A hybrid framework for effective and efficient machine unlearning

Mingxin Li[1], Yizhen Yu[1], Ning Wang[2], Zhigang Wang[*2],
Xiaodong Wang[1], Haipeng Qu[1], Jia Xu[2,3], Shen Su[2], and Zhichao Yin[4]

{limingxin,yuyizhen}@stu.ouc.edu.cn,
{wangninggzu,wangzhiganglab}@gmail.com
{wangxiaodong,quhaipeng}@ouc.edu.cn,{xujia,sushen}@gzhu.edu.cn,
yinzhichao@cvte.com

[1]Ocean University of China, China
[2]Cyberspace Institute of Advanced Technology, Guangzhou University, China
[3]Key Laboratory of Cyberspace Security Defense, Institute of Information
Engineering, Chinese Academy of Sciences, China
[4]Guangzhou ShiZhen Information Technology Co., LTD, China

Abstract. Recently machine unlearning (MU) is proposed to remove the imprints of revoked samples from the already trained model parameters, to solve users' privacy concern. Different from the runtime expensive retraining from scratch, there exist two research lines, exact MU and approximate MU with different favorites in terms of accuracy and efficiency. In this paper, we present a novel *hybrid strategy* on top of them to achieve an overall success. It implements the unlearning operation with an acceptable computation cost, while simultaneously improving the accuracy as much as possible. Specifically, it runs reasonable unlearning techniques by estimating the retraining workloads caused by revocations. If the workload is lightweight, it performs retraining to derive the model parameters consistent with the accurate ones retrained from scratch. Otherwise, it outputs the unlearned model by directly modifying the current parameters, for better efficiency. In particular, to improve the accuracy in the latter case, we propose an optimized version to amend the output model with lightweight runtime penalty. We particularly study the boundary of two approaches in our frameworks to adaptively make the smart selection. Extensive experiments on real datasets validate that our proposals can improve the unlearning efficiency by 1.5× to 8× while achieving comparable accuracy.

Keywords: Exact machine unlearning · Approximate machine unlearning · Unlearning efficiency · Model accuracy.

1 Introduction

Machine learning (ML) has gained significant success in numerous applications where model parameters are iteratively refined to learn knowledge from training samples. Many efforts have been devoted into accelerating such complex

---

*Corresponding Author.

training computations [28,29]. On the other hand, the output model can make predictions for newly fed testing samples, i.e., inference. However, with the ever growing number of end-users, some of them as attackers, attempt to illegally deduce personal information embedded in training samples. Such attacks are usually performed by analyzing the inference output [25,6,20,23]. That is possible because the model parameters are capable of memorizing sensitive samples during training. This poses a new yet heavy threat to the privacy of samples' owners, which cannot be easily handled by classic privacy protection techniques, such as differential privacy [27]. Thus, recently many laws have been emergently promulgated for privacy protection, such as the European Union's General Data Protection Regulation (GDPR) [21] and the United States' California Consumer Privacy Act (CCPA) [1]. They stipulate the model provider fulfill the obligation to eliminate any impact of a sample on the trained model, once its owner requests to be forgotten, which is known as the "right-to-be-forgotten". This spawns a new research field—machine unlearning (MU) [3], to make ML services compliant with the law constraints.

The most straightforward way to achieve the goal of MU is to delete the corresponding samples from the original dataset and then retrain from scratch. However, as the dataset size increases and the number of unlearning requests grows, retraining from scratch is significantly runtime expensive. Therefore, many efforts have been devoted into accelerating MU while simultaneously ensuring the prediction accuracy of the unlearned model.

Bourtoule et al. propose a representative MU framework called SISA (Sharded-Isolated-Sliced-Aggregated) [3] to strictly remove the impact of revoked samples (termed as **extract MU**). It firstly partitions samples into several subsets, and then sequentially uses them in training. Notably, before samples in a new subset are used, SISA archives the currently trained model parameters as a checkpoint. As a result, once any sample in this subset is revoked, SISA can rollback parameters to this checkpoint to completely remove its impact, while the training work before this checkpoint will not be repeatedly run. That effectively confines retraining workloads. However, the practical performance heavily depends on when the revoked sample begins to participate in training. In the worst case where it is used at the very beginning, SISA degrades to retraining from scratch. Although analyzing the revocation probability of different samples can guide to optimize the participation order and then improve MU efficiency, this prior knowledge cannot be easily obtained. Following SISA, many researchers study how to effectively partition data [26,12,8,7], especially utilizing existing semantic relationship like the graph structure [8]. These works can partially mitigate the negative impact mentioned above but are still far from ideal.

