

# Beyond Sharpness: A Flatness Decomposition Framework for Efficient Continual Learning

Yanan Chen<sup>1,2</sup>, Tieliang Gong<sup>1,2\*</sup>, Yunjiao Zhang<sup>4</sup>, Wen Wen<sup>1,3</sup>

<sup>1</sup>School of Computer Science and Technology, Xi'an Jiaotong University

<sup>2</sup>Ministry of Education Key Laboratory of Intelligent Networks and Network Security, Xi'an Jiaotong University

<sup>3</sup>Shaanxi Province Key Laboratory of Big Data Knowledge Engineering, Xi'an Jiaotong University

<sup>4</sup>China Telecom

{ynchen894, adidasgtl}@gmail.com

## Abstract

Continual Learning (CL) aims to enable models to sequentially learn multiple tasks without forgetting previous knowledge. Recent studies have shown that optimizing towards flatter loss minima can improve model generalization. However, existing sharpness-aware methods for CL suffer from two key limitations: (1) they treat sharpness regularization as a unified signal without distinguishing the contributions of its components, and (2) they introduce substantial computational overhead that impedes practical deployment. To address these challenges, we propose FLAD, a novel optimization framework that decomposes sharpness-aware perturbations into gradient-aligned and stochastic-noise components, and show that retaining only the noise component promotes generalization. We further introduce a lightweight scheduling scheme that enables FLAD to maintain significant performance gains even under constrained training time. FLAD can be seamlessly integrated into various CL paradigms and consistently outperforms standard and sharpness-aware optimizers in diverse experimental settings, demonstrating its effectiveness and practicality in CL.

## Introduction

Continual Learning (CL) aims to enable models to learn from a non-stationary stream of data or tasks while retaining previously acquired knowledge (Hadsell et al. 2020; De Lange et al. 2022). This setting poses a fundamental challenge known as catastrophic forgetting, a phenomenon where the learning of new tasks causes abrupt degradation in performance on previously learned ones (Parisi et al. 2019). Recent effort has been devoted to addressing this challenge by designing diverse paradigms including replay-based methods (Rebuffi et al. 2017; Lin et al. 2023), regularization-based methods (Lin et al. 2023; Sun et al. 2023) and architecture-based methods (Hu et al. 2023; Zhou et al. 2023b).

While these paradigms have proven effective, they often impose substantial storage requirements for raw data or model parameters, raising scalability concerns as task complexity increases (Zhou et al. 2023b; Wang et al. 2024). An alternative research direction focuses on enhancing the intrinsic generalization capacity of models to address these

limitations. This has motivated exploration of sharpness-aware minimization methods, which seek flat minima that are widely believed to yield superior generalization and robustness under distributional shifts (Keskar et al. 2017; Zhang et al. 2023). Notable examples include Sharpness-Aware Minimization (SAM) (Foret et al. 2021), which leverages zeroth-order sharpness to seek minima that are locally robust to parameter perturbations, and Gradient norm Aware Minimization (GAM) (Zhang et al. 2023), which exploits first-order sharpness to improve feature reuse. These approaches have shown considerable promise in transfer and few-shot learning, and have recently been adapted to CL with specialized designs (Tran Tung et al. 2023; Chen et al. 2024; Bian et al. 2024).

Despite their empirical success, existing flatness-based methods in CL suffer from two fundamental limitations that constrain their practical applicability. First, these approaches adopt an overly simplistic treatment of sharpness regularization, conceptualizing it as a monolithic objective rather than recognizing the heterogeneous nature of loss landscape curvature. Specifically, they fail to disentangle the distinct geometric properties and learning implications of different curvature components, particularly the critical distinction between gradient-aligned and gradient-orthogonal directions in parameter space. Second, they require substantial architectural modifications to standard optimization pipelines, introducing prohibitive computational overhead that severely limits their scalability and real-world deployment. These approaches typically require either double-gradient computations through higher-order differentiation or multiple forward-backward passes per optimization step, which significantly increases the training time.

