

# MultiDelete for Multimodal Machine Unlearning

Jiali Cheng  Hadi Amiri

University of Massachusetts Lowell  
{jiali\_cheng, hadi\_amiri}@uml.edu

**Abstract.** Machine Unlearning removes specific knowledge about training data samples from an already trained model. It has significant practical benefits, such as purging private, inaccurate, or outdated information from trained models without the need for complete re-training. Unlearning within a multimodal setting presents unique challenges due to the complex dependencies between different data modalities and the expensive cost of training on large multimodal datasets and architectures. This paper presents the first machine unlearning approach for *multimodal* data and models, titled MULTIDELETE, which is designed to decouple associations between unimodal data points during unlearning without losing the overall representation strength of the trained model. MULTIDELETE advocates for three key properties for effective multimodal unlearning: (a): modality decoupling, which effectively decouples the association between individual unimodal data points marked for deletion, rendering them as unrelated data points, (b): multimodal knowledge retention, which retains the multimodal representation post-unlearning, and (c): unimodal knowledge retention, which retains the unimodal representation post-unlearning. MULTIDELETE is efficient to train and is not constrained by using a strongly convex loss—a common restriction among existing baselines. Experiments on two architectures and four datasets, including image-text and graph-text datasets, show that MULTIDELETE gains an average improvement of 17.6 points over best performing baseline in unlearning multimodal samples, can maintain the multimodal and unimodal knowledge of the original model post unlearning, and can provide better protection to unlearned data against adversarial attacks<sup>1</sup>.

**Keywords:** Multimodal · Machine Unlearning

## 1 Introduction

Multimodal models are used in a variety of real-world applications [32, 35, 36, 53], which require extensive training on large datasets. However, the underlying training data can be subject to change due to various reasons such as copyright issues, users revoking consent, and data becoming outdated. Continuing to use a model trained on such data poses significant risks to privacy and questions the model’s relevance and accuracy. Machine Unlearning addresses these challenges

---

<sup>1</sup> Code and data is available at <https://github.com/CLU-UML/MultiDelete>

by removing specific knowledge of training samples from an already trained model, while preserving the model’s functionality on its downstream tasks.

Despite recent advances in machine unlearning in *unimodal* settings [24, 25, 29, 40, 48], unlearning in multimodal settings is largely unexplored. Multimodal unlearning is a challenging task due to the complex relationship and dependency among individual data modalities and the complexity inherent in multimodal data and model architectures. To the best of our knowledge, no existing approach is specifically designed for unlearning in the context of multimodal data, and existing unimodal approaches may not be directly applicable or effective on multimodal data. Specifically, existing *weight-scrubbing* methods add noise to model weights [24, 48, 49], which may fall short in full unlearning of the inter-modality dependencies, which is a key requirement for eliminating any residual data traces. *Certified removal* methods assume convexity of training objectives [13, 16], which often does not hold in multimodal settings. *Optimization-based* methods focus on unimodal settings or last layer models [29, 40, 43], which are less effective on multimodal data because of the cross-modality interactions that can occur beyond just the last layer. Finally, *efficient retraining methods* require significant training cost and can lead to overfitting [2, 7, 20].

In this paper, we take the initial step to investigate the multimodal unlearning problem. Our approach to formulating multimodal unlearning is titled `MultiDelete` and centers on developing a model that satisfies three key properties: (a) modality decoupling which reduces the dependencies between modalities for data samples marked for deletion, (b) multimodal knowledge retention which preserves the previously learned multimodal knowledge of the model during unlearning, and (c) unimodal knowledge retention which retains the previously learned unimodal knowledge. We formally define these properties and design specific loss functions for effective and efficient multimodal unlearning. We summarize our contributions as follows:

- we conceptualize multimodal unlearning through three pivotal properties: modality decoupling, multimodal knowledge retention, and unimodal knowledge retention, which collectively enable multimodal unlearning while preserving the essential knowledge previously learned by the model, and
- through extensive experiments across image-text and graph-text datasets, and architectures, we show the efficacy of the proposed properties and approach in providing protection to deleted data and maintaining model’s robustness against adversarial attacks.

Experimental results show the efficacy of `MultiDelete` measured by standard performance metrics across different multimodal tasks, as well as efficiency versus retraining from scratch. Compared to the best-performing baseline, `MultiDelete` obtains superior performance advantage of 17.6 in forgetting deleted data, while maintains the previously learned knowledge with 0.3 performance gain across a wide-range of tasks. In addition, `MultiDelete` can better protect the deleted data against membership inference attacks post-unlearning by 0.13 absolute points, and is substantially more efficient than retraining the model from scratch.

## 2 MULTIDELETE

Without loss of generality, we center the presentation of `MultiDelete` on dual data modalities, such as image-text. We note that `MultiDelete` can be applied to more than two modalities or other types of modalities. We demonstrate this broader applicability in experiments using graph-text data.

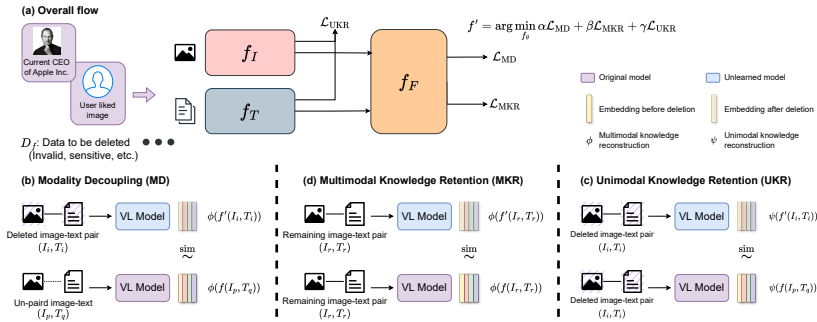
**Notation** Consider a vision-language model  $f$  trained on a dataset of  $N$  image-text pairs,  $D_{\text{train}} = \{(I_i, T_i)\}_{i=1}^N$ . We denote  $D_f \in D_{\text{train}}$  as the subset of training data that should be unlearned from  $f$ . Conversely,  $D_r = D_{\text{train}} \setminus D_f$  denotes the remaining data after removing  $D_f$ . We denote  $f'$  as the desired unlearned model, which functions as if  $D_f$  had never been used in training of  $f$ . In addition, we assume that the original model  $f$  can be decomposed into a vision feature extractor  $f_I$ , a language feature extractor  $f_T$ , and a modality fusion module  $f_F$ .

**Problem Formulation** Given a vision-language model  $f$  trained on  $D_{\text{train}}$ , we aim to unlearn a subset of training data  $D_f$  from  $f$  and obtain a corresponding unlearned model  $f'$ . For this purpose, our core objective is to eliminate the influence that any  $(I_i, T_i) \in D_f$  has on the parameters of  $f$ , so that it effectively “forgets” the patterns learned from  $D_f$ . Crucially, we aim to maintain the performance of the unlearned model  $f'$  on the test data  $D_{\text{test}}$  as close as possible to that of the original model  $f$ . This will ensure that while  $f'$  discards specific data knowledge from  $D_f$ , it retains the overall task effectiveness of  $f$  on  $D_r$  samples. Ideally,  $f'$  should effectively forget  $D_f$  without compromising its performance on  $D_{\text{Test}}$ . Here, we assume that the model only needs to unlearn the *relationship* between image-text pairs  $(I_i, T_i) \in D_f$ , but not necessarily the individual unimodal elements,  $I_i$  and  $T_i$  themselves. This approach ensures that the model retains its foundational knowledge of individual modalities, which is essential for effective learning of the target task and prevents the unnecessary loss of information. For example, it allows for forgetting of a user’s “like” on a social media post without requiring to remove either the user or the post from the platform.

