This is the peer reviewd version of the followng article:

Multi-Class Unlearning for Image Classification via Weight Filtering / Poppi, Samuele; Sarto, Sara; Cornia, Marcella; Baraldi, Lorenzo; Cucchiara, Rita. - In: IEEE INTELLIGENT SYSTEMS. - ISSN 1541-1672. - (2024), pp. 1-8. [10.1109/MIS.2024.3412742]

Terms of use:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

02/10/2024 13:54

Multi-Class Unlearning for Image Classification via Weight Filtering

Samuele Poppi University of Modena and Reggio Emilia, Modena, 41125, Italy

Sara Sarto University of Modena and Reggio Emilia, Modena, 41125, Italy

Marcella Cornia University of Modena and Reggio Emilia, Reggio Emilia, 42121, Italy

Lorenzo Baraldi University of Modena and Reggio Emilia, Modena, 41125, Italy

Rita Cucchiara University of Modena and Reggio Emilia, Modena, 41125, Italy

Abstract—Machine Unlearning is an emerging paradigm for selectively

removing the impact of training datapoints from a network. Unlike existing methods that target a limited subset or a single class, our framework unlearns all classes in a single round. We achieve this by modulating the network's components using memory matrices, enabling the network to demonstrate selective unlearning behavior for any class after training. By discovering weights that are specific to each class, our approach also recovers a representation of the classes which is explainable by design. We test the proposed framework on smalland medium-scale image classification datasets, with both convolution- and Transformer-based backbones, showcasing the potential for explainable solutions through unlearning.

DEEP LEARNING MODELS and web datasets necessitate protecting user privacy and supporting the "right to be forgotten" [1]. Among other techniques which can alter the internal knowledge representation of a model, like knowledge distillation, model pruning, or incremental learning, Machine Unlearning [2] emerges as the sole technique which can delete private training data. Machine unlearning approaches [3]–[5], indeed, remove traces left by specific datapoints from trained models. Their goal is to "untrain" the model, eliminating the impact of unwanted datapoints and reaching weights similar to those of models trained without such data. For instance, entire classes can be removed from classification networks [6]–[8], which is useful in cases like face recognition. While existing machine unlearning approaches are limited to unlearning one or few classes jointly, our work advances by unlearning *all* classes orthogonally and in a *single fine-tuning round*. In particular, our approach can unlearn *all* classes in an orthogonal fashion, ensuring the user a customizable, final model which can behave as a model unlearned on a single class of choice, imposed by the user at runtime. This is obtained by exploiting the mapping between inner

1



Figure 1. The proposed single-round multi-class unlearning setting, which can unlearn *any* class in a *single* untraining round. WF-Net requires less computational resources and supports explainability by-design.

network components (e.g., convolutional filters in a Convolutional Neural Network (CNN) or attentive projection in a Vision Transformer (ViT)) and output classes [9], and builds a Weight-Filtering layer which can selectively turn on and off those inner components to accomplish the desired unlearning behavior on a class of choice (Figure 1). This is achieved by encapsulating existing network operators, avoiding alterations to the network structure. By training the proposed Weight-Filtering layers, we essentially uncover the connections between the inner network components and output classes. This not only enables effective unlearning but also provides an inherently interpretable representation, finally gaining explainability properties through unlearning. Experimentally, we validate the proposed approach, named WF-Net, on small-scale and medium-scale image classification datasets and demonstrate its applicability to a variety of image classification architectures, including both CNNs and ViTs and to the more challenging case of unlearning without having access to the training set.

Contributions. To sum up, the contributions of this work are as follows:

- Our framework enables the simultaneous unlearning of multiple classes in a single round for an image classification network. This approach significantly reduces computational requirements during untraining and testing, offering greater flexibility compared to existing methods.
- Our approach encapsulates inner network components like convolutional filters or attentive

projections into Weight-Filtering layers, which can selectively activate or deactivate these components to achieve the desired unlearning behavior.