In this paper, we propose FLAD, a novel Flatness Decomposition framework for continual learning that addresses the aforementioned limitations. Specifically, FLAD decomposes perturbation directions in sharpness-aware minimization by isolating and penalizing only the stochastic-noise component, thereby preventing counterproductive penalties on directions essential for optimization. Our comprehensive empirical analysis demonstrates that this targeted approach enables models to escape sharp valleys more effectively, converging to flatter minima with enhanced generalization properties. Furthermore, recognizing the computational inefficiencies of existing methods, particularly under time-

\*Corresponding author.

constrained training scenarios, we introduce a lightweight optimization scheme that maximizes learning efficacy without incurring additional computational overhead.

In summary, our main contributions are as follows:

- We propose a CL framework that leverages a decomposed flatness signal to escape sharp minima and preserve past knowledge with negligible computational overhead.
- We provide a novel interpretation of sharpness-aware optimization by decomposing perturbation directions into gradient-aligned and stochastic-noise components, and demonstrate the critical role of noise-aligned perturbations in improving generalization for CL.
- Extensive experimental results demonstrate that our framework outperforms state-of-the-art competitors across multiple CL benchmarks. We further show that even partial application of our strategy yields significant gains, enabling substantial reductions in training time without sacrificing effectiveness.

## Related Works

**Continual learning.** CL methods are broadly categorized into three groups (Wang et al. 2024; van de Ven, Tuytelaars, and Tolias 2022; Parisi et al. 2019): replay-based, regularization-based, and architecture-based methods. Replay-based methods mitigate forgetting by extending data space with selected exemplars from old tasks (Rebuffi et al. 2017; Rolnick et al. 2019; Sun, Mu, and Hua 2023). Regularization-based methods introduce explicit constraints into the loss function to balance performance between new and old tasks (Lin et al. 2023; Saha, Garg, and Roy 2021; Sun et al. 2023), which typically act on the parameter space or the gradients to preserve past knowledge during optimization. Architecture-based methods allocate task-specific sub-networks or expand the model, thereby minimizing interference across tasks (Lu et al. 2024; Zhou et al. 2023b). Alternatively, recent efforts explicitly modify the training dynamics to preserve generalization across tasks (Wu et al. 2024; Lee et al. 2024). One direction is to modify the gradients of different tasks to overcome forgetting (Chaudhry et al. 2019; Lopez-Paz and Ranzato 2017). Another emerging line of research investigates CL from the perspective of the loss landscape, aiming to improve generalization by steering training towards flatter minima (Shi et al. 2021; Deng et al. 2021). Several works have empirically shown the benefits of SAM in mitigating catastrophic forgetting and preserving performance on previous tasks (Mehta et al. 2023; Chen et al. 2024), including DFGP (Yang et al. 2023), FS-DGPM (Deng et al. 2021), and SAM-CL series (Tran Tung et al. 2023). Recent work C-Flat combines zeroth- and first-order sharpness to achieve globally flatter solutions and improve CL performance (Bian et al. 2024). While promising, such sharpness-aware methods remain underexplored in CL.

**Sharpness-aware minimization.** Sharpness-aware minimization was originally proposed to improve generalization by guiding optimization toward flat minima in standard supervised learning (Foret et al. 2021; He, Huang, and Yuan

2019). This has motivated increasing interest in understanding the theoretical foundations of its generalization behavior, especially those of SAM, which is based on zeroth order sharpness. These effort include PAC-Bayes generalization analyses (Foret et al. 2021), studies of implicit bias and optimization dynamics (Andriushchenko and Flammarion 2022; Chen et al. 2024, 2023), and recent investigations dissecting the core components responsible for empirical effectiveness in SAM (Li et al. 2024). These empirical and theoretical developments motivate our work, propose a principled and efficient framework for CL that leverages the core mechanism of sharpness-aware optimization to improve generalization.

## Method

In this section, we propose a framework for continual learning that explicitly incorporates flatness-promoting perturbations. By leveraging both zeroth- and first-order sharpness information, the method aims to mitigate forgetting while enabling effective knowledge transfer across tasks.