Unlike exact MU, recently some researchers attempt to directly deduce the unlearned model which is similar to that by retraining from scratch [14,15,11,18,19]. They achieve this goal by estimating the contribution from revoked data and then directly subtracting it from already trained model. That significantly improves the MU efficiency but only provides a statistical guarantee on the privacy pro-

tection, since the contribution estimation perhaps is inaccurate. These methods are called as **approximate MU**.

Thereby paper proposes a *hybrid strategy* on top of exact and approximate MU methods, to strike a good balance between accuracy and efficiency. It selects appropriate unlearning strategies dynamically with an acceptable computation cost. Specifically, when the scale of the training process affected by the revoked sample is small, we adopt the *partial retraining strategy*(PRS) [3] directly. We load the previously preserved parameters and then retrain after removing the target data. This ensures the accuracy of the model. If the scale is beyond the acceptable computation cost, we resort to the *direct parameters update strategy*(DPUS) [13] to achieve approximate unlearning. It subtracts the incremental changes caused by the batch of target data from the final parameters. The scale of the training process affected by the revoked sample is determined by the position at which the data participates in the model training process. Through the position, one can infer the magnitude of the computational cost for retraining the affected parts. Computational cost refers to the amount of data read during the training process. HS can dynamically choose the appropriate strategy based on the position of the revoked data in the training process and the computational overhead of the model.

However, model providers typically need to frequently respond to unlearning requests, which may trigger DPUS multiple times, and each time requires subtracting the increment of an entire batch of data from the final parameters. The accumulated error definitely causes heavy accuracy degradation. Thus, we further optimize our *hybrid strategy*. Specifically, we first subtract the influence of the revoked data from the specified parameters which is calculated by the training cost. And then retrain the model based on the basis of the modified parameters to fine-tune and correct the model's output, ensuring that the final model closely approximates one that was never trained on the revoked data.

Lastly, We conduct extensive experiments on four datasets, using unlearning time and model accuracy as evaluation metrics. We demonstrate that our methods can improve unlearning efficiency by $1.5\times$ to $8\times$ across different datasets while maintaining model accuracy. Then we validate the effectiveness of our unlearning method using membership inference attacks [23].

Organization. Section 2 reviews existing studies about the problem this paper focuses on. Section 3 introduces the detailed design of our proposals for *hybrid strategy* and the corresponding *optimized hybrid strategy*. Section 4 evaluates the usefulness of our proposals, and Section 5 finally concludes the paper.

2 Related Work

The concept of machine unlearning was initially proposed by [5], referring to erasing traces of specified data from a trained model. Machine unlearning algorithms are generally classified into exact machine unlearning(**extract MU**) and approximate machine unlearning(**approximate MU**), both of which will be elaborated on below.

Exact MU. It aims to completely remove the instances of requested deletion from the learned model. The simplest approach is to retrain the entire model from scratch after removing the revoked data. However, as the number of requests and the size of the dataset increase, this method incurs significant resource overhead. Hence, existing works focus on accelerating the unlearning process. Ref. [5] investigated machine unlearning within the framework of statistical query learning. It efficiently eliminates the influence of revoked samples by transforming the learning algorithm into a summation form and subtracting the corresponding summation terms associated with the revoked samples. Ref. [10] proposed an efficient unlearning method based on the divide-and-conquer principle for the K-means algorithm. Ref. [4] studied unlearning for random forests. However, the above studies are relatively specific and cannot be extended to other machine learning models. The famous unlearning framework proposed by [3] is SISA, inspired by ensemble learning. It is specifically designed for learning algorithms and operates iteratively. It divides the training dataset into multiple disjoint shards, each used to train a sub-model. Furthermore, each shard is further divided into slices to participate in multiple iterations of the learning algorithm. During the inference and prediction, test samples are inputted into each sub-model, and the responses of all trained models are aggregated. This process is similar to the currently distributed strategies [26,12]. The main advantage of SISA lies in that a sample only affects one sub-model. This means that it's only necessary to repeat training from the iteration containing the revoked samples to the final iteration. It only retrains the affected sub-models, effectively reducing the workload of retraining.

Approximate MU. It aims to relax effectiveness and authentication requirements [24] of unlearning. It provides a statistical guarantee that an unlearned model cannot be distinguished from a model that has never been trained on the deleted data [14]. They usually employ gradient-based update strategies to approximately eliminate the influence of the data to be revoked [22]. Currently, many works [14,22,11,15,30] link approximate unlearning to differential privacy to provide unlearning guarantees. Ref. [14] introduced certified unlearning, which first introduced the statistical influence function [19] into machine unlearning. Specifically, the impact of training data on model parameters can be represented by the product between the model's first-order derivative concerning the revoked data and the model's Hessian inverse matrix for the remaining data. When unlearning, it is only to subtract this term from the model parameters. Additionally, to ensure that the unlearned model has no difference in the output space compared to the retrained model, they defined ε-certified unlearning based on the concept of differential privacy [18]. They proved that their method satisfies certified unlearning. Some other studies propose unlearning by caching parameters of the model training process, such as caching the contribution of the data to be revoked [13]. Then subtract the increment of the batch containing the revoked data from the final parameters.

