

Regularizing Second-Order Influences for Continual Learning

Zhicheng Sun¹, Yadong Mu^{1,2*}, Gang Hua³

¹Peking University, ²Peng Cheng Laboratory, ³Wormpex AI Research

{sunzc,myd}@pku.edu.cn, ganghua@gmail.com

Abstract

Continual learning aims to learn on non-stationary data streams without catastrophically forgetting previous knowledge. Prevalent replay-based methods address this challenge by rehearsing on a small buffer holding the seen data, for which a delicate sample selection strategy is required. However, existing selection schemes typically seek only to maximize the utility of the ongoing selection, overlooking the interference between successive rounds of selection. Motivated by this, we dissect the interaction of sequential selection steps within a framework built on influence functions. We manage to identify a new class of second-order influences that will gradually amplify incidental bias in the replay buffer and compromise the selection process. To regularize the second-order effects, a novel selection objective is proposed, which also has clear connections to two widely adopted criteria. Furthermore, we present an efficient implementation for optimizing the proposed criterion. Experiments on multiple continual learning benchmarks demonstrate the advantage of our approach over state-of-the-art methods. Code is available at <https://github.com/feifeiobama/InfluenceCL>.

1. Introduction

The ability to continually accumulate knowledge is a hallmark of intelligence, yet prevailing machine learning systems struggle to remember old concepts after acquiring new ones. *Continual learning* has hence emerged to tackle this issue, also known as *catastrophic forgetting*, and thereby enable the learning on long task sequences which are frequently encountered in real-world applications [15, 17]. Amongst a variety of valid methods, the replay-based approaches [33, 39] achieve evidently strong results by buffering a small *coreset* of previously seen data for rehearsal in later stages. Due to the rigorous constraints on memory overhead, the replay buffer needs to be carefully maintained such that only the samples deemed most critical

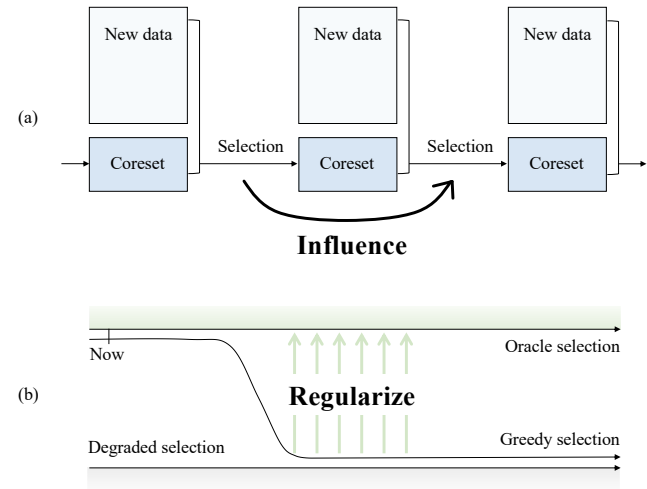


Figure 1. (a) In continual learning, earlier coreset selection exerts a profound influence on subsequent steps through the data flow. (b) By ignoring this, a greedy selection strategy can degrade over time. Therefore, we propose to model and regularize the influence of each selection on the future.

for preserving overall performance are selected during the learning procedure.

Despite much effort in sophisticated design of coreset selection criteria [3, 7, 47, 52], existing works are mostly developed for an oversimplified scenario, where each round of selection is considered isolated and primarily focused on refining single-round performance. For example, Borsos *et al.* [7] greedily minimize the empirical risk on the currently available subset of data. However, as shown in Fig. 1a, the actual selection process within continual learning is rather different in that each prior selection defines the input for subsequent selection and therefore has an inevitable impact on the later decision. Neglecting such interactions between successive selection steps, as in previous practice, may result in a degraded coreset after the prolonged selection process. To maximize overall performance, an ideal continual learning method should take into account the future influence of each round of selection, which remains unresolved due to the obscure role of intermediate selection results.

*Corresponding author.

This work models the interplay among multiple selection steps with the classic tool of influence functions [21, 29], which estimates the effect of each training point by perturbing sample weights. Similarly, we begin by upweighting two samples from consecutive time steps, and uncover that the latter one’s influence evaluation is biased due to the consequent perturbations in test gradient and coreset Hessian. This gives rise to a new class of second-order influences that interfere with the initial influence-based selection strategy. To be specific, the selection tends to align across different rounds, favoring those samples with a larger gradient projection along the previously upweighted ones. It further suggests that some unintended adverse effects of the early selection steps, such as incidental bias introduced into the replay buffer, will be amplified after rounds of selection and impair the final performance.

To address the newly discovered disruptive effects, we propose to regularize each round of selection beforehand based on its expected second-order influence, as illustrated in Fig. 1b. Intuitively, the selection with a large magnitude of second-order influence will substantially interfere with future influence estimates, and therefore should be avoided. However, the magnitude itself cannot be precalculated for direct guidance of the selection, since it is related to some unknown future data. We instead derive a tractable upper bound for the magnitude which results in the form of ℓ_2 -norm, and integrated it with the vanilla influence functions to serve as the final selection criterion. The proposed regularizer can be interpreted with clarity as it equates to a couple of existing criteria ranging from gradient matching to diversity, under varying simplifications. Finally, we present an efficient implementation that tackles the technical challenges in selecting with neural networks.

Our contributions are summarized as below:

- We investigate the previously-neglected interplay between consecutive selection steps within an influence-based selection framework. A new type of second-order influences is identified, and further analysis states its harmfulness in continual learning.
- A novel regularizer is proposed to mitigate the second-order interference. The regularizer is associated with two other popular selection criteria and can be efficiently optimized with our implementation.
- Comprehensive experiments on three continual learning benchmarks (Split CIFAR-10, Split CIFAR-100 and Split *mini*ImageNet) clearly illustrates that our method achieves new state-of-the-art performance.

2. Related works

Continual learning studies the training of models on a sequence of tasks with potential data distribution shift. It is known for suffering from catastrophic forgetting [34],

where the model abruptly forgets past knowledge after being updated on new tasks. To overcome this effect, three main streams of methods have been developed: weight regularization [27, 32, 54], parameter isolation [2, 41, 43, 53] and memory replay [33, 39, 44]. However, the first two approaches exhibit over-regularization or unconstrained parameter growth on long task sequences [20]. In this work, we focus on the replay-based approach, which stores a small subset of previous data in a memory buffer, to be revisited when learning new tasks.