- Additionally, our method implicitly discovers the underlying relationships between convolutional filters or attentive projections and output classes and therefore allows to obtain a representation that can be employed for explainability purposes.
- We conduct experiments on small and medium-scale image classification datasets, employing both CNN-based and ViT-based architectures. The results demonstrate the effectiveness of our proposed approach.

RELATED WORK

Machine unlearning seeks to eliminate specific or sensitive data from pre-trained models, ideally without necessitating a full retraining process. In this realm, Cao *et al.* [3] were pioneers in tackling the machine unlearning problem within traditional machine learning algorithms. However, their approach exhibited limited advantages compared to retraining from scratch.

More recent efforts have been concentrated on incorporating unlearning into deep neural networks. While Izzo *et al.* [4] introduced a technique for removing datapoints from linear models, other works involve a data grouping during training, enabling seamless unlearning by restricting the impact of datapoints on the model learning process [10]. Another approach, instead, is that of monitoring the impact of training datapoints on model parameters, thereby minimizing



Figure 2. Application of Weighted-Filter layers for single-shot multiple class unlearning on CNN-based and ViT-based architectures.

the amount of retraining required when a deletion request is received [5]. However, all the aforementioned methods lead to high storage costs. Another line of research proposed to employ efficient partitioning of training data. Among these works, Ginart *et al.* [11] introduced a method for removing points from clustering by dividing the data into independent partitions, eventually using a distinct trained model for each partition.

Recent unlearning paradigms involve the removal of entire classes from a pre-trained model [6]–[8]. For example, Golatkar *et al.* [12] presented a technique for removing information from intermediate layers of deep learning networks via stochastic gradient descent, that can be also extended to the final activation of the model [13]. Tarun *et al.* [14], instead, introduced an effective class-wise unlearning framework in which an error-maximizing noise matrix for the forget class is learned and then used along with a subset of data the model should retain to update network parameters.

Differently from previous works, we propose to unlearn all classes of a pre-trained model, orthogonally and in a single untraining stage, thus avoiding the need to save multiple unlearned models.

PROPOSED METHOD

Preliminaries

We consider a set of input data $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$ with *n* number of samples, where x_i is the *i*-th sample and y_i is the corresponding label belonging to a set of classes \mathcal{C} . Traditional training aims at identifying a set of weights

 $\hat{w} \in \mathcal{W}$ via an iterative update rule $w_{t+1} = w_t - g(w, x')$, where $g(\cdot, x')$ is a stochastic gradient of a fixed loss function. Once the model has reached convergence and given a set of datapoints \mathcal{D}_f drawn from the same distribution of \mathcal{D} , machine unlearning aims at identifying an update to \hat{w} in the form

$$\hat{w}_{t+1} = \hat{w}_t + g(z'_f), \text{ with } \hat{w}_0 := \hat{w} \text{ and } z'_f \sim \mathcal{D}_f$$
(1)

so that \mathcal{D}_f is to be unlearned. In other words, the final set of weights of the unlearned model should be close to a model trained from scratch on $\mathcal{D} \setminus \mathcal{D}_f$. Depending on the use case, \mathcal{D}_f can be conceived as containing a single datapoint (*i.e.*, item removal), a group of data with similar features or labels (*i.e.*, feature removal), or an entire class (class removal) [2].

Class removal. Forgetting data from single or multiple classes in a trained model presents challenges. Unlearning a class typically involves multiple passes over its datapoints, which is computationally costly compared to item or feature removal. Further, when unlearning different classes is required, the training procedure described in Eq. 1 must be repeated for each class as unlearning one class inevitably shifts \hat{w} towards a configuration in which recovering a previously unlearnt class and then unlearning a new one becomes both cumbersome and inefficient. In the case where one wants a model to selectively unlearn all the N_c classes on which it has been trained, this requires repeating the procedure N_c times starting from \hat{w} and store as many independent checkpoints.