### 2.1 Key Properties for Multimodal Unlearning

We outline key properties essential for successful multimodal unlearning. These properties and the overall work flow of the model are depicted in Fig. 1.

**Modality Decoupling** This property requires that for the unlearned model  $f'$ , the relationship between any image-text pair  $(I_i, T_i) \in D_f$  should be indistinguishable from the relationship of any non-related image-text pair in the dataset. In other words,  $f'$  should be unable to discern a removed pair  $(I_i, T_i) \in D_f$  from those that are deemed unassociated in the dataset. This decoupling is pivotal because it effectively prevents the reconstruction or extraction of any specific information about the removed pair  $(I_i, T_i) \in D_f$ . It also helps model’s generalizability and robustness, because residual associations from the removed data



**Fig. 1:** Summary of the proposed approach, MULTIDELETE. **(a)** given a trained multimodal model (e.g. a vision-language model) and a subset of its training data  $D_f$  with image-text *relations* marked for unlearning or deletion, MULTIDELETE decouples the inter-modality dependency on  $D_f$ , while maintaining the multimodal and unimodal knowledge on the remaining training data  $D_r = D_{\text{train}} \setminus D_f$ . **(b)** modality decoupling: ensures that individual modalities in the deleted data pairs  $D_f$  are treated as unrelated by the model, **(c)** multimodal knowledge retention: preserves the model’s overall multimodal knowledge, and **(d)** unimodal knowledge retention: preserves the model’s understanding of individual modalities.

could potentially skew the model’s performance on new data [24]. We formally define this notion of unlearning in the multimodal context as follows:

**Definition 1 (modality decoupling).** Let  $(I_i, T_i) \in D_f$  denote an image-text pair marked for deletion from model  $f$ . The unlearned model  $f'$  achieves *e* effective modality decoupling when  $(I_i, T_i) \in D_f$  becomes indistinguishable from any unrelated image-text pair  $(I_p, T_q), p \neq q$  sourced from  $D_r$ :

$$\mathbb{E}_{(I_i, T_i) \in D_f, (I_p, T_q)_{p \neq q}} \left[ \phi(f'(I_i, T_i)) - \phi(f(I_p, T_q)) \right] = \epsilon, \quad (1)$$

where  $f(\cdot)$  and  $f'(\cdot)$  generate multimodal representations of their inputs,  $\phi$  is a readout function (such as the concatenation operator, applied to a set of representations), and  $\epsilon$  is an infinitesimal constant.

To realize this property, we randomly draw unrelated image text pairs  $(I_p, T_q)$  from  $D_r$ , and minimize the difference in multimodal associations between the image-text pairs  $(I_i, T_i) \in D_f$  and the unassociated image-text pairs  $(I_p, T_q)$ . This is achieved by minimizing the following distance:

$$\mathcal{L}_{\text{MD}} = \text{Dis} \left( \left\{ f'(I_i, T_i) \mid (I_i, T_i) \in D_f \right\}, \left\{ f(I_p, T_q) \mid (I_p, T_p) \in D_r, (I_q, T_q) \in D_r, p \neq q \right\} \right), \quad (2)$$

where  $\text{Dis}(\cdot)$  can be mean squared error. By minimizing this loss, the model is trained to forget or unlearn specific relationships in deleted pairs, making them indistinguishable from unrelated or random data pairs. This is a crucial step in ensuring that the unlearned model does not retain any knowledge about the data it is supposed to forget.

**Multimodal Knowledge Retention** This property requires that, the process of unlearning  $D_f$  does not adversely affect the learned multimodal knowledge on the remaining dataset  $D_r$ . In other words, the multimodal knowledge related to image-text pairs in  $D_r$ , i.e.  $f'(I_r, T_r), \forall (I_r, T_r) \in D_r$ , should preserve the corresponding original knowledge,  $f(I_r, T_r)$ , after the unlearning process. This approach ensures that while specific data pairs are being unlearned, the overall multimodal knowledge and capability of the model remain robust. Formally, we define retention of multimodal knowledge as follows:

**Definition 2** (multimodal knowledge retention). *Let  $(I_r, T_r) \in D_r$  denote an image-text pair that is "not" marked for deletion. The unlearning approach is effective in retaining multimodal knowledge if it minimizes the deviation in the multimodal knowledge between the unlearned model  $f'$  and the original model  $f$ :*

$$\mathbb{E}_{(I_r, T_r) \in D_r} \left[ \phi(f'(I_r, T_r)) - \phi(f(I_r, T_r)) \right] = \epsilon, \quad (3)$$

where the readout function  $\phi$  is a vector combination operator (such as concatenation). We realize this property by minimizing the gap in the multimodal knowledge between  $f'$  and  $f$  as follows:

$$\mathcal{L}_{\text{MKR}} = \text{Dis}\left(f'(I_r, T_r), f(I_r, T_r)\right), (I_r, T_r) \in D_r. \quad (4)$$

**Unimodal Knowledge Retention** This property requires that the individual unimodal representations of the data points  $(I_i, T_i) \in D_f$  remain intact post unlearning. The rationale is that although the inter-modal relationships are unlearned through modality decoupling,  $I_i$  and  $T_i$  are still valid standalone image and text data. Therefore, it is important that their unimodal representations are preserved to retain the unimodal knowledge initially learned by  $f$ , i.e.  $f_I(I)$  for images and  $f_T(T)$  for texts. This property helps maintain the core knowledge of individual modalities, and prevents unnecessary loss of information or the need to relearn basic features from scratch post-unlearning. Formally, we define unimodal knowledge retention as follows:

**Definition 3** (unimodal knowledge retention). *The unlearning process effectively retains the unimodal knowledge if it minimizes the discrepancy between the unimodal representations produced by the unlearned model  $f'$  and the original model  $f$ :*

$$\mathbb{E}_{(I_i, T_i) \in D_f} \left[ \psi(f'_I(I_i), f'_T(T_i)) - \psi(f_I(I_i), f_T(T_i)) \right] = \epsilon, \quad (5)$$

where  $f_I(\cdot)$  and  $f_T(\cdot)$  generate unimodal representations for image and text data respectively, the readout function  $\psi$  is a vector combination operator (such as concatenation), and  $\epsilon$  is an infinitesimal constant. Note that, although Eq. (5) can be applied to all training data  $D_{\text{train}}$ , we only apply it to  $D_f$  samples for efficiency purpose. In fact, we expect unimodal knowledge on remaining data  $D_r$  to be preserved in  $f'$  due to multimodal knowledge retention, discussed above.

Thus, we minimize the following gap to realize unimodal knowledge retention:

$$\mathcal{L}_{\text{UKR}} = \text{Dis} \left( \left\{ [f'_I(I_i); f'_T(T_i)] \mid (I_i, T_i) \in D_f \right\}, \left\{ [f_I(I_i); f_T(T_i)] \mid (I_i, T_i) \in D_f \right\} \right), \quad (6)$$

where  $[\cdot]$  denotes vector concatenation. This loss aims to retain the core unimodal knowledge during training even after unlearning certain relationships.

We note that an alternative approach is to use the fusion module  $f_F$ , while freezing the unimodal encoders  $f_I$  and  $f_T$ . However, this can be limiting, especially for models like CLIP [53], which use nonparametric fusion modules (e.g., dot product) for modality interaction. There is also a risk that an adversarial agent might exploit the original  $f_F$  and can take advantage of the frozen image and text representations. Therefore, we advocate for the strategy in  $\mathcal{L}_{\text{UKR}}$  but encourage the adjustments to be minimal.