One of the key components in replay-based methods is the sample selection strategy, which undertakes the task of summarizing representative data into the replay buffer. The pioneering works [33, 39] featured a ring buffer for exemplar management, followed by reservoir sampling [40, 49] and its class-balanced variants [13, 38]. More sophisticated selection criteria have been proposed as well, such as matching the overall gradient [4, 47] or enhancing the diversity of samples [3, 5, 52]. Recently, Borsos *et al.* [7] and Sun *et al.* [45] introduced influence functions [29] to continual learning, providing a fresh and interpretable perspective on sample selection. Our method builds on an influence-based framework but also has clear connections to gradient matching and diversity.

Influence functions, known from robust statistics [21], were recently advocated in machine learning by Koh and Liang [29] to estimate the effect of upweighting a training sample to the model parameter and test loss. Since then, influence functions have served as sample selection criteria for various scenarios such as data-efficient learning [46, 51], class imbalanced learning [36] and noisy label learning [30]. Typically, samples with the highest influence scores are identified and then downweighted or discarded. Closely related to our work are their applications in continual learning [7, 45, 58] and domain adaptation [14], but they were still developed for the single-round selection scenario described earlier and overlooked the high-order interactions between multiple rounds of selection, leading to suboptimal results in the continual learning setup.

There are also efforts made toward more accurate and efficient calculation of influence functions. For example, Basu *et al.* [6] introduced second-order approximations to measure group effects [6, 28] of large, coherent groups of training points. However, their analysis is limited to jointly optimized samples and cannot be applied to continuously incoming data. In the other direction, various speedup tricks [7, 19, 42, 56] have been proposed for over-parameterized neural networks, among which a representative approach is to use neural tangent kernels [23] as the proxy model [7, 56]. We adapt some of these techniques to the proposed second-order influence functions and derive an efficient selection criterion for continual learning.

3. Influence functions for continual learning

3.1. Problem formulation

We consider learning on a continuous stream of data $\mathcal{Z} = \{z_i\}_{i=1}^n$ assuming that only a fraction of samples $\mathcal{Z}_t \subset \mathcal{Z}$ can be accessed at each time step t . The main challenge in such a learning paradigm is to retain performance on the seen data $\mathcal{Z}_{1:t} = \bigcup_{i=1}^t \mathcal{Z}_i$ in later stages. To achieve this, we adopt a most straightforward and effective approach that stores a few samples \mathcal{C}_t in a replay buffer as the representatives of previous data $\mathcal{Z}_{1:t}$, where the memory size $|\mathcal{C}_t| \leq m$ and $m \ll n$. In consequence, it is crucial to maintain a high-quality subset throughout the training process [3, 52]. Formally, let θ be the model parameters and $L(z_i, \theta)$ be the loss on training point z_i . Then, our selection goal is to preserve model performance on $\mathcal{Z}_{1:t}$ by replaying on \mathcal{C}_t , which is formulated as:

$$\begin{aligned} \min_{\mathcal{C}_t \subset \mathcal{C}_{t-1} \cup \mathcal{Z}_t, |\mathcal{C}_t| \leq m} \sum_{z_i \in \mathcal{Z}_{1:t}} L(z_i, \hat{\theta}) \\ \text{s.t. } \hat{\theta} = \arg \min_{\theta} \sum_{z_i \in \mathcal{C}_t} L(z_i, \theta). \end{aligned} \quad (1)$$

Though its main form is consistent with that of the well-studied coreset selection problem [7], note that the outer objective is intractable since $\mathcal{Z}_{1:t}$ is partly unavailable.

In the following sections, we first present a baseline selection strategy produced by the vanilla influence functions, then showcase its limitation in handling continual selection and propose our improved version.

3.2. Influence-based selection

Influence functions [29] provide an efficient approximation for solving the coreset selection problem by perturbing sample weights, for which the previous exact solution requires expensive leave-one-out retraining.

Before diving in, we first convert Eq. (1) into a tractable problem by leveraging the empirical risk on $\mathcal{C}_{t-1} \cup \mathcal{Z}_t^1$ as a proxy for the original test loss on $\mathcal{Z}_{1:t}$, presuming a close correlation between them:

$$\begin{aligned} \min_{\mathcal{C}_t \subset \mathcal{C}_{t-1} \cup \mathcal{Z}_t, |\mathcal{C}_t| \leq m} \sum_{z_i \in \mathcal{C}_{t-1} \cup \mathcal{Z}_t} L(z_i, \hat{\theta}) \\ \text{s.t. } \hat{\theta} = \arg \min_{\theta} \sum_{z_i \in \mathcal{C}_t} L(z_i, \theta). \end{aligned} \quad (2)$$

Solving this involves uncovering the effect of selecting or discarding some training sample $z \in \mathcal{C}_{t-1} \cup \mathcal{Z}_t$, for which a small weight ϵ is added to the interested sample, and then the optimal point in the inner optimization becomes:

$$\hat{\theta}_{\epsilon, z} = \arg \min_{\theta} \sum_{z_i \in \mathcal{C}_t} L(z_i, \theta) + \epsilon L(z, \theta). \quad (3)$$

¹The loss on \mathcal{Z}_t needs to be reweighted by a constant to balance with the loss on \mathcal{C}_{t-1} . We omit the weight for simplicity unless otherwise stated.

Let $\hat{\theta}_t = \hat{\theta}_{\epsilon, z}|_{\epsilon=0}$ denote the initial optimal parameters and $H_{\hat{\theta}_t} = \sum_{z_i \in \mathcal{C}_t} \nabla_{\theta}^2 L(z_i, \hat{\theta}_t)$ denote the Hessian which by assumption is positive definite. A classic result by Cook and Weisberg [16] yields the change in model parameters as $\left. \frac{d\hat{\theta}_{\epsilon, z}}{d\epsilon} \right|_{\epsilon=0} = -H_{\hat{\theta}_t}^{-1} \nabla_{\theta} L(z, \hat{\theta}_t)$, from which we can derive the influence of upweighting z on the outer loss in Eq. (2):

$$\begin{aligned} \mathcal{I}(z) &= \sum_{z_i \in \mathcal{C}_{t-1} \cup \mathcal{Z}_t} \left. \frac{dL(z_i, \hat{\theta}_{\epsilon, z})}{d\epsilon} \right|_{\epsilon=0} \\ &= - \sum_{z_i \in \mathcal{C}_{t-1} \cup \mathcal{Z}_t} \nabla_{\theta} L(z_i, \hat{\theta}_t)^{\top} H_{\hat{\theta}_t}^{-1} \nabla_{\theta} L(z, \hat{\theta}_t). \end{aligned} \quad (4)$$

We denote the inverse Hessian-vector product therein as $s_t = H_{\hat{\theta}_t}^{-1} \sum_{z_i \in \mathcal{C}_{t-1} \cup \mathcal{Z}_t} \nabla_{\theta} L(z_i, \hat{\theta}_t)$ for future use.