**Loss landscape sharpness.** Let  $\mathcal{D}$  denote the training distribution on  $\mathcal{X} \times \mathcal{Y}$  and  $S = \{(x_i, y_i)\}_{i=1}^n$  denote the training dataset with  $n$  data-points drawn independently from  $\mathcal{D}$ . Let a family of models parameterized by  $w \in \mathcal{W} \in \mathbb{R}^d$  and a per-data-point loss function  $\ell(\cdot) : \mathcal{W} \times \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}^+$ . The empirical training loss is defined as

$$L(w) = \frac{1}{n} \sum_{i=1}^n \ell_i(w). \quad (1)$$

Zeroth-order sharpness is the worst-case loss within a defined neighborhood, defined over training set  $S$  as (Foret et al. 2021):

$$R_\rho^0(w) = \max_{\|\delta\|_2 < \rho} L(w + \delta) - L(w), \quad (2)$$

where  $\rho > 0$  denotes the neighborhood radius.

First-order sharpness is a measurement of the maximal neighborhood gradient norm, reflecting landscape curvature and defined over training set  $S$  as (Zhang et al. 2023):

$$R_\rho^1(w) = \max_{\|\delta\|_2 < \rho} \|\nabla L(w + \delta)\|, \quad (3)$$

where  $\nabla L(w)$  is the gradient of  $L(w)$  with respect to  $w$ .

**Optimization with sharpness.** To efficiently optimize the zeroth-order flatness in Eq. 2, it is approximated via first-order expansion and compute the adversarial perturbation  $\delta$  as follows (Foret et al. 2021):

$$\delta = \rho \cdot \frac{\nabla L(w)}{\|\nabla L(w)\|_2}. \quad (4)$$

Subsequently, one can compute the gradient at the perturbed point  $w + \delta$ , and then use the updating step of the base optimizer such as SGD to update

$$w_{t+1} = w_t - \lambda \nabla L(w)|_{w_t + \delta_t}. \quad (5)$$

The zeroth-order sharpness minimization addresses the loss landscape curvature through the gradient norm  $\|\nabla L_B(w)\|$  in Eq. 4, which serves as a proxy for local

sharpness. To enhance generalization, we aim to remove the shared, dominated structure from this direction and retain only its batch-specific stochastic variation (Li et al. 2024). Denote the batch gradient by  $\nabla L_B(w)$  as  $\hat{g}_B$ , and the full-batch gradient  $\nabla L(w)$  by  $g$ , we have

$$\hat{g}_B = \text{Proj}_g(\hat{g}_B) + \text{Proj}_g^\perp(\hat{g}_B), \quad (6)$$

where  $\text{Proj}_g(\hat{g}_B)$  is the projection of  $\hat{g}_B$  onto  $g$ ,  $\text{Proj}_g^\perp(\hat{g}_B)$  is the orthogonal component. In practice, we approximate the orthogonal component as

$$\text{Proj}_g^\perp \hat{g}_B = \hat{g}_B - \sigma m_t, \quad (7)$$

where  $\sigma = \cos(g, \hat{g}_B)$  denotes the cosine similarity between  $g$  and  $\hat{g}_B$ . To improve training efficiency,  $\sigma$  is fixed as a constant during optimization.

Note that computing  $g$  over the full dataset in each iteration is computationally prohibitive. We therefore maintain an exponential moving average (EMA) of the per-step gradient norm to approximate the full-batch gradient:

$$m_t = \lambda_0 m_{t-1} + (1 - \lambda_0) \hat{g}_{B_t}, \quad (8)$$

where  $\lambda_0 \in (0, 1)$  is a hyper-parameter. It is proved that EMA provides a reliable approximation of the full gradient (Li et al. 2024). We can use  $\text{Proj}_g^\perp \hat{g}_B$  Eq. 7 to substitute  $g$  in Eq. 4 to get  $\delta$  then update in Eq. 5.

Similarly, the minimization of the first-order sharpness in Eq.3 is approximated and compute the adversarial perturbation, follow by the updating step

$$\delta = \rho \frac{\nabla \|\nabla L(w)\|}{\|\nabla \|\nabla L(w)\|_2\|_2}, \quad (9)$$

$$w_{t+1} = w_t - \lambda \nabla \|\nabla L(w)\|_2|_{w_t + \delta_t},$$

where  $\nabla \|\nabla L(w)\|$  denotes the gradient of  $\|\nabla L(w)\|$  with respect to  $w$ . Note that

$$\nabla \|\nabla L(w)\| = \nabla^2 L(w) \cdot \frac{\nabla L(w)}{\|\nabla L(w)\|}. \quad (10)$$

