnik. GNNDelete: A general strategy for unlearning in graph neural networks. In *The Eleventh International Conference on Learning Representations*, 2023.
- [9] E. Chien, C. Pan, and O. Milenkovic. Certified graph unlearning. In *NeurIPS 2022 New Frontiers in Graph Learning Workshop (NeurIPS GLFrontiers 2022)*, 2022.
- [10] W. Cong and M. Mahdavi. Grapheditor: An efficient graph representation learning and unlearning approach, 2023.
- [11] D. R. Cox. The regression analysis of binary sequences. *Journal of the Royal Statistical Society: Series B (Methodological)*, 20(2):215–232, 1958.
- [12] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284, 2006.
- [13] 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.
- [14] A. Gohilkar, A. Achille, A. Ravichandran, M. Polito, and S. Soatto. Mixed-privacy forgetting in deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 792–801, 2021.
- [15] A. Gohilkar, 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*, pages 9304–9312, 2020.
- [16] C. Guo, T. Goldstein, A. Hannun, and L. Van Der Maaten. Certified data removal from machine learning models. In *International Conference on Machine Learning*, pages 3832–3842, 2020.
- [17] V. Gupta, C. Jung, S. Neel, A. Roth, S. Sharifi-Malvajerdi, and C. Waites. Adaptive machine unlearning. *Advances in Neural Information Processing Systems*, 34, 2021.
- [18] W. Hamilton, Z. Ying, and J. Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.
- [19] X. He, J. Jia, M. Backes, N. Z. Gong, and Y. Zhang. Stealing links from graph neural networks. In *30th USENIX Security Symposium*, pages 2669–2686, 2021.
- [20] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *5th International Conference on Learning Representations, ICLR*, 2017.
- [21] P. W. Koh and P. Liang. Understanding black-box predictions via influence functions. In *International Conference on Machine Learning*, pages 1885–1894, 2017.
- [22] C. Kwak, J. Lee, K. Park, and H. Lee. Let machines unlearn—machine unlearning and the right to be forgotten. In *America's Conference on Information Systems (AMCIS)*, 2017.
- [23] A. Mahadevan and M. Mathioudakis. Certifiable machine unlearning for linear models. *arXiv preprint arXiv:2106.15093*, 2021.
- [24] J. Martens et al. Deep learning via hessian-free optimization. In *ICML*, volume 27, pages 735–742, 2010.
- [25] S. Neel, A. Roth, and S. Sharifi-Malvajerdi. Descent-to-delete: Gradient-based methods for machine unlearning. In *Algorithmic Learning Theory*, pages 931–962, 2021.
- [26] C. Pan, E. Chien, and O. Milenkovic. Unlearning nonlinear graph classifiers in the limited training data regime. *arXiv preprint arXiv:2211.03216*, 2022.
- [27] S. L. Pardo. The california consumer privacy act: Towards a european-style privacy regime in the united states. *J. Tech. L. & Pol'y*, 23:68, 2018.
- [28] C. Parliament. Personal information protection and electronic documents act. *Consolidated Acts, SC 2000*, c. 5:13, 2000.
- [29] B. A. Pearlmutter. Fast exact multiplication by the hessian. *Neural computation*, 6(1):147–160, 1994.
- [30] R. Pytlak. *Conjugate gradient algorithms in nonconvex optimization*, volume 89. Springer Science & Business Media, 2008.
- [31] G. D. P. Regulation. General data protection regulation (gdpr). *Intersoft Consulting, Accessed in October*, 24(1), 2018.
- [32] 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*, 34, 2021.
- [33] P. Sen, G. M. Namata, M. Bilgic, L. Getoor, B. Gallagher, and T. Eliassi-Rad. Collective classification in network data. *AI Magazine*, 29(3):93–106, 2008.
- [34] O. Shchur, M. Mumme, A. Bojchevski, and S. Günnemann. Pitfalls of graph neural network evaluation. *arXiv preprint arXiv:1811.05868*, 2018.
- [35] T. Steihaug. The conjugate gradient method and trust regions in large scale optimization. *SIAM Journal on Numerical Analysis*, 20(3):626–637, 1983.
- [36] K. K. Thekumparampil, C. Wang, S. Oh, and L.-J. Li. Attention-based graph neural network for semi-supervised learning. *arXiv preprint arXiv:1803.03735*, 2018.
- [37] F. Wu, A. Souza, T. Zhang, C. Fifty, T. Yu, and K. Weinberger. Simplifying graph convolutional networks. In *International conference on machine learning*, pages 6861–6871. PMLR, 2019.
- [38] K. Wu, Y. Ning, J. Shen, T. Wang, and W. H. Wang. Certified edge unlearning for graph neural networks (full version). https://github.com/kunwu522/certified_edge_unlearning/blob/main/full_paper/Certified_Edge_Unlearning_for_Graph_Neural_Network_full_version.pdf, 2023.
- [39] Y. Wu, E. Dobriban, and S. Davidson. Deltagrad: Rapid retraining of machine learning models. In *International Conference on Machine Learning*, pages 10355–10366, 2020.
- [40] K. Xu, W. Hu, J. Leskovec, and S. Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019.
- [41] Z. Yang, W. Cohen, and R. Salakhudinov. Revisiting semi-supervised learning with graph embeddings. In *International conference on machine learning*, pages 40–48, 2016.
- [42] C. Zhang, D. Song, C. Huang, A. Swami, and N. V. Chawla. Heterogeneous graph neural network. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 793–803, 2019.