## 2.2 Optimization

The above loss functions correspond to different key properties for multimodal unlearning. We integrate them through the following aggregate loss function, which is optimized through stochastic gradient descent:

$$\mathcal{L} = \alpha \mathcal{L}_{\text{MD}} + \beta \mathcal{L}_{\text{MKR}} + \gamma \mathcal{L}_{\text{UKR}}, \quad (7)$$

where  $\mathcal{L}_{\text{MD}}$  realizes modality decoupling by ensuring that relations between data pairs marked for deletion are unrelated by the model,  $\mathcal{L}_{\text{MKR}}$  realizes multimodal knowledge retention by preserving the model’s overall multimodal knowledge, and  $\mathcal{L}_{\text{UKR}}$  realizes unimodal knowledge retention by preserving the model’s knowledge of individual modalities. This aggregated loss function effectively unlearns specific data points while maintaining the general functionality and knowledge of the original model. This balanced approach is crucial for the practical application of machine unlearning, particularly in settings where both data unlearning and model performance are of importance.

## 3 Experimental Setup

**Tasks and Datasets** Mul tiDelete is flexible and broadly applicable to a wide range of tasks, including generative ones. In this paper, we focus our evaluation on several image-text and graph-text tasks across several datasets.

- Image-Text Retrieval (TR) and (IR) are the tasks of retrieving the top-k relevant texts for a given image query (TR), and, vice versa, retrieving the top-k relevant images for a given text query (IR). We use Flickr30K dataset [1] for IR and TR.

**Table 1:** Statistics of evaluated datasets.

Dataset	Flickr30k	SNLI-VE	NLVR <sup>2</sup>	PGR
# images or graphs ( $I$ )	29.0K	29.8K	51.6K	4.0K
# texts ( $T$ )	144.5K	462.7K	22.8K	4.0K
# $I - T$ pairs	145.0K	529.5K	86.4K	4.0K

- Visual Entailment (VE) is an image-text entailment task, where the objective is to determine whether a given text hypothesis  $T_i$  entails, contradicts, or is neutral with respect to a given image premise  $I_i$ . We use SNLI-VE dataset [73] for VE.
- Natural Language for Visual Reasoning (NLVR) is the binary classification task of predicting whether a given text  $T_i$  accurately describes a given pair of images  $(I_{i,1}, I_{i,2})$ . We use NLVR<sup>2</sup> dataset [60] for NLVR.
- Graph-Text Classification is the task of classifying whether a text indicates a specific (e.g. causal) relationship between two given entities in a subgraph. We use PGR dataset [58], in which the target entities are phenotypes and genes, and the task is to determine if their relationship, as described by the accompanying text, is causal or non-causal.

Baselines We compare MultiDelete to the following models:

- Retrain is retraining a new model  $f$  of same architecture from scratch with the remaining data  $D_f$ .
- FineTune [24] is an optimization-based and modality-agnostic approach that unlearns data through continued fine-tuning. Specifically, it fine-tunes  $f$  on  $D_f$  with a larger learning rate, similar to catastrophic forgetting.
- NegGrad [24] is an optimization-based and modality-agnostic approach that unlearns data using negative gradient. Specifically, it optimizes the original loss function of training  $f$  on  $D_f$  but reverses the direction of gradients to unlearn these samples.
- DtD [49], Descent to Delete is a weight scrubbing-based and modality-agnostic approach to unlearning. It assumes that the weights of  $f'$  are close to the weights of  $f$ , trains  $f$  for a few more steps while adding Gaussian noise to scrub the weights.
- L-codec [48] is a weight scrubbing-based and uni-modal (vision only or text only) approach that approximates the Hessian matrix and performs a Newton update step to scrub the parameters while adding noise to them.
- Erm-K tp [40] is a retraining-based and uni-modal (vision only) approach that unlearns data by retraining the model with extra parameters inserted after visual feature maps to entangle correlations between classes. This method has been developed for machine unlearning in image classification.
- UL [29] is an optimization-based and unimodal (text only) approach that unlearns data by maximizing the log likelihood of samples in  $D_f$ . This method has been developed for machine unlearning in language models.

For each method, we also consider a variant where only the parameters of the fusion module  $f_F$  are updated during unlearning, while the rest of the model remains frozen; this is inspired by recent works on Parameter-Efficient Tuning [5, 11, 59]. We denote this setting by adding ‘-F’ to method names.

Evaluation Following previous works [11, 21, 30, 67], we employ several standard metrics to evaluate the unlearning efficacy of different methods:

- Test Set Performance ( $D_{\text{Test}} \uparrow$ ), which evaluates the unlearned model on the original test set  $D_{\text{Test}}$ . We follow previous work [35, 36, 66] to use mean recall (recall@1, recall@3, recall@10) for retrieval tasks and accuracy for other tasks. Higher values indicate that the model maintains better performance on the test set post unlearning.
- Deletion Set Performance ( $D_f \downarrow$ ), which evaluates the unlearned model on the forget set  $D_f$ . The metrics are the same as those on  $D_{\text{Test}}$ , with the difference that lower values indicate better performance.
- Membership Inference Vulnerability (MI  $\uparrow$ ), which measures vulnerability against membership inference (MI) attacks. Following previous work [11, 24], we evaluate the unlearned model  $f'$  in a blackbox MI setting, where the adversarial agent only has access to the output distribution of  $f'$ . An SVM classifier is trained using validation data as negative samples and a similarly sized subset of training data as positive samples [24]. We probe the existence probability of the deleted data  $D_f$  with the MI attacker before and after unlearning, and report the ratio of prior-to-post existence probabilities. A lower MI ratio indicates higher robustness to MI attacks and better protection to the data marked for deletion.

Settings We assume that only the associations between  $(I_i, T_i)$  are deleted, while individual  $I_i$  and  $T_i$  are not removed. For each dataset, we first train a corresponding model with the full training set ( $D_{\text{Train}}$ ). Then we randomly sample 5K data points from  $D_{\text{Train}}$  to create our deletion set  $D_f$ . We evaluate the unlearning methods across a range of deletion volumes, from 1K–5K samples with step size of 1K; this  $D_f/D_{\text{Train}}$  ratio matches the ratios used in previous studies [11, 12]. Tab. 1 shows statistics of the datasets and deleted data. We set  $\alpha, \beta, \gamma = 1$  in Eq. (7). The original models  $f$  are trained until convergence before being used for deletion experiments. For deletion, we select the best checkpoint using validation set of each dataset. Supplementary materials provides additional details.

Multimodal Architectures For image-text tasks, we use two popular pre-trained vision-language transformers, ALBEF [36] and BLIP [35], and follow their training and evaluation settings. For PGR, we employ the GCN [33] and BERT [17] models used in [66] to obtain unimodal subgraph and text representations respectively. These representations are then fused using a feed-forward network to obtain multimodal representations [66]. The unimodal and multimodal representations are concatenated and fed into a classifier for prediction.



**Table 2:** Experimental results on image-text and graph-text datasets on ALBEF [36]. Performance shows average of recall@1, recall@3, recall@10 on Flickr30K-TR and Flickr30K-IR, and accuracy on other datasets. ERM-KTP inserts trainable parameters to the vision encoder, causing the ERM-KTP-F variant (‘-F’ suffix in model titles denotes variants where only fusion module parameters are updated during unlearning) inapplicable. The best results are in **bold** and the second best results are underlined. The **RETRAIN** performance is provided *for reference purpose only*. See supplementary materials for additional results.