Algorithm 1 General forward step of a WF model

Require: a WF Model M comprising a sequence of layers \mathcal{L} , an input tensor X_{in} , a label Y_{unl}

Ensure: a predicted Label $Y_{out} \leftarrow M(X_{in}, Y_{unl})$

Initialize $X_0 \leftarrow X_{in}$ for $i \in \mathbb{N}_{|\mathcal{L}|}$ do $l \leftarrow \mathcal{L}[i]$ $\hat{w}_l \leftarrow \boldsymbol{\alpha}_l[Y_{unl}] \odot \hat{w}_l$ $X_{i+1} \leftarrow l(X_i, \hat{w}_l)$ end for return $\operatorname{argmax}(X_{|\mathcal{L}|})$

Single-shot multiple class unlearning

Leveraging the known association between internal components and classes [9], our unlearning procedure spurs the model to recover these connections and store them in a learnable memory matrix. This is obtained by encapsulating internal network components in a Weight-Filtering (WF) layer, enabling orthogonal removal of each class in a single untraining step and resulting in a single final model checkpoint. At test time, one can simply select one of the N_c classes and instruct the model to behave as if it was unlearned on the class of choice. An overview is shown in Figure 2.

Intuition and single-class scenario. We introduce WF-Net with the description of the case in which we restrict to unlearning a single class only. Given a model $M : \mathcal{X} \to \mathcal{Y}$, a layer l and its set of learned parameters \hat{w}_l , we encapsulate l into our WF-layer. The WF then incorporates the original layer and a learnable memory vector α_l having a shape of $1 \times \mathcal{K}$, where \mathcal{K} is the cardinality of the original learnable weights. The memory vector is meant to mask the learned weights \hat{w}_l and filter the information reaching the following deeper layers. The parameters of α_l are real numbers, which are squeezed into the interval [0, 1] by applying a sigmoid function.

Concretely, α_l modulates the contributions of the inner weights of the layer, thus obtaining a modified layer $\hat{l}(\cdot)$ which behaves as

$$\hat{l}(\cdot) = l(\cdot, \alpha_l \odot \hat{w}_l).$$
⁽²⁾

The modified layer equals the original layer when all the elements of α_l are set to 1. Analogously, α_l can act as an unlearned layer, by dropping some of the α scores to 0. During this procedure, the model's original parameters are kept frozen (only α_l is actually trained). Recalling the dependency between inner components and classes, a good unlearning procedure should then drop the minimum, yet optimal, set of α scores to 0, to completely forget the objective, unlearning class. While **Multi-class** scenario. the abovementioned approach can provide an alternative to classic unlearning for the single-class case, it does not solve the time and storage issues arising when dealing with more classes is required. However, since our procedure exclusively relies on learnable memory vectors, it can be easily extended to the multi-class scenario. By extending α_l to $N_c \times \mathcal{K}$ matrices, each row vector can memorize class-specific information. At forward time, given an unlearning class label, it will be sufficient to index the corresponding row in α_l and apply the layer (see Algorithm 1). Finally, coupled with a suitable untraining protocol, this allows performing a single procedure to unlearn all classes, orthogonally. Finally, it is noteworthy to mention that our approach could be extended to multi-label datasets as well, by selecting a random groundtruth label as row selectors for the α_l matrices when computing the unlearning loss.

Applications to image architectures

To prove its generalization capabilities, we apply our Weight-Filtering to both CNN- and Transformer-based architectures.

CNNs. In CNN-based architectures, we mask both the convolution kernels and biases, using the output channels as the masking granularity. Given a convolutional layer l with kernels $\hat{w}_l^c \in \mathbb{R}^{C_{\text{out}} \times C_{\text{in}} \times kH \times kW}$ and biases $\hat{w}_l^b \in \mathbb{R}^{C_{\text{out}}}$, we apply two weighting matrices over each output channel, $\alpha_l^c \in \mathbb{R}^{N_{\text{classes}} \times C_{\text{out}}}$ over \hat{w}_l^c and $\alpha_l^b \in \mathbb{R}^{N_{\text{classes}} \times C_{\text{out}}}$ over the biases. This results in masking the C_{out} convolutive filters, with shape $C_{\text{in}} \times kH \times kW$ and having C_{out} activations, weighted by a selected row of α_l^c . As well, each of the C_{out} biases is weighted by a selected row of α_l^b .