The best-known exact unlearning method [3] is extracted but still relatively expensive to unlearn. On the other hand, ref. [13] is the most efficient option
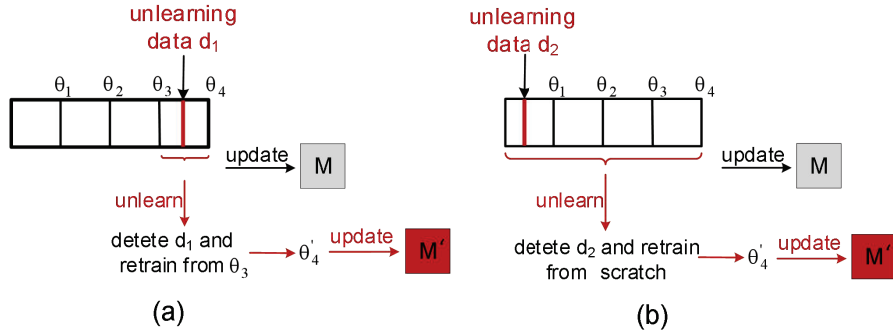
Fig. 1. An example of SISA.

in approximate unlearning, although it may have some impact on the model performance. After considering the advantages and disadvantages of these two methods, we propose a *hybrid strategy* and *optimized hybrid strategy* that aim to fully utilize their advantages.

3   Proposed Methods

In this paper, we first propose a novel machine unlearning method called *hybrid strategy* (HS), to improve the efficiency of SISA [3]. SISA [3] is described in detail in the section 2. To simplify the model, the number of shards is set to 1, and only the internal slice division of the dataset is considered. We refer to the unlearning method in SISA as the *partial retraining strategy* (PRS), as shown in Fig. 1. During the training, the corresponding parameters are saved after the training of each slice. When the unlearning request arrives, its position is first determined. In Fig. 1(a), d_1 is located in the last slice, and then retraining from θ_3 after deleting d_1, saves the retraining overhead. However, as shown in Fig. 1(b), when the deleted data d_2 is located in the first slice, triggering retraining for

---

Algorithm 1 Training Algorithm in Hybrid Strategy

Input:Training dataset D, the number of slices S, threshold for triggering partial retraining slice l.

Output: Training model M.

1: Partition datasets into S disjoint slices $D_1, D_2, ...D_S$.
2: **for** each D_i **do**
3:    Train model based on D_i to get parameters θ_i.
4:    **if** $i < l$ **then**
5:       Save the increments Δ_j for the corresponding batch.
6:    **end if**
7: **end for**
8: Return M with its parameter θ_S.

---

the entire model becomes necessary. Although SISA considers partitioning the dataset into multiple shards, this still comes at the cost of compromising model accuracy. With a low number of shards, triggering retraining for the entire sub-model remains a considerable overhead.

Therefore, we introduce another strategy, the *direct parameters update strategy*(DPUS) [13], in cases where partial retraining cost is unacceptable. It achieves approximate unlearning by directly subtracting the parameter changes caused by the batch in which the revocation data is located, thereby improving unlearning efficiency. When the retraining cost is acceptable, i.e., when unlearning requests occur towards the end of a slice, we directly adopt the PRS. *Hybrid strategy* (see Fig. 2) combines the PRS and the DPUS, dynamically selecting the appropriate strategy to reduce unlearning overhead.

Furthermore, if DPUS is frequently triggered, it requires multiple updates to the model parameters, which can degrade the model's accuracy. We further optimize the *hybrid strategy* and propose the *optimized hybrid strategy* (OHS). In cases of unacceptable retraining overhead, after applying the DPUS, we still perform retraining to achieve convergence. Section 3.1 introduces algorithm design for *hybrid strategy* and section 3.2 shows our optimizations for it.

3.1   Hybrid Strategy

The *hybrid strategy* consists of a training algorithm and an unlearning algorithm. Alg. 1 shows the detailed training process. Firstly, the dataset D is partitioned into S disjoint slices. During training, incremental training is sequentially performed on each slice. When the completion of training on each slice, the corresponding parameters θ_i are saved (Lines 2-3). We set a slice threshold l where the current slice index is less than l, meaning that the computation cost is unacceptable. Conversely, the computation cost is acceptable. l is related to the acceptable retraining overhead, which is calculated as follows.