Method	Image-Text								Graph-Text		Avg.	
	Flickr30K				SNLI-VE		NLVR <sup>2</sup>		PGR			
	IR		TR		$D_{\text{Test}}$	$D_f$	$D_{\text{Test}}$	$D_f$	$D_{\text{Test}}$	$D_f$	$D_{\text{Test}}$	$D_f$
Retrain	97.8	50.4	93.5	50.4	79.4	50.2	80.3	50.3	67.5	50.2	83.4	50.3
FineTune	96.7	50.4	94.1	50.4	79.1	50.5	80.3	49.8	67.4	49.9	83.5	50.2
FineTune-F	97.1	49.9	94.6	49.9	79.5	49.9	81.2	50.0	67.5	50.1	<u>83.8</u>	49.9
NegGrad	92.4	50.5	91.7	50.5	77.8	48.6	77.3	50.6	63.4	49.6	80.5	50.0
NegGrad-F	93.3	50.2	90.6	50.2	<u>79.6</u>	50.6	80.8	50.0	63.5	49.9	81.5	50.2
DtD	10.3	51.4	8.9	51.4	45.2	50.1	50.8	49.8	50.0	50.2	33.0	50.5
DtD-F	22.5	50.9	20.7	50.9	48.6	49.8	50.9	49.8	53.6	50.2	39.2	50.2
L-codec	83.5	50.0	78.5	50.0	56.7	49.9	55.3	52.7	57.8	48.8	66.3	50.3
L-codec-F	87.4	49.4	50.6	48.2	57.4	48.4	56.8	53.1	59.1	46.9	62.2	49.2
Erm-K tp	57.4	48.7	56.2	49.0	53.2	48.9	52.9	50.8	N/A		54.9	49.3
Erm-K tp-F	N/A											
UL	95.1	50.4	90.3	50.4	75.7	49.8	76.3	50.4	64.8	49.7	80.4	50.2
UL-F	94.4	50.2	94.1	50.2	79.1	49.7	76.8	50.4	66.1	48.8	82.1	49.8
MultiDelete	97.1	33.2	<u>94.3</u>	33.2	79.8	35.3	80.8	23.5	68.5	18.6	84.2	28.7
MultiDelete-F	<u>96.8</u>	34.4	94.1	<u>34.5</u>	79.5	<u>36.3</u>	<u>80.4</u>	<u>26.4</u>	67.7	19.5	83.7	<u>30.2</u>

## 4 Main Results

The results in Tab. 2 show that on average across all *image-text* tasks, MultiDelete achieves 88.0 on  $D_{\text{Test}}$ , outperforming all baselines by 17.4 absolute points. Furthermore, it achieves 31.1 on  $D_f$ , outperforming all baselines by 17.6 absolute points. In addition, MultiDelete effectively reduces the likelihood of deleted data ( $D_f$ ) being identified, resulting in an average MI ratio of 1.3 across all tasks. These results indicate that MultiDelete can accomplish effective and targeted unlearning, while maintaining strong capability and utility on downstream tasks. In addition, on PGR dataset, MultiDelete outperforms baselines on  $D_{\text{Test}}$ , outperforming FineTune by +1.1 points, NegGrad by +5.1 points, DtD by +38.5 points, L-codec by +10.7 points, and UL by +3.7 points. In addition, MultiDelete achieves 18.6 on deleted data ( $D_f$ ), significantly outperforming the top-performing baseline model, L-codec (48.8). The performance of MultiDelete on the original test set  $D_{\text{Test}}$  is better than that of Retrain by +1.0, see Tab. 2.

Comparison to Modality-agnostic Approaches Results in Tab. 2 show that none of the existing modality-agnostic approaches is sufficient to unlearn multimodal samples from trained models. Specifically, on  $D_{\text{Test}}$ , MultiDelete

outperforms FineTune, NegGrad, DtD, L-codec by +0.6, +3.3, +49.3 and +19.6 average points respectively. The corresponding improvements on  $D_f$  are +19.0, +18.8, +19.4 and +19.4 points respectively. The lower performance of these approaches show that they can't remove learned multimodal dependencies.

**Comparison to Unimodal Approaches** Results in Tab. 2 show that unimodal unlearning approaches do not effectively translate to multimodal contexts. Specifically, MultiDelete outperforms Erm-Ktp and UL by substantial margins of +33.2 and +3.8 absolute points on  $D_{\text{Test}}$ , +18.0 and 19.0 absolute points on  $D_f$ . Updating the knowledge on one of the modalities results in drop on both test set performance and model's ability in forgetting  $D_f$ . Therefore, merely unlearning a single modality is inadequate for comprehensive unlearning in multimodal settings, where removal of inter-modality association is anticipated.

**Limitations of Scrubbing Methods and Retrain** Our results in Tab. 2 show that scrubbing methods, despite with theoretical guarantee, fall short in multimodal unlearning in practice; the scrubbing methods DtD and L-codec achieve an average performance of 37.6 and 59.5, respectively, which are considerably lower than that of FineTune (68.5) and NegGrad (67.5). They also result in extremely low performance of 26.15 and 68.5 on  $D_{\text{Test}}$  respectively. In case of multimodal settings, we argue that scrubbing or noise addition disrupts the original learned dependencies, particularly when model parameters are shared, e.g. by nodes in graphs [11] or by different fused modalities. In unimodal settings, where scrubbing methods are tested, since the encoder encompasses most of the model parameters, scrubbing methods do not show strong influence on the performance of downstream tasks. MultiDelete even outperforms Retrain by  $xx$  and  $xx$  absolute point on  $D_{\text{Test}}$  and  $D_f$ , respectively. As several previous works have observed [11, 21, 23, 34], Retrain does not necessarily serve as the gold standard for unlearning. These results indicate that matching model parameters does not necessarily mean successful unlearning due to potential distribution discrepancy in model parameters, as noted by many previous works [11, 21, 23, 34, 64].

**Membership Inference Attack** MultiDelete achieves a reduced probability of detecting deleted data ( $D_f$ ) compared to before unlearning (see the prior-to-post MI ratio in Tab. 2). This indicates that MultiDelete can better protect the deleted data and is less susceptible to MI attacks. Specifically, MultiDelete outperforms non-scrubbing baselines (FineTune, NegGrad, Erm-Ktp, UL) by 0.19 absolute points in MI ratio. We note that, although scrubbing methods like DtD and L-codec show a significant decrease in existence probability (higher MI ratios) than non-scrubbing methods, the drop applies to all data including both  $D_r$  and  $D_f$ . This shows that the unlearning of scrubbing methods is not targeted at a specific subset of data, but the entire data, which signals a failed deletion.

**Table 3:** Prior-to-post MI ratio on image-text datasets with the best-performing baseline. The RETRAIN performance is provided *for reference purpose only*. See supplementary materials for additional results.