Vision Transformers. In the case of a Transformer-based network, we instead mask the weights $\hat{w}_l^c \in \mathbb{R}^{F_{\text{in}} \times F_{\text{out}}}$ and the biases $\hat{w}_l^b \in \mathbb{R}^{F_{\text{out}}}$

of the linear projection generating queries, keys, and values of each attentive layer l.

Training

Taking inspiration from previous works [15], we employ two classification loss functions, namely an *unlearning* loss L_f and a *retaining* loss L_r . The first one encourages the model to forget a specific class c_f , while the second one measures the capacity to retain the information about all the other classes $(c_r \in C \setminus c_f)$. We implement both of them as cross-entropy loss functions.

To avoid the shortcomings of the negative gradient approach [12], obtained by inverting the sign of the gradient, we suggest minimizing the following objective, in which we employ the reciprocal of the forget loss:

$$L = \lambda_0 \sum_{(x,y)\in\mathcal{D}_r} L_r \left(M(x), y \right) + \lambda_1 \sum_{(x,y)\in\mathcal{D}_f} \frac{1}{L_f \left(M(x), y \right)}.$$
 (3)

Consequently, minimizing Eq 3 towards zero, implies $L_f(\cdot)$ to be maximized, and $L_r(\cdot)$ to be minimized.

Regularization. To ensure that only a few elements of α_l are dropped to zero during untraining, we add a regularization term, enforcing the elements of α_l to be nearly all active, except for a closed number of its parameters, which will serve as gates. We, therefore, add a regularizer $R(\hat{\alpha})$ to the final loss function.

Label expansion. To achieve simultaneous untraining of all classes, we avoid physically splitting the dataset \mathcal{D} into two retaining and unlearning subsets \mathcal{D}_r and \mathcal{D}_f , as each sample of the dataset could be employed to unlearn its designated class as well as to retain all the others. Instead, given a randomly sampled mini-batch (of length B), we split it into two halves, to obtain B/2 images and labels which are employed for unlearning and B/2 which are employed for retaining. Samples from the first half are employed, together with their ground-truth label, as row selectors for the α_l matrices, and used to compute the unlearning loss. The second half, instead, is paired with randomly selected rows of α_l , and employed to compute the retaining loss with its ground-truth label.

It shall be noted that during optimization, every image of class c will unlearn over the same row of α_l (e.g., class 0 will unlearn over the 0-th row). When retaining, every class will randomly retain on one of the other rows, thanks to the random selection in α_l . This implies that every class will have full impact over its row while having considerably less impact for the retaining rows. To compensate for this disparity, we replicate the "retaining" component of the mini-batch χ times, pairing it with differently sampled labels, so that every retaining image has a χ -times greater impact. As a result, the retaining loss has a shape of $(\chi, B/2)$, while the untraining loss has a shape of (B/2,). The last step is to average both losses.

EXPERIMENTAL EVALUATION

Datasets

We perform experiments on three well-known image classification datasets: MNIST, CIFAR-10, and ImageNet. MNIST consists of 60,000 training and 10,000 test images, each corresponding to one of the 10 classes representing handwritten digits. CIFAR-10 contains 60,000 images across 10 classes, with 6,000 images per class, split into training (50,000) and validation (10,000) sets. ImageNet comprises 1.28 million training images and 50,000 validation images, each associated with one of the 1,000 classes of the dataset. Evaluation on ImageNet involves a randomly selected subset of the validation set with 5,000 images, featuring five images for each class.

Evaluation Metrics

We evaluate WF-Net in terms of evaluation metrics that measure the unlearning capabilities of the model as well as metrics that estimate its level of explainability.

Accuracy on retain and forget sets. We employ accuracy scores to measure the correctness of predictions for both retain and forget sets. Ideally, the accuracy on the retain set should be close to the one of the original model before unlearning, while the accuracy on the forget set should be close to zero or equal to the accuracy of a model re-trained without samples of the forget class.