Assuming that the size of dataset is n and the tolerable retraining overhead of the model is ϕ. The number of slices is S, then the amount of data contained in each slice is n/S. Since model training is incremental, we assume the retraining starts from the i-th slice. The maximum value of i can be calculated by (1) is m, and $l = \lfloor S - m \rfloor$.

$$i \cdot n/S + (i+1) \cdot n/S + ... + S \cdot n/S \leq \phi \tag{1}$$

When the index i of the training slice is less than l, the corresponding parameter increments Δ_j for each batch are saved (Lines 4-5). Finally, the trained model M is returned.

The unlearning algorithm is shown in Alg. 2, assuming that the revocation data index is denoted as d. Iterating through the data indices D_i contained in each slice, if d belongs to D_i, then take the corresponding unlearning measures. If $i < l$ (i.e., the overhead of retraining from D_i is unacceptable), the DPUS is adopted. It involves subtracting the increment Δ_j caused by the batch containing d from the final parameters θ_S to update θ'_S. If $i \geq l$ (i.e., the overhead of
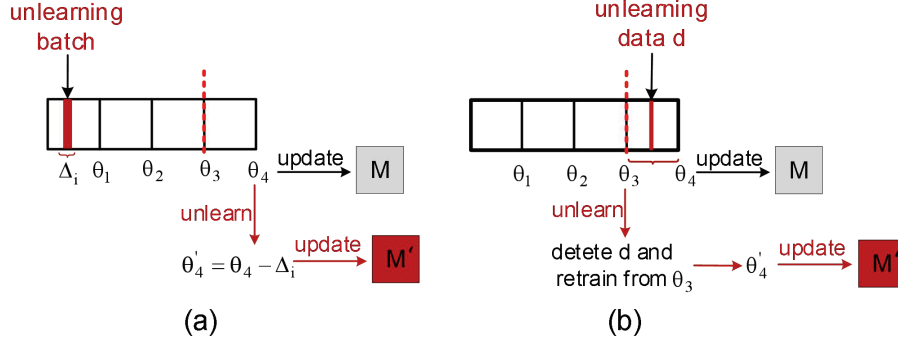
Fig. 2. An example of *hybrid strategy*.

retraining from D_i is acceptable), the parameters trained from the previous slice are directly loaded. Subsequently, d is deleted, and the final model parameters θ'_S are obtained by retraining based on $\theta_{(i-1)}$ (Lines 3-6). Finally, the model is updated with the θ'_S to obtain the model M'.

Fig. 2 shows a specific example of the *hybrid strategy*, where the dataset is divided into four slices, and l is three. After training each slice, the corresponding parameters are saved, and the parameter increments for each batch are saved for the first three slices. The parameters θ_4 obtained from the last slice are used to update the model M. As shown in Fig. 2(a), when the revocation data is located in the first three slices, using PRS would incur excessive retraining overhead. Therefore, the DPUS is employed, subtracting the parameter increments caused by the revocation data batch directly from the final parameters to obtain θ'_4. And use θ'_4 to update the model M'. Fig. 2(b) demonstrates that when the revocation data is in the last slice, *partial retraining* can be directly applied by deleting d and retraining from θ_3 to obtain the final unlearning model. From this

---

Algorithm 2 Unlearning Algorithm in *Hybrid Strategy*

Input:sub-datasets$\{D_1, D_2, ..., D_S\}$, revocation data d.

Output: Unlearning model M'.

1: **for** i from 1 to S **do**
2:     **if** $d \in D_i$ **then**
3:         **if** $i < l$ **then**
4:             $\theta'_S \leftarrow \theta_S - \Delta_j$ where Δ_j is the increment of corresponding batch of d.
5:         **else**
6:             delete d from D_i and retrain from $\theta_{(i-1)}$ to get θ'_S.
7:         **end if**
8:     **end if**
9: **end for**
10: Return M' with its parameter θ'_S.

Algorithm 3 Unlearning Algorithm in *Optimized Hybrid Strategy*

Input:sub-datasets$\{D_1, D_2, ..., D_S\}$, revocation data d.

Output: Unlearning model M'.

1: **for** i from 1 to S **do**
2:    **if** $d \in D_i$ **then**
3:       **if** $i < l$ **then**
4:          $\theta'_{(S-l)} \leftarrow \theta_{(S-l)} - \Delta_j$.
5:          retrain from $\theta'_{(S-l)}$ to get θ'_S.
6:       **else**
7:          delete d from D_i and retrain from $\theta_{(i-1)}$ to get θ'_S.
8:       **end if**
9:    **end if**
10: **end for**
11: Return M' with its parameter θ'_S.

---

example, we can observe that the *hybrid strategy* dynamically selects appropriate strategies for unlearning based on the retraining overhead.