To minimize the outer loss, the coreset \mathcal{C}_t should contain the samples with the lowest negative influence scores. It can be efficiently solved by a greedy strategy, which starts from the full set $\mathcal{C}_t = \mathcal{C}_{t-1} \cup \mathcal{Z}_t$ and discards the sample z with the largest influence at each iteration until the memory constraint $|\mathcal{C}_t| \leq m$ is met. While such a strategy is qualified in a variety of evaluations [7, 50, 51, 58], we argue that it is not the answer to continual learning by dissecting the joint effect of influence-based selection in a sequence.

3.3. Second-order influences in continual selection

In continual learning, each prior selection determines the input for subsequent steps and thus influences the effectiveness of future selection, *i.e.*, it may interfere with the later evaluation of the selection criterion.

To model such an interaction, we consider two samples z and z' from consecutive steps t and $t+1$. With the previous sample z upweighted by ϵ , the next round of selection is also affected, by a drift in the influence score of z' . Specifically, there are two cases depending on whether z and z' are jointly optimized in the following step:

(1) The samples z and z' are not jointly optimized, meaning that the previous sample z is excluded from the coreset in the next round and only serves as a test point. As a result, only the outer loss in Eq. (2) is affected, with the influence score of the subsequent candidate z' changed to:

$$\begin{aligned} \mathcal{I}_{\epsilon, z}(z') &= - \left(\sum_{z_i \in \mathcal{C}_t \cup \mathcal{Z}_{t+1}} \nabla_{\theta} L(z_i, \hat{\theta}_{t+1}) + \epsilon \nabla_{\theta} L(z, \hat{\theta}_{t+1}) \right)^{\top} \\ &\quad H_{\hat{\theta}_{t+1}}^{-1} \nabla_{\theta} L(z', \hat{\theta}_{t+1}). \end{aligned} \quad (5)$$

By taking the derivative w.r.t. ϵ , it yields the influence of upweighting z on the influence score of z' . Alternatively, it can be viewed as the second-order influence of two samples

from consecutive selection steps on the outer loss:

$$\begin{aligned} \mathcal{I}^{(2)}(z, z') &= \left. \frac{d\mathcal{I}_{\epsilon, z}(z')}{d\epsilon} \right|_{\epsilon=0} \\ &= -\nabla_{\theta} L(z, \hat{\theta}_{t+1})^{\top} H_{\hat{\theta}_{t+1}}^{-1} \nabla_{\theta} L(z', \hat{\theta}_{t+1}). \end{aligned} \quad (6)$$

We will soon analyze its effect using the geometric view of projection.

(2) The samples z and z' are jointly optimized by the inner objective in Eq. (2). In this case, the inner Hessian is also changed, in the direction of $H_{\hat{\theta}_{t+1}, z} := \nabla_{\theta}^2 L(z, \hat{\theta}_{t+1})$. Hence, the influence score of the following sample z' is modified to:

$$\begin{aligned} \mathcal{I}_{\epsilon, z}(z') &= - \left(\sum_{z_i \in \mathcal{C}_t \cup \mathcal{Z}_{t+1}} \nabla_{\theta} L(z_i, \hat{\theta}_{t+1}) + \epsilon \nabla_{\theta} L(z, \hat{\theta}_{t+1}) \right)^{\top} \\ &\quad (H_{\hat{\theta}_{t+1}} + \epsilon H_{\hat{\theta}_{t+1}, z})^{-1} \nabla_{\theta} L(z', \hat{\theta}_{t+1}). \end{aligned} \quad (7)$$

Similarly, the second-order influence of z' and z can be obtained as follows (see Supplementary for the derivation):

$$\begin{aligned} \mathcal{I}^{(2)}(z, z') &= \left. \frac{d\mathcal{I}_{\epsilon, z}(z')}{d\epsilon} \right|_{\epsilon=0} \\ &= -(\nabla_{\theta} L(z, \hat{\theta}_{t+1}) - H_{\hat{\theta}_{t+1}, z} s_{t+1})^{\top} H_{\hat{\theta}_{t+1}}^{-1} \nabla_{\theta} L(z', \hat{\theta}_{t+1}). \end{aligned} \quad (8)$$

Compared to the first case, there is an additional term related to the Hessian perturbation, whose effect will be discussed in Sec. 3.5. Nevertheless, the prior term in Eq. (6) is fully retained.

Second-order influences compromise sample diversity.

For an intuitive understanding of the new influences, let $\langle u, v \rangle = u^{\top} H_{\hat{\theta}_{t+1}}^{-1} v$ define an inner product over the gradient space and induce the notion of projection. In both cases, the sample z' with a larger gradient projection onto the previous z is estimated with a lower influence and thus more likely to be selected. This drives the influence-based selection strategy to include more similar samples to the buffer and expel dissimilar ones, leading to a lack of diversity.

Second-order influences amplify memory bias.

While the above analysis is based on a real perturbation to sample weights, it could also be a simulated perturbation on the gradient statistics of the coreset. Such perturbations are common in the learning process, as each selection step may introduce an incidental bias to the buffer. Like real samples, they can produce the same second-order effect of attracting similar ones into the buffer. Although these perturbations are expected to have zero mean like white noises, the earlier perturbations have taken a greater effect in accumulating memory bias, which results in inaccurate influence estimates and degradation of future selection.

3.4. Regularizing second-order influences

To mitigate the harmful second-order effects without imposing additional memory overhead, we suggest regularizing each coreset selection step in advance, so that future influence estimates will be less error-prone.