Relearn Time (RT). It measures the number of epochs taken by the model to regain full accuracy on a class from the forget set. Following previous works [14], we train the model on 500 random samples from the training set in each epoch, stopping training if the model does not reach the original accuracy after 100 epochs. In our experiments, we report the relearn time averaged over all classes of the dataset, showing in parentheses the number of classes the model was unable to relearn after 100 epochs.

Zero Retrain Forgetting (ZRF) score. It estimates the randomness of the unlearned model by comparing it with a randomly initialized network. In particular, this metric compares the output probabilities of the two models using a Jensen-Shannon divergence (JS). Following previous literature [15], we compute the ZRF score as follows:

$$ZRF = 1 - \frac{1}{N_f} \sum_{i=0}^{N_f} JS(M(x_i), M^*(x_i)), \quad (4)$$

where M is the unlearned model, M^* is its randomly initialized version, x_i is the *i*-th sample from the forget set, and N_f is the number of elements of the forget set. The final score lies between 0 and 1, where a score near 1 corresponds to a completely random behavior of M.

Insertion score. The proposed metric measures explainability map goodness [17]. In a few words, it computes the rise in the confidence score as the image pixels are iteratively reactivated, by relevance, from a baseline (usually a 0 tensor). In our case, α_l matrices contain key information on forgotten classes. Unlike the original implementation, at each iteration, we reactivate a percentage of elements in each α_l sorted by relevancy, evaluate the network's confidence score, and normalize it by baseline model's confidence score, *i.e.* the corresponding model before untraining. Eventually, the average AUC score is taken. The closer the score is to 1, the more significant the features selected by the α_l matrices are.

Deletion score. It represents the counterpart of the insertion score, as it measures the drop in score as the image pixels are zeroed, by relevance, at each iteration until the initial picture becomes a baseline tensor. In our case, we again maneuver our α scores to show the correlation between their values and the information about the unlearned classes. To compute the score, we start from a non-unlearned WF-Net and, iteratively, drop a percentage of them to 0, from most relevant to least relevant ones. Each iteration records the normalized confidence score of the WF-Net for unlearned classes, relative to the baseline model's confidence score. The average AUC of the resulting curve represents the final score.

Implementation and Training Details

We conduct our investigation with various backbones with a model size compatible with both middle and medium-scale datasets, namely VGG-16, ResNet-18, and ViT in its Tiny (ViT-T) and Small (ViT-S) versions.

During untraining, we initialize all α_l to 3, so that $\sigma(\alpha) \simeq 1$ and have a non-zero gradient during the initial stages of learning. We employ a fixed (un)learning rate of 100, which allows us to escape from the minimum reached during the pre-training, and a fixed mini-batch size of 128. In the CNNs case, we employ $\lambda_0 = 1$, $\lambda_1 =$ 10, $\lambda_2 = 1$. In the case of ViT we raise λ_1 to 100. The label expansion factor $\chi = 3$ in all the cases, and eventually, we early stop the untraining procedure with a patience of 10, evaluating the validation loss 5 times per epoch.

Experimental Results

Unlearning performance. We start by assessing the unlearning performance of our approach. Table 1 shows results on MNIST, CIFAR-10, and ImageNet training sets. We compare with the performances of a model trained on the full training set from scratch (termed as "original model"), and with a collection of models trained without each of the classes. We train these models only for MNIST and CIFAR-10, due to the higher number of classes in the ImageNet dataset. The average performance of these models is termed "retrained model". Additionally, we implement two singleclass unlearning methods, namely random labels [16] and negative gradients [12]. In the former, we fine-tune the model using randomly assigned labels for samples from the forget set, while in the latter, the model is fine-tuned on the forget set using negative gradients (i.e., fine-tuned in the direction of gradient ascent). It is worth noting that the unlearning phase of these two methods is performed separately for each class. For this reason, we do not report the results of these methods on the ImageNet dataset.