3.2 Optimized Hybrid Strategy

The *hybrid strategy* chooses unlearning strategies based on computational overhead through a simple dynamic selection. However, in real life, there is unlikely to be only one unlearning request, which means that service providers have to respond to numerous unlearning requests. Adopting the HS strategy will frequently trigger DPUS. It essentially records the contribution of the corresponding batches of data to the model. Multiple unlearning requests cannot be concentrated in the same batch, which means that during the unlearning process, the parameters corresponding to numerous batches will be subtracted, approximating the process as if the model was never trained on those batches of data. This can potentially lead to a decrease in model accuracy. Therefore, we further optimize the DPUS within *hybrid strategy*, proposing the *optimized hybrid strategy*.

    However, it is very difficult to reduce the unlearning time as much as possible and also ensure the model's usability. To ensure that the model accuracy is not compromised, we make some concessions on the unlearning time. The training algorithm is the same as HS, with the difference lying in the unlearning algorithm. As shown in Alg. 3, OHS does not involve subtracting the corresponding increments from the final trained parameters. Instead, based on the l, it subtracts the corresponding increments from the parameters at specified positions and then retrains l slices based on $\theta'_{(S-l)}$ to bring the model to a converged state (Lines 3-5). We ensure that the computation is acceptable and the model accuracy is accurate. Because retraining can correct the model's outputs to a certain extent, it ensures the model's usability.

    Fig. 3 illustrates a concrete example where the size of l is one. When the overhead of partial retraining is high, we subtract the corresponding increments

Fig. 3. An example of *optimized hybrid strategy*

from θ_3 to obtain θ_3' and then retrain the last slice based on θ_3' to obtain the unlearning model. When the retraining overhead is acceptable, the *optimized hybrid strategy* behaves the same as the *hybrid strategy*, directly adopting the PRS. The optimization aims to ensure the usability of the model regardless of the chosen strategy while maintaining unlearning efficiency.

4 Experiments

In this section, we evaluate our proposals to show their effectiveness. Section 4.1 introduces four datasets we used. Section 4.2 shows the two comparative methods. Section 4.3 describes experimental equipment, model architecture, and model parameters. Section 4.4 evaluates our methods.

4.1 Datasets

We use 4 real datasets *Adult* [2], *Purchase* [17], *Census* [9] and *Olympics* [16] to test our proposals and competitors.

- *Adult* [2] has 48,842 samples and 111 features. It aims to predict whether an individual's annual income exceeds $50,000 based on their attributes.
- *Purchase* [17] has 311,519 samples and 600 features, aiding in the understanding of consumer behavior and purchasing patterns.
- *Census* [9] has 299,285 samples and 432 features, serves as a powerful resource for analyzing demographic trends and socioeconomic patterns
- *Olympics* [16] has 206,165 samples and 1,016 features, and serves as a comprehensive resource for analyzing trends, patterns, and outcomes within the realm of Olympic sports.

4.2 Competitors

There are currently many methods for machine unlearning, and this paper selects the most representative methods from exact unlearning and approximate

unlearning for comparison, namely SISA[3] and DPUS[13]. SISA is one of the most efficient methods for achieving exact unlearning. Many other methods are further extensions based on SISA. A significant advantage of SISA is its wide applicability, as it is not limited to specific algorithms and is therefore widely adopted in various scenarios. SISA implements partial retraining through sharding and incremental training, significantly reducing training overhead and thus achieving efficient unlearning. DPUS (Direct Parameter Update Strategy) is one of the most efficient methods in approximate unlearning. It requires only simple subtraction operations to execute the unlearning process, avoiding complex computational procedures. In contrast, many approximate unlearning methods require computing the information of the Hessian inverse matrix, which incurs substantial computational overhead. Although SISA performs exceptionally well in exact unlearning, its unlearning operations have higher computational overhead compared to approximate unlearning methods. Its greatest advantage lies in its high accuracy. On the other hand, while DPUS is very efficient, it can have some impact on model accuracy. Therefore, this paper chooses these two representative methods as baselines.

4.3   Setup

All experiments were performed on a Dell Precision server equipped with a NVIDIA GeForce RTX 3090 GPU (10496 CUDA Cores @1.70GHz, 24GB HBM2), and two Intel(R) Xeon(R) Silver 4216 Processors (2.10GHz, 64 GB memory). For all datasets, we train a binary classification task based on a Multilayer Perceptron with two hidden layers, each containing 128 neurons. During training and testing, the batch size is 128, the learning rate is 0.005 and the optimizer is adam. Both our methods and competitors utilize the same model architecture and the same hyperparameters.