Method	Image-Text				Graph-Text	Avg.
	Flickr-IR	Flickr-TR	SNLI-VE	NLVR <sup>2</sup>	PGR	
Retrain	1.10	1.10	1.05	1.07	1.09	1.08
L-codec	1.21	1.21	1.23	1.23	1.07	1.19
L-codec-F	1.22	1.22	1.26	1.26	1.09	1.21
MultiDelete	1.27	1.27	1.30	1.25	1.24	1.27
MultiDelete-F	<u>1.25</u>	<u>1.25</u>	<u>1.26</u>	<u>1.21</u>	<u>1.20</u>	<u>1.24</u>

All Key Properties Contribute to Unlearning Through an ablation study, we assess the individual contributions of the key properties proposed in MultiDelete: modality decoupling (MD), multimodal knowledge retention (UKR), and unimodal knowledge retention (MKR). Tab. 4 shows that excluding MD results in a significant decline in model’s ability in distinguishing between  $D_f$  and  $D_r$ , with performance dropping from 76.5 to 50.3 (−26.2). Both UKR and MKR serve as objectives for maintaining the original knowledge acquired by the model, targeting at  $D_f$  and  $D_r$  respectively. The exclusion of UKR and MKR lead to performance drops of 0.5 and 6.6 on downstream tasks respectively. The more substantial impact observed by removing MKR can be attributed to two factors: (1)  $D_r$  usually has a much larger size than  $D_f$ , leading to a much larger influence for MKR; and (2) downstream tasks tend to rely more heavily on multimodal knowledge than unimodal knowledge, making MKR crucial for maintaining model performance.

Utility of Unimodal Knowledge We highlight that the multimodal unlearning in MultiDelete can retain the utility of unimodal embeddings. We train unimodal image and text classifiers  $g_I$  on  $f(I)$ , and test the classifiers with the unimodal embeddings after unlearning, namely  $f'(I)$ . This experiment can show how much unlearning preserves the utility of unimodal embeddings  $f'(I)$  after unlearning, where an optimal unlearning method should yield same performance for  $g_I(f(I))$  and  $g_I(f'(I))$ . Tab. 5 MultiDelete outperforms the best baseline (FineTune) by +0.4 in accuracy, while removing UKR in optimization leading to dramatic drop by −5.7.

Updating All Parameters vs. Fusion Module Only We experiment whether updating all model parameters or just those of the modality fusion module is more effective for unlearning. We denote these two approaches as Method (updating all parameters) and Method-F (updating only the fusion module parameters). For MultiDelete, focusing solely on updating  $f_F$  (the fusion module) is somewhat akin to bypassing the optimization for  $\mathcal{L}_{UKR}$ , though not exactly the same. We find that this version of MultiDelete exhibits less fluctuation in performance on  $D_{Test}$  during training but tends to converge more slowly on  $D_f|D_r$ .

**Table 4:** Ablation study of unlearning objectives of MULTIDELETE. All three objectives contributes to both downstream performance on  $D_{\text{Test}}$  and forgetting  $D_f$ .

	NLVR <sup>2</sup>		PGR	
	$D_{\text{Test}}$	$D_f$	$D_{\text{Test}}$	$D_f$
Retrain	80.3	50.3	67.5	50.2
Full model	80.8	23.5	67.8	18.6
- MD	80.3	50.3	67.5	49.3
- UKR	79.2	25.8	66.3	22.6
- MKR	77.1	25.6	64.8	23.7


**Table 5:** Accuracy of image classification using unimodal embeddings after unlearning.

	FINETUNE	NEGGRAD	DtD	L-CODEC	UL	MULTIDELETE	w/o UKR
Acc.	83.2	81.7	43.8	55.2	82.7	83.6	77.9

compared to the full version of MultiDelete. For scrubbing-based methods (DtD, L-codec), updating all the parameters results in a complete loss previously acquired knowledge, resulting in random performance across all tasks. Conversely, targeting only the fusion modules for scrubbing helps retain performance on downstream tasks. This suggests that (a) robust unimodal knowledge plays a critical role in multimodal tasks, and (b) the fusion module is more resilient to noise or minor perturbations than the unimodal encoders. However, neither approach significantly aids the model in distinguishing  $D_f$  from  $D_r$  or in protecting  $D_f$  against MI attacks. For modality-agnostic approaches, we observe negligible differences between the two strategies, with a marginal performance gap of less than 0.6 absolute points. This indicates that for these methods, the strategy chosen for parameter updating has minimal impact on overall performance.

*Efficiency For training time*, Fig. 2 presents a comparative analysis of the training times for MultiDelete and Retrain across datasets with an increasing size of  $|D_f|$ . The results indicate that MultiDelete is efficient to run, and exhibits a linear growth in running time as  $|D_f|$  increases. This trend illustrates MultiDelete’s scalability and effectiveness in managing larger volumes of data marked for deletion. *For trainable parameters*, several works have optimized unlearning in a parameter-efficient manner [5, 11, 21], similar to parameter-efficient fine-tuning [19, 59]. We argue that MultiDelete-F only optimizes a small portion of the parameters, while delivering comparable performance as the full model with trivial gap.

*Why MultiDelete works* We attribute the superior performance of MultiDelete to the three proposed properties. MD aims at relaxing the dependencies between  $(I_i, T_i)$  pairs, while MKR and UKR preserve the acquired knowledge. Compared to existing unimodal approaches, MD can remove the relationships



**Fig. 2:** Training time of unlearning methods.

between data modalities. Compared to existing modality-agnostic approaches, MKR and UKR maintains the capability of model on multimodal tasks. Collectively, all three properties contribute to the success of unlearning, where strong downstream performance and successful targeted forgetting are desired in the same time.

## 5 Related work

Machine Unlearning Existing machine unlearning research can be categorized into four classes. (a): *Efficient retraining methods* retrain models on partitioned data and learn to combine the predictions from each sharded model to maintain performance [2, 7, 20, 44, 71, 72]. Training on partitioned multimodal data may inversely affect the dependencies between modalities, lead to overfitting, be inefficient and not scalable. (b) *Weight scrubbing methods* adopt a one-shot weight update followed by added noise to model weights, whose probability distribution is indistinguishable from that of a model retrained from scratch with theoretical guarantee [3, 24, 25, 28, 42, 49, 61, 70]. However, the assumption of strongly convex loss cannot be guaranteed in multimodal setting. Added noise may affect the modality dependencies of all data points and poor empirical performance [64]. (c) *Teacher-student unlearning*: centers on making the student (unlearned model) to follow the teacher on deleted data. The teacher can be a separate model [68], the opposite direction of the original model [34], an untrained model [15], random node pairs [11], a random label [21], an error matrix [62].

Unlearning in Vision Certified data removal [25] views image classifiers as frozen feature extractors and linear classifiers and derives a Newton update step to remove information. Boundary Unlearning [6] tackles the problem of unlearning an entire class by shifting the decision boundary, either shrinking or expanding the original boundary. ERM-KTP [40] inserts entanglement-reduced mask (ERM) layers and retrains the model, incorporated during training and doesn't handle existing trained models. MUter [43] proposes a close-form solution for unlearning on adversarially trained vision models based on influence measure.

Unlearning in Language KGA [67] aligns the knowledge gap between the deleted data and remaining data of the unlearned model with a separate model

trained with extra data. However, obtaining high quality data from the same distribution may not be easy, especially in multimodal setting. Other methods include negative gradient [29], encouraging model to generate dissimilar text to ground truth [31], making samples unlearnable [37]. EUL [5] maximizes the difference between the unlearning model and original model on the deleted samples.

**Unlearning in Graphs** GraphEraser [8] re-trains multiple models on partitioned graphs, which removes structural information and hurt performance. CGU [13], CEU [70] and Projector [16] assume linear underlying models, which limits their practical performance and applicability. GNNDelete [11] performs local update to fulfill Deleted Edge Consistency and Neighborhood Influence. D2DGN [56] formulates unlearning as distillation from the original model.