		MNIST				CIFAR-10				ImageNet-1k		
		$\operatorname{Acc}_r \uparrow$	$\mathrm{Acc}_f\downarrow$	$ZRF\uparrow$	RT ↑	$\operatorname{Acc}_r \uparrow$	$\mathrm{Acc}_f\downarrow$	$ZRF\uparrow$	RT ↑	$\operatorname{Acc}_r\uparrow$	$\mathrm{Acc}_f\downarrow$	$\text{ZRF}\uparrow$
VGG-16	Original Model Retrained Model	99.6 99.4	99.6 0.0	48.0 48.7	74.5 (5)	93.0 89.9	93.0 0.0	48.3 50.1	65.3 (5)	71.2	71.3	0.35
	Random Labels [†] Negative Gradients [†]	99.6 99.4	0.0 0.0	53.3 48.3	40.3 (0) 24.6 (0)	88.7 89.4	0.0 0.0	52.5 51.6	56.5 (5) 76.7 (7)	-	-	-
	WF-Net	73.2	0.0	79.2	27.6 (1)	80.2	18.3	57.4	54.2 (5)	64.1	1.89	0.53
ResNet-18	Original Model Retrained Model	99.6 99.4	99.6 0.0	47.0 48.7	43.6 (0)	93.9 90.5	94.0 0.0	48.0 51.4	> 100 (10)	70.5	70.3	0.35
	Random Labels [†] Negative Gradients [†]	99.4 99.2	0.0 0.0	57.1 52.1	15.7 (0) 20.1 (0)	90.6 92.4	0.0 0.0	59.4 55.2	21.6 (0) 31.3 (3)	-	-	-
	WF-Net	94.0	9.68	63.1	22.7 (0)	79.7	9.25	63.9	94.3 (9)	64.4	1.47	0.40
ViT-T	Original Model Retrained Model	98.9 99.0	98.9 0.0	47.2 50.2	71.4 (7)	78.0 71.2	78.0 0.0	49.8 67.6	52.2 (4)	75.6	75.5	0.34
	Random Labels [†] Negative Gradients [†]	98.3 98.4	0.0 0.0	48.7 54.6	81.9 (8) 71.6 (7)	73.2 74.4	0.0 0.0	57.7 50.1	83.5 (8) 81.7 (8)	-	-	-
	WF-Net	93.5	0.0	47.4	> 100 (10)	73.5	0.0	59.8	> 100 (10)	68.0	2.51	0.44
ViT-S	Original Model Retrained Model	99.0 99.0	98.9 0.0	47.1 49.5	91.5 (8)	85.2 75.5	85.2 0.0	53.6 61.3	51.1 (4)	82.3	82.2	0.34
	Random Labels [†] Negative Gradients [†]	99.5 99.6	0.0 0.0	50.4 52.3	74.1 (7) 44.1 (4)	74.8 75.5	0.0 0.0	50.8 51.7	75.0 (6) 61.7 (5)	-	-	-
	WF-Net	94.0	0.0	48.0	91.0 (9)	74.7	0.0	59.7	81.1 (8)	69.2	9.46	0.45

 Table 1. Machine Unlearning results on MNIST, CIFAR-10 and ImageNet-1k, in comparison with original and retrained models. The † marker indicates single-class unlearning methods.

As it can be observed, WF-Net exhibits a proper unlearning behavior across all classes and models. In particular, we notice that on MNIST and CIFAR-10 it showcases a 0.0 accuracy on the forget sets when trained with VGG-16, ViT-S and ViT-T, underlying that the proposed strategy is effective in removing the knowledge of a class of choice, even without directly updating the model weights and while managing the unlearning of more classes concurrently. In terms of accuracy on the retain set, instead, we observe a restrained loss with respect to the performances of both the full model and of the retrained model. For example, on the ImageNet-1k dataset the CNNbased versions of WF-Net only lose a few accuracy points on the retain set compared to the accuracy of the original model (e.g., from 70.5 to 64.4 with ResNet-18)¹. With Transformer-based architectures, instead, the loss in terms of retain accuracy is slightly more evident (e.g., from 82.3) to 69.2 with ViT-S). When considering the ZRF score, computed on the unlearned classes, we can notice that the performance reached by our approach is always higher than that of the original

model and generally in-line or higher than that of the retrained model, underlying again that our approach can properly unlearn.