4.4   Experimental Evaluation

We use unlearning time and model accuracy after unlearning as evaluation metrics. Specifically, both our methods and the comparative methods sequentially handle 100 randomly generated unlearning requests to simulate multiple users submitting revocation requests to the service provider. Then calculate the average unlearning time for each request as the unlearning efficiency – clearly, shorter times indicate higher efficiency. Model accuracy after processing all unlearning requests serves as another metric. Finally, we validate method effectiveness using membership inference attacks [23], verifying whether traces of unlearning data are thoroughly erased from the model.

Fig. 4 displays the difference in accuracy between two baselines and our methods after unlearning. It's clear that model accuracy is the poorest when using DPUS. This is because DPUS directly subtracts the corresponding batch's parameter increments with each unlearning request, leading to significant degradation in model accuracy as more requests accumulate. The subtraction of these
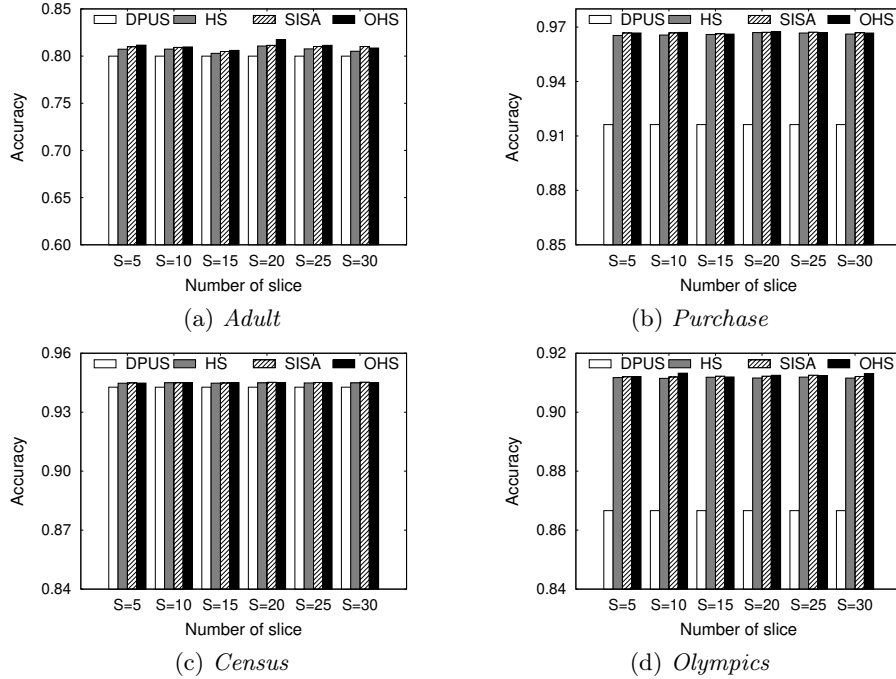
Fig. 4. The comparison of accuracy among *direct parameters update strategy*(DPUS), *hybrid strategy*(HS), *partial retraining*(SISA) and *optimized hybrid strategy*(OHS) across the four datasets.

increments means that the model loses not only the specific unlearned data's contributions but also the overall batch information, which cumulatively impacts the model's accuracy.

In contrast, the HS method performs better than DPUS in terms of accuracy. Instead of subtracting increments directly for each unlearning, HS intermittently triggers partial retraining. This retraining allows the model to recalibrate and reach a more stable state, which mitigates the negative impact on accuracy. However, HS's accuracy is slightly lower than SISA because SISA ensures partial retraining with every unlearning request, maintaining a higher level of model stability and accuracy.

Specifically, DPUS can be up to 5% less accurate than HS on different datasets, while HS can be up to 0.5% less accurate than SISA. To address the accuracy gap observed with HS, we propose OHS, which improves accuracy to levels comparable with SISA. This demonstrates that OHS effectively balances the need for efficient unlearning with the maintenance of high model accuracy.

Fig. 5 illustrates the comparison of unlearning time among DPUS, SISA, HS, and OHS. The unlearning time of DPUS is the shortest because it doesn't need to retrain the model. Both HS and OHS show significantly shorter unlearning times than SISA. HS improves unlearning efficiency by a factor of 6 × to 8 × compared to SISA, while OHS shows an improvement of 1.5 × to 6 ×. This