**Other Unlearning Work** Other approaches include sparsity [21, 30, 46] and operations on gradient [27, 65], training with ability to unlearn in mind [40, 74], in continual learning setting [41], for recommender systems [4, 38, 39], in Bayesian models [50], without accessing training data [14], for shallow models [22, 54], handling a sequence of deletion requests [26], for regression models [63], verification [57], vulnerability to attack [75], under federated setting [51], trade-off with reverting decisions [52], and benchmarking [10, 47]. Applications of unlearning include removing bias [9, 55], text generation [45], alleviating backdoor attack [69], and conducting data poison attack [18].

## 6 Conclusion

This work formulates the problem of multimodal unlearning and introduces MultiDelete—the first multimodal unlearning approach that is task and architecture agnostic, and is efficient to use. It implements three key properties for effective multimodal unlearning: (a): modality decoupling, which aims to decouple the association between individual unimodal data points marked for deletion, rendering them as unrelated data points, (b): multimodal knowledge retention, which retains the multimodal representation capability of the model post-unlearning, and (c): unimodal knowledge retention, which retains the unimodal representation capability of the model post-unlearning. Through extensive experiments across image-text and graph-text tasks and several datasets and architectures, we show that MultiDelete reliably unlearn multimodal relationships, and outperforms existing modality-agnostic and unimodal unlearning methods in maintaining downstream performance, and distinguishing between deleted data and remaining data. In addition, we show that the model can maintain the multimodal and unimodal knowledge post unlearning, can provide better protection to unlearned data, and is robust against adversarial attacks.

## References

1. Bojchevski, A., Günnemann, S.: Deep gaussian embedding of graphs: Unsupervised inductive learning via ranking. In: International Conference on Learning Representations (2018)
2. Bourtole, L., Chandrasekaran, V., Choquette-Choo, C.A., Jia, H., Travers, A., Zhang, B., Lie, D., Papernot, N.: Machine unlearning. In: IEEE Symposium on Security and Privacy (SP) (2021)
3. Brophy, J., Lowd, D.: Machine unlearning for random forests. In: Meila, M., Zhang, T. (eds.) Proceedings of the 38th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 139, pp. 1092–1104. PMLR (18–24 Jul 2021), <https://proceedings.mlr.press/v139/brophy21a.html>
4. Chen, C., Sun, F., Zhang, M., Ding, B.: Recommendation unlearning. In: Proceedings of the ACM Web Conference 2022 (2022)
5. Chen, J., Yang, D.: Unlearn what you want to forget: Efficient unlearning for LLMs. In: Bouamor, H., Pino, J., Bali, K. (eds.) Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing. pp. 12041–12052. Association for Computational Linguistics, Singapore (Dec 2023). <https://doi.org/10.18653/v1/2023.emnlp-main.738>, <https://aclanthology.org/2023.emnlp-main.738>
6. Chen, M., Gao, W., Liu, G., Peng, K., Wang, C.: Boundary unlearning: Rapid forgetting of deep networks via shifting the decision boundary. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 7766–7775 (2023), <https://api.semanticscholar.org/CorpusID:257636742>
7. Chen, M., Zhang, Z., Wang, T., Backes, M., Humbert, M., Zhang, Y.: Graph unlearning. Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (2021), <https://api.semanticscholar.org/CorpusID:232404451>
8. Chen, M., Zhang, Z., Wang, T., Backes, M., Humbert, M., Zhang, Y.: Graph unlearning. In: Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (2022)
9. Chen, R., Yang, J., Xiong, H., Bai, J., Hu, T., Hao, J., FENG, Y., Zhou, J.T., Wu, J., Liu, Z.: Fast model debias with machine unlearning. In: Thirty-seventh Conference on Neural Information Processing Systems (2023), <https://openreview.net/forum?id=BL9Pc7xsdX>
10. Cheng, J., Amiri, H.: Mu-bench: A multitask multimodal benchmark for machine unlearning. arXiv preprint arXiv:2406.14796 (2024)
11. Cheng, J., Dasoulas, G., He, H., Agarwal, C., Zitnik, M.: GNNDelete: A general strategy for unlearning in graph neural networks. In: The Eleventh International Conference on Learning Representations (2023), <https://openreview.net/forum?id=X9yCkmT5Qr1>
12. Chien, E., Pan, C., Milenkovic, O.: Certified graph unlearning. In: NeurIPS 2022 Workshop: New Frontiers in Graph Learning (2022)
13. Chien, E., Pan, C., Milenkovic, O.: Efficient model updates for approximate unlearning of graph-structured data. In: International Conference on Learning Representations (2023), <https://api.semanticscholar.org/CorpusID:259298766>
14. Chundawat, V.S., Tarun, A.K., Mandal, M., Kankanhalli, M.: Zero-shot machine unlearning. arXiv (2022)
15. Chundawat, V.S., Tarun, A.K., Mandal, M., Kankanhalli, M.S.: Can bad teaching induce forgetting? unlearning in deep networks using an incompetent teacher.

- ArXiv [abs/2205.08096](https://arxiv.org/abs/2205.08096) (2022), <https://api.semanticscholar.org/CorpusID:248834527>
16. Cong, W., Mahdavi, M.: Efficiently forgetting what you have learned in graph representation learning via projection. In: Ruiz, F., Dy, J., van de Meent, J.W. (eds.) *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research*, vol. 206, pp. 6674–6703. PMLR (25–27 Apr 2023), <https://proceedings.mlr.press/v206/cong23a.html>
  17. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: Burstein, J., Doran, C., Solorio, T. (eds.) *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. pp. 4171–4186. Association for Computational Linguistics, Minneapolis, Minnesota (Jun 2019). <https://doi.org/10.18653/v1/N19-1423>, <https://aclanthology.org/N19-1423>
  18. Di, J.Z., Douglas, J., Acharya, J., Kamath, G., Sekhari, A.: Hidden poison: Machine unlearning enables camouflaged poisoning attacks. In: *Thirty-seventh Conference on Neural Information Processing Systems (2023)*, <https://openreview.net/forum?id=Isy7gl1Hqc>
  19. Ding, N., Qin, Y., Yang, G., Wei, F., Yang, Z., Su, Y., Hu, S., Chen, Y., Chan, C.M., Chen, W., Yi, J., Zhao, W., Wang, X., Liu, Z., Zheng, H.T., Chen, J., Liu, Y., Tang, J., Li, J., Sun, M.: Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nature Machine Intelligence* **5**(3), 220–235 (Mar 2023). <https://doi.org/10.1038/s42256-023-00626-4>, <https://doi.org/10.1038/s42256-023-00626-4>
  20. Dukler, Y., Bowman, B., Achille, A., Golatkar, A., Swaminathan, A., Soatto, S.: Safe: Machine unlearning with shard graphs. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 17108–17118 (October 2023)
  21. Fan, C., Liu, J., Zhang, Y., Wong, E., Wei, D., Liu, S.: Salun: Empowering machine unlearning via gradient-based weight saliency in both image classification and generation. In: *The Twelfth International Conference on Learning Representations (2024)*, <https://openreview.net/forum?id=gn0mIhQGNN>
  22. Ginart, A., Guan, M., Valiant, G., Zou, J.Y.: Making ai forget you: Data deletion in machine learning. In: *Advances in Neural Information Processing Systems (2019)*
  23. Goel, S., Prabhu, A., Sanyal, A., Lim, S.N., Torr, P., Kumaraguru, P.: Towards adversarial evaluations for inexact machine unlearning. *arXiv* (2022)
  24. Golatkar, A., Achille, A., Soatto, S.: Eternal sunshine of the spotless net: Selective forgetting in deep networks. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (2020)*
  25. Guo, C., Goldstein, T., Hannun, A., Van Der Maaten, L.: Certified data removal from machine learning models. In: *Proceedings of the 37th International Conference on Machine Learning. Proceedings of Machine Learning Research, PMLR* (2020), <https://proceedings.mlr.press/v119/guo20c.html>
  26. Gupta, V., Jung, C., Neel, S., Roth, A., Sharifi-Malvajerdi, S., Waites, C.: Adaptive machine unlearning. In: *Advances in Neural Information Processing Systems (2021)*
  27. Hoang, T., Rana, S., Gupta, S., Venkatesh, S.: Learn to unlearn for deep neural networks: Minimizing unlearning interference with gradient projection. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. pp. 4819–4828 (January 2024)