Suppose we are at the t -th step, the goal is to minimize its interference with some sample z' of the next step, namely the total effect of discarding the subset $\bar{\mathcal{C}}_t = \mathcal{C}_{t-1} \cup \mathcal{Z}_t - \mathcal{C}_t$, which can be computed by summing Eq. (6) or Eq. (8)². However, it cannot be predetermined which equation to be used, so a hyperparameter $\mu \in [0, 1]$ is introduced as the ratio for the latter case, allowing us to consider a weighted sum of the two second-order influences:

$$\begin{aligned} \Delta \mathcal{I}(z') &\approx - \sum_{z \in \bar{\mathcal{C}}_t} \mathcal{I}^{(2)}(z, z') \cdot 1 \\ &= \sum_{z \in \bar{\mathcal{C}}_t} (\nabla_{\theta} L(z, \hat{\theta}_{t+1}) - \mu H_{\hat{\theta}_{t+1}, z} s_{t+1})^{\top} H_{\hat{\theta}_{t+1}}^{-1} \nabla_{\theta} L(z', \hat{\theta}_{t+1}). \end{aligned} \quad (9)$$

Ideally, the magnitude of total influence $\Delta \mathcal{I}(z')$ should be very small so that future selection will receive less interference. However, this term is intractable since it is associated with an unknown sample z' , so we turn to optimize its upper bound given by the Cauchy–Schwarz inequality:

$$\begin{aligned} |\Delta \mathcal{I}(z')| &\leq \left\| \sum_{z \in \bar{\mathcal{C}}_t} (\nabla_{\theta} L(z, \hat{\theta}_{t+1}) - \mu H_{\hat{\theta}_{t+1}, z} s_{t+1}) \right\| \\ &\quad \times \left\| H_{\hat{\theta}_{t+1}}^{-1} \nabla_{\theta} L(z', \hat{\theta}_{t+1}) \right\|, \end{aligned} \quad (10)$$

We omit the second norm that depends mainly on the upcoming data, and employ the first one to regularize the ongoing selection, during which the yet unknown θ_{t+1} and s_{t+1} are approximated with θ_t and s_t , respectively:

$$\mathcal{R}(\mathcal{C}_t) = \left\| \sum_{z \in \bar{\mathcal{C}}_t} (\nabla_{\theta} L(z, \hat{\theta}_t) - \mu H_{\hat{\theta}_t, z} s_t) \right\|. \quad (11)$$

Proposed selection criterion. A weighted sum of the first-order influence and second-order regularizer balanced by a hyperparameter ν is adopted to guide the selection process. The final objective at the t -th step is as follows:

$$\min_{\mathcal{C}_t \subset \mathcal{C}_{t-1} \cup \mathcal{Z}_t, |\mathcal{C}_t| \leq m} \sum_{z \in \mathcal{C}_t} \mathcal{I}(z) + \nu \mathcal{R}(\mathcal{C}_t). \quad (12)$$

Next, we will interpret it by unveiling its connection to two widely adopted selection criteria, and then present an efficient implementation for deep models in practical use.

²With first-order approximations here, the group effect [6] of dropped datapoints can be fairly neglected.

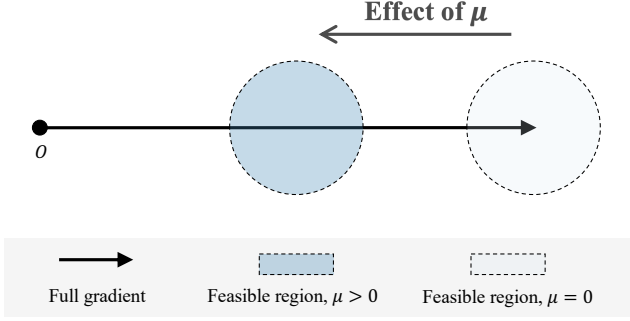


Figure 2. Geometric interpretation of our proposed regularizer. As μ grows larger from zero, the regularizer shifts from the initial gradient matching to introducing more gradient diversity along the predominant direction.

3.5. Connection to gradient matching and diversity

The following case studies of Eq. (11) allow to establish clear connections between the proposed regularizer and two commonly used selection criteria: gradient matching [4, 26, 57] and sample diversity [3, 5, 52].

Start by setting $\mu = 0$, which corresponds to having full confidence that the currently discarded samples will not be favored in the future, and then the regularizer becomes:

$$\mathcal{R}(\mathcal{C}_t) = \left\| \sum_{z \in \mathcal{C}_{t-1} \cup \mathcal{Z}_t} \nabla_{\theta} L(z, \hat{\theta}_t) - \sum_{z \in \mathcal{C}_t} \nabla_{\theta} L(z, \hat{\theta}_t) \right\|. \quad (13)$$

It turns out to minimize the Euclidean distance between the full gradient and the coreset gradient, which is equivalent to gradient matching. For an intuitive look, we consider a hard constraint $\mathcal{R}(\mathcal{C}_t) < \delta$ and mark the feasible region of the coreset gradient with light blue in Fig. 2.

In the more general case $\mu > 0$, we temporarily ignore the sample information carried in Hessian for simplicity, by assuming that the Hessian is identical for all training samples. Then there is:

$$\mathcal{R}(\mathcal{C}_t) = \left\| (1 - \alpha\mu) \sum_{z \in \mathcal{C}_{t-1} \cup \mathcal{Z}_t} \nabla_{\theta} L(z, \hat{\theta}_t) - \sum_{z \in \mathcal{C}_t} \nabla_{\theta} L(z, \hat{\theta}_t) \right\|, \quad (14)$$

where α is a coefficient related to the coreset size only. As illustrated in Fig. 2, the feasible region is moved in the opposite direction of the total gradient. This promotes samples whose gradients are less aligned or even conflicting with the main direction to be included in the coreset, thus encouraging diversity. In a holistic view, the hyperparameter μ plays the role of balancing gradient matching and diversity.

While these simplifications equate our proposed regularizer to a couple of intuitive criteria, the original form excels in the additional incorporation of Hessian-related information, which will be empirically demonstrated in the experiments.

3.6. Implementation for neural networks

This section presents an efficient implementation of optimizing the selection criterion with deep neural networks. Since similar issues for the vanilla influence functions have been intensively addressed [7, 29, 45], the focus here is on the optimization of the newly proposed regularizer.