Furthermore, our solution performs well also against single-class unlearning methods, although they require separate training for each class to be unlearned. In particular, while in terms of accuracy scores WF-Net performs slightly worse than random labels and negative gradients, it obtains significantly superior results in terms of ZRF score in most cases further confirming the unlearning capabilities of our solution. Considering instead the relearn time, our approach takes a comparable or higher number of epochs to relearn unlearned classes than all other models especially when a ViT-based model is used as backbone.

Insertion and deletion scores. In Table 2 we report the insertion and deletion scores, according to the four model architectures. Modifying the α values has an impact on the scores given to images belonging to unlearned classes, and not on images belonging to other classes. This highlights that the proposed approach can identify filters and attention weights which are responsible of the identification of each class.

Filter visualizations. Finally, in Figure 3 we qualitatively visualize the association between fil-

¹We noticed similar results also for larger backbones like ResNet-50 and ResNet-101, with a retaining accuracy respectively of 72.6 and 74.4 and a forget accuracy equal to 0.

 Table 2. Insertion (Ins) and Deletion (Del) scores of the

 WF-Net model.

	MN	IST	CIFA	AR-10	ImageNet		
	Ins ↑	$\text{Del}\downarrow$	Ins \uparrow	$Del\downarrow$	Ins ↑	$\text{Del}\downarrow$	
VGG-16	0.78	0.09	0.79	0.10	0.73	0.05	
ResNet-18	0.79	0.05	0.69	0.06	0.75	0.04	
ViT-T	0.90	0.07	0.88	0.05	0.60	$\begin{array}{c} 0.04 \\ 0.04 \end{array}$	
ViT-S	0.90	0.08	0.88	0.14	0.75		

ters (or attentive projections) and output classes, as discovered by the α values after unlearning, for randomly selected layers. We proceed to select the top-10 filters for each class and remove those which do not appear in at least two classes, in order to spot relationships between classes. As can be noticed, the proposed approach discovers proper relationships between filters and classes. For example, we can notice how the filter #332 contributes to two different classes (i.e., dogs and *horses*), while the filter #3 contributes to the classes cats, deer, and dogs. These patterns can help to better understand the model behavior and can highlight the underlying relationships within the network. Additional visualizations are shown in Figure 4.

CONCLUSION

We presented a novel approach for singleround multi-class unlearning. Given an image classification network, our approach can unlearn all classes simultaneously, in a single unlearning round. After training, the resulting network can be instructed to behave as if it has been untrained on any of the classes. Experimentally, we validated its effectiveness on classical image datasets and across different architectures, also showing that our approach can discover relationships between convolutional filters or attentive projections and output classes.

ACKNOWLEDGMENTS

This work has been supported by the the EU Horizon project "ELIAS - European Lighthouse of AI for Sustainability" (No. 101120237), funded by the European Commission.

REFERENCES

 Quang-Vinh Dang. Right to be forgotten in the age of machine learning. In *ICADS*, 2021.



Figure 3. Relationships highlighted by WF-Net between the weights of a VGG-16 layer and the CIFAR-10 classes.