(a) *Adult*   (b) *Purchase*
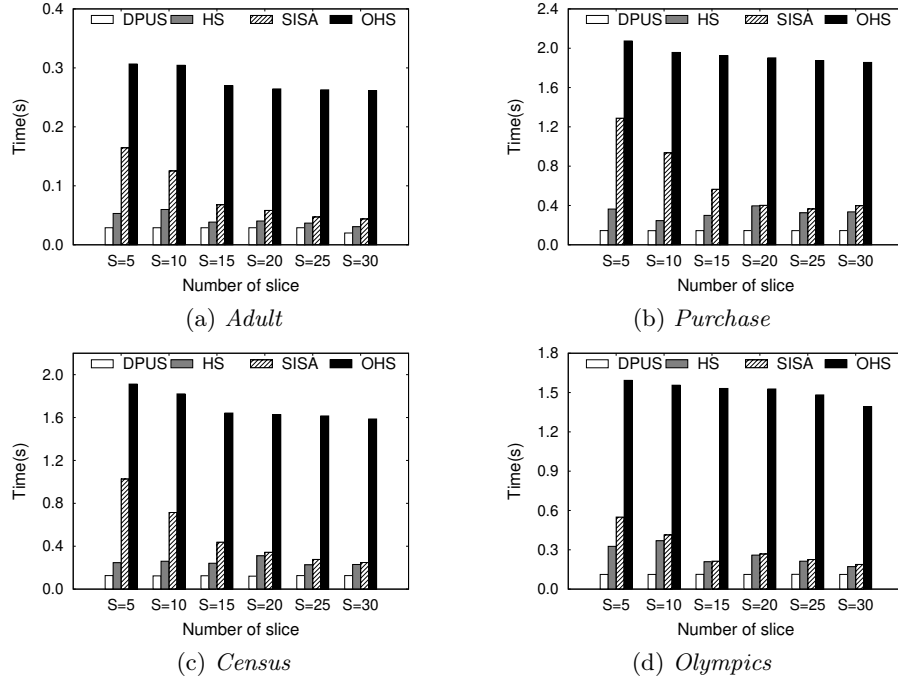
(c) *Census*   (d) *Olympics*

Fig. 5. The comparison of unlearning time among *direct parameters update*(DPUS) , *partial retraining*(SISA), *hybrid strategy*(HS) and *optimized hybrid strategy*(OHS) across the four datasets.

efficiency gain arises because our methods do not always trigger the retraining of affected parts as SISA does with each unlearning request. Instead, HS and OHS dynamically choose the most efficient strategy, minimizing retraining overhead when it is too costly.

Additionally, we observe a decreasing trend in retraining time as the number of slices increases. With more slices, there are more intermediate parameters saved, which means during retraining, it's unnecessary to load parameters unrelated to the unlearning data. This selective loading process further enhances the unlearning efficiency, allowing the model to effectively and quickly remove the contributions of specific data without a full retraining cycle.

Finally, we validate whether revocation data participates in model training using membership inference attacks [23]. Specifically, multiple shadow models are trained by randomly dividing the dataset, where half of the data participates in training while the other half does not. The models perform differently on trained and untrained data. By using the predictions generated by inputting all data into the shadow models, along with whether they participated in model training, as labels (with participation labeled as "1" and non-participation labeled as "0"), we train an attack model. Subsequently, the attack model can use the predicted results of the revocation data as inputs to infer whether the revocation data participated in training the original model. Finally, through validation, we find

that our methods achieve effective unlearning, as indicated by the attack model predicting "0" for the revocation data, indicating that it does not participate in the training of the model.

5 Conclusion

In this paper, we focus on improving unlearning efficiency while ensuring model utility. We propose a *hybrid strategy* that dynamically selects unlearning strategies to better reduce unlearning overhead. Additionally, considering the shortcomings of the *direct parameters update strategy*, we further optimize it and propose an *optimized hybrid strategy*. We validate that our methods can significantly enhance unlearning rates without compromising model performance through experiments. Finally, the effectiveness of our methods is confirmed through membership inference attacks.

References

1. California consumer privacy act (ccpa). State of California Office of the Attorney General (2024), https://oag.ca.gov/privacy/ccpa.
2. Becker, B., Kohavi, R.: Adult. `http://archive.ics.uci.edu/ml` (1996)
3. Bourtoule, L., Chandrasekaran, V., Choquette-Choo, C.A., Jia, H., Travers, A., Zhang, B., Lie, D., Papernot, N.: Machine unlearning. In: 2021 IEEE Symposium on Security and Privacy (SP). pp. 141–159. IEEE (2021)
4. Brophy, J., Lowd, D.: Machine unlearning for random forests. In: International Conference on Machine Learning. pp. 1092–1104. PMLR (2021)
5. Cao, Y., Yang, J.: Towards making systems forget with machine unlearning. In: 2015 IEEE symposium on security and privacy. pp. 463–480. IEEE (2015)
6. Carlini, N., Liu, C., Erlingsson, Ú., Kos, J., Song, D.: The secret sharer: Evaluating and testing unintended memorization in neural networks. In: 28th USENIX security symposium (USENIX security 19). pp. 267–284 (2019)
7. Chen, C., Sun, F., Zhang, M., Ding, B.: Recommendation unlearning. In: Proceedings of the ACM Web Conference 2022. pp. 2768–2777 (2022)
8. Chen, M., Zhang, Z., Wang, T., Backes, M., Humbert, M., Zhang, Y.: Graph unlearning. In: Proceedings of the 2022 ACM SIGSAC conference on computer and communications security. pp. 499–513 (2022)
9. Dua, D., Graff, C.: Uci machine learning repository. `https://archive.ics.uci.edu/ml/machine-learning-databases/census-income-mld`
10. Ginart, A., Guan, M., Valiant, G., Zou, J.Y.: Making ai forget you: Data deletion in machine learning. Advances in neural information processing systems **32** (2019)
11. Golatkar, A., Achille, A., Soatto, S.: Eternal sunshine of the spotless net: Selective forgetting in deep networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9304–9312 (2020)