A key difference of our second-order regularizer is that it cannot be partitioned into independent terms associated to different samples, which hinders application of the aforementioned greedy selection strategy. We overcome this challenge by leveraging the first-order Taylor expansion of the regularizer. To this end, the original variable \mathcal{C}_t is replaced by a continuous vector w_t , where each element $w_{t,i} = \mathbb{1}(z_i \in \mathcal{C}_t)$ indicates whether the corresponding sample is included in the coreset. After that, the regularizer can be transformed into an equivalent form:

$$\mathcal{R}(w_t) = \left\| \sum_{z_i \in \mathcal{C}_{t-1} \cup \mathcal{Z}_t} (1 - w_{t,i}) (\nabla_{\theta} L(z_i, \hat{\theta}_t) - \mu H_{\hat{\theta}_t, z_i} s_t) \right\|, \quad (15)$$

whose first-order Taylor approximation w.r.t. w_t is a linear combination of uncorrelated terms (see Supplementary for details) and can therefore be addressed with greedy heuristics. We then implement an iterative optimization algorithm for selection that starts from an exterior point $w_t = \mathbf{1}$, and modifies w_t during each iteration to drop the sample z_i with the largest $I(z_i) + \nu \nabla_{w_{t,i}} \mathcal{R}(w_t)$, until the memory constraint $\|w_t\|_0 \leq m$ is satisfied. Such a strategy of greedily minimizing the proposed criterion is found to be sufficient in later experiments.

To further enable efficient inference on non-convex and over-parameterized deep neural networks, we are inspired by Borsos *et al.* [7] and employ neural tangent kernels [23] as the proxy model for computing the regularizer in Eq. (15). The inverse Hessian-vector product s_t therein can be effectively computed using the conjugate gradient method, with a small damping term added to the Hessian to ensure that it is positive definite.

Computational cost. The performance bottleneck lies in computing the inverse Hessian-vector product s_t , for which the conjugate gradient method takes $O(m'p^2)$ time to yield an exact solution [29], where $m' = |\mathcal{C}_{t-1} \cup \mathcal{Z}_t|$ and p is the number of parameters. Since this part of cost is shared with the first-order influence function, our regularizer introduces only a small overhead of $O(m'p)$ in evaluating the Hessian-vector product $H_{\hat{\theta}_t, z_i} s_t$ each time. In addition, the adopted neural tangent kernel approximation significantly eases the burden by reducing the number of parameters.

Our computational efficiency will be further verified by the runtime comparison in Sec. 4.3 and the per-step overhead analysis in the supplementary material.

Method	Class-incremental				Task-incremental				
	$m = 300$		$m = 500$		$m = 300$		$m = 500$		
	ACC (%)	BWT (%)	ACC (%)	BWT (%)	ACC (%)	BWT (%)	ACC (%)	BWT (%)	
Non-IF	GEM [33]	37.51	-70.48	36.95	-69.76	89.34	-9.09	90.42	-7.88
	A-GEM [12]	20.02	-95.68	20.01	-95.69	85.52	-14.07	86.45	-12.83
	ER [13]	34.19	-78.18	40.45	-70.36	88.97	-9.95	90.60	-7.74
	GSS [3]	35.89	-75.80	41.96	-68.24	88.05	-10.63	90.38	-7.73
	ER-MIR [1]	38.53	-72.72	42.65	-67.50	88.50	-10.33	90.63	-7.62
	GDUMB [38]	36.92	-	44.27	-	73.22	-	78.06	-
	HAL [11]	24.45	-83.56	27.94	-80.01	79.90	-14.39	81.84	-12.73
	GMED [24]	38.12	-73.16	43.68	-66.21	88.91	-9.76	89.72	-8.75
IF	Vanilla IF	41.76	-68.59	47.14	-62.20	90.67	-7.65	91.06	-7.36
	MetaSP [45]	43.76	-66.37	50.10	-58.39	89.91	-9.00	91.41	-7.36
	Ours	48.62	-60.24	53.07	-54.44	91.52	-6.94	92.53	-5.46

Table 1. Comparison with influence function (IF)-based methods and non-IF-based methods on Split CIFAR-10 under different memory constraints. For the compared methods, we report the results summarized in [45] in chronological order. The task-incremental setting asks the model to classify within each task, while the class-incremental setting requires predicting both task identity and class label.

4. Experiments

4.1. Experimental setup

Datasets. We conduct experiments on three continual learning benchmarks: (1) Split CIFAR-10 [55] splits the original CIFAR-10 [31] dataset into 5 disjoint subsets, where each subset comprises 2 classes. (2) Split CIFAR-100 [55] is constructed from the CIFAR-100 [31] dataset, containing 10 tasks with disjoint class labels. (3) Split *miniImageNet* [13, 45] derives from the few-shot learning dataset *miniImageNet* [48], a subset of ImageNet [18] with 100 classes and 600 images per class. The dataset is divided equally to create 5 sequential tasks, with each image resized to 32×32 . For all three benchmarks, we follow their original papers in splitting training and test sets.

Metrics. Two evaluation metrics are employed, including Average Accuracy (ACC) and Backward Transfer (BWT) [33], where ACC is the average accuracy after the model has been trained on all tasks and BWT indicates the average forgetting of all previous tasks. Formally, they are defined as:

$$\text{ACC} = \frac{1}{T} \sum_{i=1}^T R_{T,i}, \quad \text{BWT} = \frac{1}{T-1} \sum_{i=1}^{T-1} R_{T,i} - R_{i,i}, \quad (16)$$

where T is the number of tasks and $R_{i,j}$ is the accuracy of the model on the j -th task after learning i tasks. As done in [9, 45], we measure the metrics under both task-incremental and class-incremental settings. The latter is particularly difficult as it does not provide task identity for each sample at test time.

Baselines. Our method is compared against nine replay-based competitors: ER [13], GEM [33], A-GEM [12], GSS [3], ER-MIR [1], GDUMB [38], HAL [11], GMED [24] and MetaSP [45]. Additionally, a base strategy that uses the vanilla influence functions (IF) introduced in Sec. 3.2 is considered. To make a fair comparison, we first reproduce the main experiments of [45] and then report the results therein.

Implementation details. We adopt ResNet-18 [22] as the backbone architecture. The model is optimized by SGD for 50 epochs per task, with fixed batch size and replay batch size of both 32, as in [10]. The learning rate is set to 0.1 on Split CIFAR-10 and Split CIFAR-100, and 0.03 on Split *miniImageNet*. The other hyperparameters are empirically set as $\mu = 0.5$ and $\nu = 0.01$ by default, whose sensitivity analysis will be given in Sec. 4.3. In the training, we use cross entropy loss for replay samples and apply standard data augmentations of random cropping and horizontal flipping. To limit the computational overhead of our method, the replay buffer is only updated during the last epoch of each task. The continual learning scenarios are provided by Mammoth [8, 9] which relies on PyTorch [37]. For calculating Neural Tangent Kernels, the library of [35] is employed.