28. Izzo, Z., Smart, M.A., Chaudhuri, K., Zou, J.: Approximate data deletion from machine learning models. In: Proceedings of The International Conference on Artificial Intelligence and Statistics (2021)
29. Jang, J., Yoon, D., Yang, S., Cha, S., Lee, M., Logeswaran, L., Seo, M.: Knowledge unlearning for mitigating privacy risks in language models. In: Rogers, A., Boyd-Graber, J., Okazaki, N. (eds.) Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 14389–14408. Association for Computational Linguistics, Toronto, Canada (Jul 2023). <https://doi.org/10.18653/v1/2023.acl-long.805>, <https://aclanthology.org/2023.acl-long.805>
30. Jia, J., Liu, J., Ram, P., Yao, Y., Liu, G., Liu, Y., Sharma, P., Liu, S.: Model sparsity can simplify machine unlearning. In: Thirty-seventh Conference on Neural Information Processing Systems (2023), <https://openreview.net/forum?id=0jZH883i34>
31. Kassem, A., Mahmoud, O., Saad, S.: Preserving privacy through dememorization: An unlearning technique for mitigating memorization risks in language models. In: Bouamor, H., Pino, J., Bali, K. (eds.) Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing. pp. 4360–4379. Association for Computational Linguistics, Singapore (Dec 2023). <https://doi.org/10.18653/v1/2023.emnlp-main.265>, <https://aclanthology.org/2023.emnlp-main.265>
32. Kim, W., Son, B., Kim, I.: Vilt: Vision-and-language transformer without convolution or region supervision. In: Meila, M., Zhang, T. (eds.) Proceedings of the 38th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 139, pp. 5583–5594. PMLR (18–24 Jul 2021), <https://proceedings.mlr.press/v139/kim21k.html>
33. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: International Conference on Learning Representations (2017)
34. Kurmanji, M., Triantafillou, P., Hayes, J., Triantafillou, E.: Towards unbounded machine unlearning. In: Oh, A., Neumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S. (eds.) Advances in Neural Information Processing Systems. vol. 36, pp. 1957–1987. Curran Associates, Inc. (2023), [https://proceedings.neurips.cc/paper\\_files/paper/2023/file/062d711fb777322e2152435459e6e9d9-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/062d711fb777322e2152435459e6e9d9-Paper-Conference.pdf)
35. Li, J., Li, D., Xiong, C., Hoi, S.: BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., Sabato, S. (eds.) Proceedings of the 39th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 162, pp. 12888–12900. PMLR (17–23 Jul 2022), <https://proceedings.mlr.press/v162/li22n.html>
36. Li, J., Selvaraju, R., Gotmare, A., Joty, S., Xiong, C., Hoi, S.C.H.: Align before fuse: Vision and language representation learning with momentum distillation. In: Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (eds.) Advances in Neural Information Processing Systems. vol. 34, pp. 9694–9705. Curran Associates, Inc. (2021), [https://proceedings.neurips.cc/paper\\_files/paper/2021/file/505259756244493872b7709a8a01b536-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/505259756244493872b7709a8a01b536-Paper.pdf)
37. Li, X., Liu, M.: Make text unlearnable: Exploiting effective patterns to protect personal data. In: Ovalle, A., Chang, K.W., Mehrabi, N., Pruksachatkun, Y., Galystan, A., Dhamala, J., Verma, A., Cao, T., Kumar, A., Gupta, R. (eds.) Proceedings of the 3rd Workshop on Trustworthy Natural Language Processing (TrustNLP 2023). pp. 249–259. Association for Computational Linguistics,

- Toronto, Canada (Jul 2023). <https://doi.org/10.18653/v1/2023.trustnlp-1.22>, <https://aclanthology.org/2023.trustnlp-1.22>
38. Li, Y., Chen, C., Zhang, Y., Liu, W., Lyu, L., Zheng, X., Meng, D., Wang, J.: UltraRE: Enhancing receraser for recommendation unlearning via error decomposition. In: Thirty-seventh Conference on Neural Information Processing Systems (2023), <https://openreview.net/forum?id=93NLxUojvc>
  39. Li, Y., Zheng, X., Chen, C., Liu, J.: Making recommender systems forget: Learning and unlearning for erasable recommendation (2022)
  40. Lin, S., Zhang, X., Chen, C., Chen, X., Susilo, W.: Erm-ktp: Knowledge-level machine unlearning via knowledge transfer. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 20147–20155 (2023), <https://api.semanticscholar.org/CorpusID:260082364>
  41. Liu, B., Liu, Q., Stone, P.: Continual learning and private unlearning. In: Chandar, S., Pascanu, R., Precup, D. (eds.) Proceedings of The 1st Conference on Lifelong Learning Agents. Proceedings of Machine Learning Research, vol. 199, pp. 243–254. PMLR (22–24 Aug 2022), <https://proceedings.mlr.press/v199/liu22a.html>
  42. Liu, J., Lou, J., Qin, Z., Ren, K.: Certified minimax unlearning with generalization rates and deletion capacity. In: Thirty-seventh Conference on Neural Information Processing Systems (2023), <https://openreview.net/forum?id=6H8Md75kAw>
  43. Liu, J., Xue, M., Lou, J., Zhang, X., Xiong, L., Qin, Z.: Muter: Machine unlearning on adversarially trained models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 4892–4902 (October 2023)
  44. Liu, Y., Xu, L., Yuan, X., Wang, C., Li, B.: The right to be forgotten in federated learning: An efficient realization with rapid retraining. In: IEEE Conference on Computer Communications (2022)
  45. Lu, X., Welleck, S., Hessel, J., Jiang, L., Qin, L., West, P., Ammanabrolu, P., Choi, Y.: QUARK: Controllable text generation with reinforced unlearning. In: Oh, A.H., Agarwal, A., Belgrave, D., Cho, K. (eds.) Advances in Neural Information Processing Systems (2022), <https://openreview.net/forum?id=5HaIds3ux50>
  46. Ma, Z., Liu, Y., Liu, X., Liu, J., Ma, J., Ren, K.: Learn to forget: Machine unlearning via neuron masking. IEEE Transactions on Dependable and Secure Computing **20**(4), 3194–3207 (2023). <https://doi.org/10.1109/TDSC.2022.3194884>
  47. Maini, P., Feng, Z., Schwarzschild, A., Lipton, Z.C., Kolter, J.Z.: Tofu: A task of fictitious unlearning for llms. arXiv preprint arXiv:2401.06121 (2024)
  48. Mehta, R.R., Pal, S., Singh, V., Ravi, S.: Deep unlearning via randomized conditionally independent Hessians. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 10412–10421 (2022), <https://api.semanticscholar.org/CorpusID:248227997>
  49. Neel, S., Roth, A., Sharifi-Malvajerdi, S.: Descent-to-delete: Gradient-based methods for machine unlearning. In: Proceedings of the International Conference on Algorithmic Learning Theory (2021)
  50. Nguyen, Q.P., Low, B.K.H., Jailliet, P.: Variational bayesian unlearning. In: Advances in Neural Information Processing Systems (2020)
  51. Pan, C., Sima, J., Prakash, S., Rana, V., Milenkovic, O.: Machine unlearning of federated clusters. In: The Eleventh International Conference on Learning Representations (2023), <https://openreview.net/forum?id=VzwfoFyYDga>
  52. Pawelczyk, M., Leemann, T., Biega, A., Kasneci, G.: On the trade-off between actionable explanations and the right to be forgotten. In: The Eleventh International Conference on Learning Representations (2023), <https://openreview.net/forum?id=HWt4BBZjVW>

53. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision (2021)
54. Schelter, S., Grafberger, S., Dunning, T.: Hedgecut: Maintaining randomised trees for low-latency machine unlearning. In: Proceedings of the International Conference on Management of Data (2021)
55. Setlur, A., Eysenbach, B., Smith, V., Levine, S.: Adversarial unlearning: Reducing confidence along adversarial directions. In: Oh, A.H., Agarwal, A., Belgrave, D., Cho, K. (eds.) Advances in Neural Information Processing Systems (2022), <https://openreview.net/forum?id=cJ006qBE8Uv>
56. Sinha, Y., Mandal, M., Kankanhalli, M.: Distill to delete: Unlearning in graph networks with knowledge distillation. arXiv preprint arXiv:2309.16173 (2023)
57. Sommer, D.M., S Ong, L., Wagh, S., Mittal, P.: Athena: Probabilistic verification of machine unlearning. In: Proceedings on Privacy Enhancing Technologies Symposium (PETS) (2022)
58. Sousa, D., Lamurias, A., Couto, F.M.: A silver standard corpus of human phenotype-gene relations. In: Burstein, J., Doran, C., Solorio, T. (eds.) Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers). pp. 1487–1492. Association for Computational Linguistics, Minneapolis, Minnesota (Jun 2019). <https://doi.org/10.18653/v1/N19-1152>, <https://aclanthology.org/N19-1152>
59. Su, Y., Chan, C.M., Cheng, J., Qin, Y., Lin, Y., Hu, S., Yang, Z., Ding, N., Sun, X., Xie, G., Liu, Z., Sun, M.: Exploring the impact of model scaling on parameter-efficient tuning. In: Bouamor, H., Pino, J., Bali, K. (eds.) Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing. pp. 15062–15078. Association for Computational Linguistics, Singapore (Dec 2023). <https://doi.org/10.18653/v1/2023.emnlp-main.931>, <https://aclanthology.org/2023.emnlp-main.931>
60. Suhr, A., Zhou, S., Zhang, A., Zhang, I., Bai, H., Artzi, Y.: A corpus for reasoning about natural language grounded in photographs. In: Korhonen, A., Traum, D., Màrquez, L. (eds.) Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. pp. 6418–6428. Association for Computational Linguistics, Florence, Italy (Jul 2019). <https://doi.org/10.18653/v1/P19-1644>, <https://aclanthology.org/P19-1644>
61. Suriyakumar, V.M., Wilson, A.C.: Algorithms that approximate data removal: New results and limitations. In: Oh, A.H., Agarwal, A., Belgrave, D., Cho, K. (eds.) Advances in Neural Information Processing Systems (2022), <https://openreview.net/forum?id=G4V0QPYxBsI>
62. Tarun, A.K., Chundawat, V.S., Mandal, M., Kankanhalli, M.: Fast yet effective machine unlearning. IEEE Transactions on Neural Networks and Learning Systems pp. 1–10 (2023). <https://doi.org/10.1109/TNNLS.2023.3266233>
63. Tarun, A.K., Chundawat, V.S., Mandal, M., Kankanhalli, M.: Deep regression unlearning. In: Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., Scarlett, J. (eds.) Proceedings of the 40th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 202, pp. 33921–33939. PMLR (23–29 Jul 2023), <https://proceedings.mlr.press/v202/tarun23a.html>
64. 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. USENIX Association, Boston, MA (Aug

- 2022), <https://www.usenix.org/conference/usenixsecurity22/presentation/thudi>
65. Ullah, E., Mai, T., Rao, A., Rossi, R.A., Arora, R.: Machine unlearning via algorithmic stability. In: Belkin, M., Kpotufe, S. (eds.) Proceedings of Thirty Fourth Conference on Learning Theory. Proceedings of Machine Learning Research, vol. 134, pp. 4126–4142. PMLR (15–19 Aug 2021), <https://proceedings.mlr.press/v134/ullah21a.html>
  66. Vakil, N., Amiri, H.: Generic and trend-aware curriculum learning for relation extraction. In: Carpuat, M., de Marneffe, M.C., Meza Ruiz, I.V. (eds.) Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. pp. 2202–2213. Association for Computational Linguistics, Seattle, United States (Jul 2022). <https://doi.org/10.18653/v1/2022.naacl-main.160>, <https://aclanthology.org/2022.naacl-main.160>
  67. Wang, L., Chen, T., Yuan, W., Zeng, X., Wong, K.F., Yin, H.: KGA: A general machine unlearning framework based on knowledge gap alignment. In: Rogers, A., Boyd-Graber, J., Okazaki, N. (eds.) Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 13264–13276. Association for Computational Linguistics, Toronto, Canada (Jul 2023). <https://doi.org/10.18653/v1/2023.acl-long.740>, <https://aclanthology.org/2023.acl-long.740>
  68. Wang, L., Chen, T., Yuan, W., Zeng, X., Wong, K.F., Yin, H.: Kga: A general machine unlearning framework based on knowledge gap alignment. In: Annual Meeting of the Association for Computational Linguistics (2023), <https://api.semanticscholar.org/CorpusID:258615571>
  69. Wei, S., Zhang, M., Zha, H., Wu, B.: Shared adversarial unlearning: Backdoor mitigation by unlearning shared adversarial examples. In: Thirty-seventh Conference on Neural Information Processing Systems (2023), <https://openreview.net/forum?id=zq0cW3R9rd>
  70. Wu, K., Shen, J., Ning, Y., Wang, T., Wang, W.H.: Certified edge unlearning for graph neural networks. In: Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. p. 2606–2617. KDD '23, Association for Computing Machinery, New York, NY, USA (2023). <https://doi.org/10.1145/3580305.3599271>, <https://doi.org/10.1145/3580305.3599271>
  71. Wu, Y., Dobriban, E., Davidson, S.: DeltaGrad: Rapid retraining of machine learning models. In: Proceedings of the International Conference on Machine Learning (2020)
  72. Wu, Y., Dobriban, E., Davidson, S.B.: Deltagrad: Rapid retraining of machine learning models. In: International Conference on Machine Learning (2020), <https://api.semanticscholar.org/CorpusID:220128049>
  73. Xie, N., Lai, F., Doran, D., Kadav, A.: Visual entailment: A novel task for fine-grained image understanding. arXiv preprint arXiv:1901.06706 (2019)
  74. Zhang, Z., Zhou, Y., Zhao, X., Che, T., Lyu, L.: Prompt certified machine unlearning with randomized gradient smoothing and quantization. In: Oh, A.H., Agarwal, A., Belgrave, D., Cho, K. (eds.) Advances in Neural Information Processing Systems (2022), <https://openreview.net/forum?id=ue4gP8ZKiWb>
  75. ZHAO, C., Qian, W., Ying, Z., Huai, M.: Static and sequential malicious attacks in the context of selective forgetting. In: Thirty-seventh Conference on Neural Information Processing Systems (2023), <https://openreview.net/forum?id=uEJfW30tUm>