12. Gratton, C., Venkategowda, N.K., Arablouei, R., Werner, S.: Privacy-preserved distributed learning with zeroth-order optimization. IEEE Transactions on Information Forensics and Security **17**, 265–279 (2021)
13. Graves, L., Nagisetty, V., Ganesh, V.: Amnesiac machine learning. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 11516–11524 (2021)
14. Guo, C., Goldstein, T., Hannun, A., Van Der Maaten, L.: Certified data removal from machine learning models. arXiv preprint arXiv:1911.03030 (2019)
15. Gupta, V., Jung, C., Neel, S., Roth, A., Sharifi-Malvajerdi, S., Waites, C.: Adaptive machine unlearning. Advances in Neural Information Processing Systems **34**, 16319–16330 (2021)
16. Kaggle.: 120 years of olympic history: Athletes and events. https://www.kaggle.com/heesoo37/120-years-of-olympic-history-athletes-and-results
17. Kaggle.: Acquire valued shoppers challenge. https://www.kaggle.com/competitions/acquire-valued-shoppers-c
18. Koh, P.W., Liang, P.: Understanding black-box predictions via influence functions. In: International conference on machine learning. pp. 1885–1894. PMLR (2017)
19. Koh, P.W., Liang, P.: Understanding black-box predictions via influence functions. In: International conference on machine learning. pp. 1885–1894. PMLR (2017)
20. Li, Z., Zhang, Y.: Membership leakage in label-only exposures. In: Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security. pp. 880–895 (2021)
21. Mantelero, A.: The EU proposal for a general data protection regulation and the roots of the 'right to be forgotten'. Comput. Law Secur. Rev. **29**(3), 229–235 (2013)
22. Neel, S., Roth, A., Sharifi-Malvajerdi, S.: Descent-to-delete: Gradient-based methods for machine unlearning. In: Algorithmic Learning Theory. pp. 931–962. PMLR (2021)
23. Shokri, R., Stronati, M., Song, C., Shmatikov, V.: Membership inference attacks against machine learning models. In: 2017 IEEE symposium on security and privacy (SP). pp. 3–18. IEEE (2017)
24. Thudi, A., Jia, H., Shumailov, I., Papernot, N.: On the necessity of auditable algorithmic definitions for machine unlearning. In: 31st USENIX Security Symposium (USENIX Security 22). pp. 4007–4022 (2022)
25. Tramèr, F., Zhang, F., Juels, A., Reiter, M.K., Ristenpart, T.: Stealing machine learning models via prediction {APIs}. In: 25th USENIX security symposium (USENIX Security 16). pp. 601–618 (2016)
26. Van Huynh, N., Hoang, D.T., Nguyen, D.N., Dutkiewicz, E.: Joint coding and scheduling optimization for distributed learning over wireless edge networks. IEEE Journal on Selected Areas in Communications **40**(2), 484–498 (2021)
27. Wang, N., Wang, Y., Wang, Z., Nie, J., Wei, Z., Tang, P., Gu, Y., Yu, G.: Privnud: Effective range query processing under local differential privacy. In: 39th IEEE International Conference on Data Engineering, ICDE 2023, Anaheim, CA, USA, April 3-7, 2023. pp. 2660–2672. IEEE (2023)
28. Wang, Z., Gu, Y., Bao, Y., Yu, G., Yu, J.X., Wei, Z.: Hgraph: I/o-efficient distributed and iterative graph computing by hybrid pushing/pulling. IEEE Trans. Knowl. Data Eng. **33**(5), 1973–1987 (2021)
29. Wang, Z., Tu, Y., Wang, N., Gao, L., Nie, J., Wei, Z., Gu, Y., Yu, G.: FSP: towards flexible synchronous parallel frameworks for distributed machine learning. IEEE Trans. Parallel Distributed Syst. **34**(2), 687–703 (2023)
30. Warnecke, A., Pirch, L., Wressnegger, C., Rieck, K.: Machine unlearning of features and labels. arXiv preprint arXiv:2108.11577 (2021)