4.2. Main results

The results on the Split CIFAR-10 benchmark are summarized in Tab. 1. Overall, our approach achieves the best continual learning performance in all metrics under various evaluation settings and memory constraints, suggesting its superiority over state-of-the-art methods. From the results, we also observe that: (1) The first-order influence-

Method	Class-incremental				Task-incremental				
	$m = 500$		$m = 1000$		$m = 500$		$m = 1000$		
	ACC (%)	BWT (%)	ACC (%)	BWT (%)	ACC (%)	BWT (%)	ACC (%)	BWT (%)	
Non-IF	GEM [33]	15.91	-77.07	22.79	-68.32	68.68	-18.72	73.71	-12.81
	A-GEM [12]	9.31	-85.18	9.27	-84.88	55.28	-34.10	55.95	-33.01
	ER [13]	13.75	-81.64	17.56	-77.52	66.82	-22.73	71.74	-17.40
	GSS [3]	14.01	-80.02	17.87	-76.04	66.80	-21.44	71.98	-16.06
	ER-MIR [1]	13.49	-82.09	17.56	-77.59	66.18	-23.60	71.20	-18.10
	GDUMB [38]	11.11	-	15.75	-	36.40	-	43.25	-
	HAL [11]	8.20	-65.70	10.59	-63.86	44.98	-25.17	50.07	-20.61
	GMED [24]	14.56	-80.68	18.67	-76.23	68.82	-20.53	73.91	-15.10
IF	Vanilla IF	17.49	-77.54	22.75	-72.56	71.74	-17.90	73.25	-17.22
	MetaSP [45]	19.28	-76.13	25.72	-68.69	70.81	-19.74	76.14	-14.32
	Ours	21.15	-73.24	27.99	-64.56	72.53	-17.22	74.27	-16.37

(a) Split CIFAR-100

Method	Class-incremental				Task-incremental				
	$m = 500$		$m = 1000$		$m = 500$		$m = 1000$		
	ACC (%)	BWT (%)	ACC (%)	BWT (%)	ACC (%)	BWT (%)	ACC (%)	BWT (%)	
Non-IF	A-GEM [12]	10.69	-49.22	10.69	-49.16	18.34	-39.65	18.78	-39.05
	ER [13]	11.00	-50.84	11.35	-50.08	28.97	-28.40	31.59	-24.95
	GSS [3]	11.09	-50.66	11.42	-49.91	28.67	-28.71	31.75	-24.56
	ER-MIR [1]	11.07	-50.46	11.32	-49.92	29.10	-27.95	31.39	-24.89
	GDUMB [38]	6.22	-	7.15	-	16.37	-	17.69	-
	GMED [24]	11.03	-50.23	11.73	-48.93	30.47	-26.02	32.85	-22.69
IF	Vanilla IF	12.08	-48.55	14.64	-47.15	33.74	-21.71	37.55	-19.28
	MetaSP [45]	12.74	-48.84	14.54	-45.52	34.36	-21.70	37.20	-17.83
	Ours	13.63	-47.94	16.15	-43.78	36.46	-19.48	39.61	-16.01

(b) Split *miniImageNet*

Table 2. Comparison with state-of-the-art methods on more challenging benchmarks, including Split CIFAR-100 and Split *miniImageNet*. While most setups follow the previous experiment, a different hyperparameter setting $\mu = 0.75$ is adopted on Split CIFAR-100 with memory size $m = 1000$ empirically.

based methods obtain significant performance gains over traditional methods, but our version almost doubles the improvement under the small memory setting of $m = 300$. By taking into account the second-order influences, it outperforms the state-of-the-art method by 4.86% in ACC and 6.13% in BWT in the challenging class-incremental evaluations, which in turn reflects the non-trivial effect of second-order influences. (2) Though the base strategy vanilla IF is slightly inferior to the state-of-the-art method MetaSP in most cases for it does not introduce any new losses in the model training, it can substantially surpass MetaSP with a simple regularization term on the selection criterion. This verifies the effectiveness of our proposed regularizer.

We further evaluate our method on the more difficult Split CIFAR-100 and Split *miniImageNet* benchmarks. As presented in Tab. 2, our approach brings consistent perfor-

mance improvement over vanilla IF on both benchmarks, by up to 5.24% and 8.0% in terms of ACC and BWT, respectively. Compared to other approaches, it obtains state-of-the-art performance on Split *miniImageNet* in all metrics, and leads ACC by about 2% on three different setups of Split CIFAR-100. Only in the simplest scenario of task-incremental learning on Split CIFAR-100 with $m = 1000$, our method fails to surpass MetaSP, for which we speculate that the high-order interference is alleviated to some extent by the large buffer itself. Nevertheless, our method remains competitive in most evaluations.

4.3. Ablation study and analysis

This section empirically justifies the design of our approach through a series of experiments conducted on Split CIFAR-10 with a fixed memory size of $m = 500$.

Method	Class-incremental		Task-incremental	
	ACC (%)	BWT (%)	ACC (%)	BWT (%)
Vanilla IF	47.14	-62.20	91.06	-7.36
+ matching	49.86	-58.24	92.30	-5.78
+ diversity	50.59	-57.06	90.69	-7.28
+ both	51.61	-56.31	91.93	-6.11
Ours	53.07	-54.44	92.53	-5.46

Table 3. Comparison with alternative regularizers on Split CIFAR-10 with memory size $m = 500$. The diversity-based regularizer is implemented by minimizing the norm of the coreset gradient.

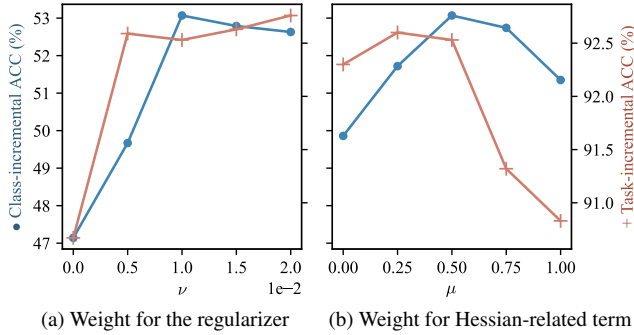


Figure 3. Sensitivity to hyperparameters on Split CIFAR-10 with $m = 500$. They are set to $\nu = 0.01$ and $\mu = 0.5$ by default.