- Thanh Tam Nguyen, Thanh Trung Huynh, Phi Le Nguyen, Alan Wee-Chung Liew, Hongzhi Yin, and Quoc Viet Hung Nguyen. A Survey of Machine Unlearning. *arXiv preprint arXiv:2209.02299*, 2022.
- Yinzhi Cao and Junfeng Yang. Towards Making Systems Forget with Machine Unlearning. In *IEEE Sympo*sium on Security and Privacy, 2015.
- Zachary Izzo, Mary Anne Smart, Kamalika Chaudhuri, and James Zou. Approximate data deletion from machine learning models. In *AISTATS*, 2021.
- Lucas Bourtoule, Varun Chandrasekaran, Christopher A Choquette-Choo, Hengrui Jia, Adelin Travers, Baiwu Zhang, David Lie, and Nicolas Papernot. Machine unlearning. In *IEEE Symposium on Security and Privacy*, 2021.
- Quoc Phong Nguyen, Bryan Kian Hsiang Low, and Patrick Jaillet. Variational Bayesian Unlearning. In *NeurIPS*, 2020.
- Seth Neel, Aaron Roth, and Saeed Sharifi-Malvajerdi. Descent-to-delete: Gradient-based methods for machine unlearning. In *ALT*, 2021.
- Thomas Baumhauer, Pascal Schöttle, and Matthias Zeppelzauer. Machine unlearning: Linear filtration for logit-based classifiers. *Machine Learning*, 111(9):3203– 3226, 2022.
- Andong Wang, Wei-Ning Lee, and Xiaojuan Qi. HINT: Hierarchical Neuron Concept Explainer. In *CVPR*, 2022.
- Yinjun Wu, Edgar Dobriban, and Susan Davidson. Deltagrad: Rapid retraining of machine learning models. In *ICML*, 2020.
- Antonio Ginart, Melody Guan, Gregory Valiant, and James Y Zou. Making Al Forget You: Data Deletion in Machine Learning. In *NeurIPS*, 2019.
- Aditya Golatkar, Alessandro Achille, and Stefano Soatto. Eternal sunshine of the spotless net: Selective forgetting in deep networks. In CVPR, 2020.



Figure 4. Application of Weighted-Filter layers for single-round multiple class unlearning on the CIFAR-10 dataset on different layers, using VGG-16 (first row), ResNet-18 (second row), and ViT-T (third row).

- Aditya Golatkar, Alessandro Achille, and Stefano Soatto. Forgetting outside the box: Scrubbing deep networks of information accessible from input-output observations. In ECCV, 2020.
- Ayush K Tarun, Vikram S Chundawat, Murari Mandal, and Mohan Kankanhalli. Fast Yet Effective Machine Unlearning. *IEEE Trans. on Neural Networks and Learning Systems*, 2023.
- Vikram S Chundawat, Ayush K Tarun, Murari Mandal, and Mohan Kankanhalli. Can Bad Teaching Induce Forgetting? Unlearning in Deep Networks using an Incompetent Teacher. In AAAI, 2023.
- Tomohiro Hayase, Suguru Yasutomi, and Takashi Katoh Selective Forgetting of Deep Networks at a Finer Level than Samples. arXiv preprint arXiv:2012.11849, 2020.
- Vitali Petsiuk, Abir Das, and Kate Saenko. RISE: Randomized Input Sampling for Explanation of Blackbox Models. In *BMVC*, 2018.

Samuele Poppi is a PhD student of the Italian National PhD Program in AI at the University of Modena and Reggio Emilia, 41125 Modena, Italy. His research interests include explainable AI and machine unlearning. Contact him at samuele.poppi@unimore.it.

Sara Sarto is currently pursuing the PhD degree in

ICT at the University of Modena and Reggio Emilia, 41125 Modena, Italy. Her research interests include image captioning, cross-modal retrieval, and attentive models. Contact her at sara.sarto@unimore.it.

Marcella Cornia is a tenure track assistant professor with the University of Modena and Reggio Emilia, 42121 Reggio Emilia, Italy. She has coauthored more than 70 publications mainly on vision-and-language integration, attentive and saliency models. Contact her at marcella.cornia@unimore.it.

Lorenzo Baraldi is a tenure track assistant professor with the University of Modena and Reggio Emilia, 41125 Modena, Italy. He has coauthored more than 100 publications, mainly working on video understanding, deep learning, and multimedia. Contact him at lorenzo.baraldi@unimore.it.

Rita Cucchiara is a full professor at the University of Modena and Reggio Emilia, 41125 Modena, Italy, where she is the director of the Artificial Intelligence Research and Innovation Center, coordinating the AlmageLab research lab and the Modena ELLIS Unit. Her research interests are related to computer vision, pattern recognition, and multimedia. Contact her at rita.cucchiara@unimore.it.