Comparison with other regularizers. To further validate the effectiveness of our regularizer, it is compared to two closely related criteria, namely gradient matching and diversity, in Tab. 3. Our regularizer draws on the strengths from both criteria, including the satisfactory task-incremental performance of gradient matching and the higher class-incremental ACC of gradient diversity. In addition, we compare it with a combination of both criteria, during which the considerable performance advantage of our Hessian-aware regularizer illustrates the benefits from incorporating Hessian-related information. More comparative studies can be found in the supplementary material.

Hyperparameter sensitivity. In order to explore the sensitivity of our strategy to hyperparameter settings, we vary the two coefficients μ and ν and plot the response of model performance. Figure 3a shows that under both evaluation settings, ACC improves consistently with the growth of weight ν on the second-order regularizer until its stability. The situation is more complicated in Fig. 3b, where the two ACC curves exhibit inconsistent trends regarding the weight μ for the Hessian-related term. Recall that in Sec. 3.5 μ is associated with the degree of diversity, so it may be interpreted that the more challenging class-incremental setting requires higher memory diversity, while the simpler task-increment setting accommodates lower diversity.




Figure 4. Evolution of the influence estimation accuracy measured by Kendall's τ [25].

Method	Runtime (h)
ER [13]	0.92
A-GEM [12]	1.22
GSS [3]	0.92
HAL [12]	1.73
Vanilla IF	1.69
MetaSP [45]	2.03
Ours	1.74

Table 4. Running time measured on Split CIFAR-10 using a single NVIDIA 2080 Ti GPU.

Accuracy of influence estimates. To measure accuracy, the previous practice of comparing with the ground-truth given by leave-one-out retraining [6,29] imposes an impractical cost under the continual learning setup. Hence, we turn to a large reservoir sampling buffer to produce unbiased influence estimates and compute their correlation with other predictions. The Kendall rank correlation [25] is adopted since we are more concerned with each sample's rank during selection. As illustrated in Fig. 4, the buffer maintained by vanilla IF dissociates from the unbiased buffer after a short period, producing inaccurate influence estimates. In contrast, our strategy yields a relatively higher agreement, which verifies its effectiveness in mitigating interference.

Running time. Table 4 compares our proposed method with several replay-based competitors in terms of efficiency. Though influence-based approaches generally incur higher time costs in the Hessian-related computations, our method imposes a limited overhead with the newly introduced second-order regularizer. This implies that the above improvements are achieved within comparable training time, demonstrating the high efficiency of our implementation.

5. Concluding remarks

We propose an effective coreset selection strategy for continual learning that addresses the interference within successive selection steps. By dissecting the interactions between consecutive rounds of influence-based selection, a new class of second-order influences is identified, with which the long selection process may gradually lose diversity and accumulate bias. To solve this problem, a novel regularizer is presented which is linked to two other popular criteria but incorporates additional Hessian-related information. Finally, we implement the proposed selection criterion with high efficiency and validate its effectiveness in a variety of comparative experiments.

This work is supported by National Key R&D Program of China (2020AAA0104401), Beijing Natural Science Foundation (Z190001) and Peng Cheng Laboratory Key Research Project No.PCL2021A07.

References

- [1] Rahaf Aljundi, Eugene Belilovsky, Tinne Tuytelaars, Laurent Charlin, Massimo Caccia, Min Lin, and Lucas Page-Caccia. Online continual learning with maximal interfered retrieval. In *NeurIPS*, pages 11849–11860, 2019. [6](#), [7](#)
- [2] Rahaf Aljundi, Punarjay Chakravarty, and Tinne Tuytelaars. Expert gate: Lifelong learning with a network of experts. In *CVPR*, pages 3366–3375, 2017. [2](#)
- [3] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. In *NeurIPS*, pages 11817–11826, 2019. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [4] Lukas Balles, Giovanni Zappella, and Cédric Archambeau. Gradient-matching coresets for continual learning. In *NeurIPS Workshops*, 2021. [2](#), [5](#)
- [5] Jihwan Bang, Heesu Kim, YoungJoon Yoo, Jung-Woo Ha, and Jonghyun Choi. Rainbow memory: Continual learning with a memory of diverse samples. In *CVPR*, pages 8218–8227, 2021. [2](#), [5](#)
- [6] Samyadeep Basu, Xuchen You, and Soheil Feizi. On second-order group influence functions for black-box predictions. In *ICML*, pages 715–724, 2020. [2](#), [4](#), [8](#)
- [7] Zalán Borsos, Mojmir Mutny, and Andreas Krause. Coresets via bilevel optimization for continual learning and streaming. In *NeurIPS*, pages 14879–14890, 2020. [1](#), [2](#), [3](#), [5](#)
- [8] Matteo Boschini, Lorenzo Bonicelli, Pietro Buzzega, Angelo Porrello, and Simone Calderara. Class-incremental continual learning into the extended DER-verse. *TPAMI*, pages 1–16, 2022. [6](#)
- [9] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. In *NeurIPS*, pages 15920–15930, 2020. [6](#)
- [10] Pietro Buzzega, Matteo Boschini, Angelo Porrello, and Simone Calderara. Rethinking experience replay: a bag of tricks for continual learning. In *ICPR*, pages 2180–2187, 2021. [6](#)
- [11] Arslan Chaudhry, Albert Gordo, Puneet Dokania, Philip Torr, and David Lopez-Paz. Using hindsight to anchor past knowledge in continual learning. In *AAAI*, pages 6993–7001, 2021. [6](#), [7](#)
- [12] Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a gem. In *ICLR*, 2019. [6](#), [7](#), [8](#)
- [13] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc’Aurelio Ranzato. On tiny episodic memories in continual learning. In *ICML Workshops*, 2019. [2](#), [6](#), [7](#), [8](#)
- [14] Hongge Chen, Si Si, Yang Li, Ciprian Chelba, Sanjiv Kumar, Duane Boning, and Cho-Jui Hsieh. Multi-stage influence function. In *NeurIPS*, pages 12732–12742, 2020. [2](#)
- [15] Zhiyuan Chen and Bing Liu. *Lifelong Machine Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2018. [1](#)
- [16] R Dennis Cook and Sanford Weisberg. *Residuals and influence in regression*. New York: Chapman and Hall, 1982. [3](#)
- [17] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *TPAMI*, 44(7):3366–3385, 2021. [1](#)
- [18] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. [6](#)
- [19] Han Guo, Nazneen Rajani, Peter Hase, Mohit Bansal, and Caiming Xiong. FastIF: Scalable influence functions for efficient model interpretation and debugging. In *EMNLP*, pages 10333–10350, 2021. [2](#)
- [20] Raia Hadsell, Dushyant Rao, Andrei A Rusu, and Razvan Pascanu. Embracing change: Continual learning in deep neural networks. *Trends in Cognitive Sciences*, 24(12):1028–1040, 2020. [2](#)
- [21] Frank R Hampel. The influence curve and its role in robust estimation. *Journal of the American Statistical Association*, 69(346):383–393, 1974. [2](#)
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [6](#)
- [23] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: convergence and generalization in neural networks. In *NeurIPS*, pages 8580–8589, 2018. [2](#), [5](#)
- [24] Xisen Jin, Arka Sadhu, Junyi Du, and Xiang Ren. Gradient-based editing of memory examples for online task-free continual learning. In *NeurIPS*, pages 29193–29205, 2021. [6](#), [7](#)
- [25] Maurice G Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938. [8](#)
- [26] Krishnateja Killamsetty, S Durga, Ganesh Ramakrishnan, Abir De, and Rishabh Iyer. Grad-match: Gradient matching based data subset selection for efficient deep model training. In *ICML*, pages 5464–5474, 2021. [5](#)
- [27] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, 2017. [2](#)
- [28] Pang Wei Koh, Kai-Siang Ang, Hubert HK Teo, and Percy Liang. On the accuracy of influence functions for measuring group effects. In *NeurIPS*, pages 5254–5264, 2019. [2](#)
- [29] Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *ICML*, pages 1885–1894, 2017. [2](#), [3](#), [5](#), [8](#)
- [30] Shuming Kong, Yanyan Shen, and Linpeng Huang. Resolving training biases via influence-based data relabeling. In *ICLR*, 2022. [2](#)
- [31] Alex Krizhevsky. Learning multiple layers of features from tiny images. *Master’s thesis, University of Toronto*, 2009. [6](#)
- [32] Zhizhong Li and Derek Hoiem. Learning without forgetting. *TPAMI*, 40(12):2935–2947, 2017. [2](#)

- [33] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. In *NeurIPS*, page 6470–6479, 2017. [1](#), [2](#), [6](#), [7](#)
- [34] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of Learning and Motivation*, 24:109–165, 1989. [2](#)
- [35] Roman Novak, Lechao Xiao, Jiri Hron, Jaehoon Lee, Alexander A Alemi, Jascha Sohl-Dickstein, and Samuel S Schoenholz. Neural tangents: Fast and easy infinite neural networks in python. In *ICLR*, 2020. [6](#)
- [36] Seulki Park, Jongin Lim, Younghan Jeon, and Jin Young Choi. Influence-balanced loss for imbalanced visual classification. In *ICCV*, pages 735–744, 2021. [2](#)
- [37] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. PyTorch: An imperative style, high-performance deep learning library. In *NeurIPS*, pages 8026–8037, 2019. [6](#)
- [38] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. GDumb: A simple approach that questions our progress in continual learning. In *ECCV*, pages 524–540, 2020. [2](#), [6](#), [7](#)
- [39] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. iCaRL: Incremental classifier and representation learning. In *CVPR*, pages 2001–2010, 2017. [1](#), [2](#)
- [40] Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. Learning to learn without forgetting by maximizing transfer and minimizing interference. In *ICLR*, 2019. [2](#)
- [41] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016. [2](#)
- [42] Andrea Schioppa, Polina Zablotskaia, David Vilar, and Artem Sokolov. Scaling up influence functions. In *AAAI*, pages 8179–8186, 2022. [2](#)
- [43] Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *ICML*, pages 4548–4557, 2018. [2](#)
- [44] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *NeurIPS*, page 2994–3003, 2017. [2](#)
- [45] Qing Sun, Fan Lyu, Fanhua Shang, Wei Feng, and Liang Wan. Exploring example influence in continual learning. In *NeurIPS*, 2022. [2](#), [5](#), [6](#), [7](#), [8](#)
- [46] Daniel Ting and Eric Brochu. Optimal subsampling with influence functions. In *NeurIPS*, pages 3654–3663, 2018. [2](#)
- [47] Rishabh Tiwari, Krishnateja Killamsetty, Rishabh Iyer, and Pradeep Shenoy. GCR: Gradient coreset based replay buffer selection for continual learning. In *CVPR*, pages 99–108, 2022. [1](#), [2](#)
- [48] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *NeurIPS*, pages 3637–3645, 2016. [6](#)
- [49] Jeffrey S Vitter. Random sampling with a reservoir. *ACM Transactions on Mathematical Software*, 11(1):37–57, 1985. [2](#)
- [50] Tianyang Wang, Jun Huan, and Bo Li. Data dropout: Optimizing training data for convolutional neural networks. In *IEEE International Conference on Tools with Artificial Intelligence*, pages 39–46, 2018. [3](#)
- [51] Zifeng Wang, Hong Zhu, Zhenhua Dong, Xiuqiang He, and Shao-Lun Huang. Less is better: Unweighted data subsampling via influence function. In *AAAI*, pages 6340–6347, 2020. [2](#), [3](#)
- [52] Jaehong Yoon, Divyam Madaan, Eunho Yang, and Sung Ju Hwang. Online coreset selection for rehearsal-based continual learning. In *ICLR*, 2022. [1](#), [2](#), [3](#), [5](#)
- [53] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. In *ICLR*, 2018. [2](#)
- [54] Guanxiong Zeng, Yang Chen, Bo Cui, and Shan Yu. Continual learning of context-dependent processing in neural networks. *Nature Machine Intelligence*, 1(8):364–372, 2019. [2](#)
- [55] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *ICML*, pages 3987–3995, 2017. [6](#)
- [56] Rui Zhang and Shihua Zhang. Rethinking influence functions of neural networks in the over-parameterized regime. In *AAAI*, pages 9082–9090, 2022. [2](#)
- [57] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. Dataset condensation with gradient matching. In *ICLR*, 2021. [5](#)
- [58] Fan Zhou and Chengtai Cao. Overcoming catastrophic forgetting in graph neural networks with experience replay. In *AAAI*, pages 4714–4722, 2021. [2](#), [3](#)