While the zeroth-order method targets the norm of the gradient, the first-order variant considers the sharpness of the gradient itself, that is, the curvature of the function  $\|L(w)\|$ . As before, we use  $\hat{g}_B$  and  $g$  to denote  $\nabla L_B(w)$  and  $\nabla L(w)$ . To extract meaningful stochastic directions, we decompose  $\nabla \|\hat{g}_B\|$  into components aligned with and orthogonal to the EMA-estimated full-batch direction:

$$\nabla \|\hat{g}_B\| = \text{Proj}_{\nabla \|g\|} \nabla \|\hat{g}_B\| + \text{Proj}_{\nabla \|g\|}^\perp \nabla \|\hat{g}_B\|, \quad (11)$$

$$\text{Proj}_{\nabla \|g\|}^\perp \nabla \|\hat{g}_B\| = \nabla \|\hat{g}_B\| - \sigma \nabla \|g\|,$$

where  $\sigma$  is a constant during training,  $n_t$  is the estimation of the gradient sharpness direction:

$$n_t = \lambda_1 n_{t-1} + (1 - \lambda_1) \nabla \|\hat{g}_{B_t}\|, \quad (12)$$

where  $\lambda_1 \in (0, 1)$  is a hyperparameter.

We unify both zeroth- and first-order sharpness minimization in a single optimization framework, where the perturbation directions in each are respectively decomposed into stochastic-noise components. The optimization is detailed in

---

#### Algorithm 1: FLAD Algorithm

---

**Input:** Training phase  $T$ , training data  $S^T$ , model  $f^{T-1}$  with parameter  $w^{T-1}$  from last phase if  $T > 1$ , batchsize  $b$ , oracle loss function  $L$ , learning rate  $\eta > 0$ , neighborhood size  $\rho$ , trade-off coefficient  $\gamma$ , hyper-parameter  $\lambda_0, \lambda_1, \sigma$ , small constant  $c$ .

**Output:** Model trained at the current time  $T$ .

- 1: **Initialize:** if  $T = 1$ : Randomly Initialize parameter  $w^{T=1}, t \leftarrow 0, m_{-1} = n_{-1} = 0$
  - 2: **while**  $w^T$  not converge, **do**
  - 3:   Sample batch  $B^T$  of  $b$  from  $S^T$
  - 4:   Compute batch's loss gradient  $\hat{g}_{B_t} = \nabla L_{B^T}(w)$
  - 5:   Compute  $m_t$  in Eq. 8
  - 6:   Compute perturbation  $\delta_0 = \rho \frac{\hat{g}_{B_t} - \sigma m_t}{\|\hat{g}_{B_t} - \sigma m_t\| + c}$
  - 7:   Approximate  $g_0 = \nabla L_{B^T}(w + \delta_0)$
  - 8:   Compute Hessian vector product in Eq. 10
  - 9:   Compute  $n_t$  in Eq. 12
  - 10:   Compute  $\delta_1 = \rho \frac{\nabla \|\hat{g}_{B_t}\| - \sigma n_t}{\|\nabla \|\hat{g}_{B_t}\| - \sigma n_t\| + c}$
  - 11:   Approximate  $g_1 = \nabla^2 L_{B^T}(w + \delta_1) \cdot \frac{\nabla L_{B^T}(w + \delta_1)}{\|\nabla L_{B^T}(w + \delta_1)\| + c}$
  - 12:   Update: model parameter:  $w = w - \eta(g_0 + \gamma g_1)$
  - 13: **end while**
  - 14: **return**  $w^T$  and  $f^T$
- 

Algorithm 1. By using Hessian-vector product in Eq. 10, we significantly reduce both time and memory complexity, the overall strategy requires 2 forward and 4 backward passes per iteration. The framework introduces minimal additional computational overhead and can be seamlessly integrated into standard optimizers.

**Continual learning.** We now integrate our noise-aware optimizer into a general continual learning framework, focusing on the class-incremental learning (CIL) setting, which is arguably the most challenging CL scenario, where the model must continuously acquire knowledge from sequentially presented, class-agnostic data without task identity at inference.

Our approach is broadly compatible with the three major paradigms in CIL: replay-based, regularization-based, and expansion-based methods. Replay-based and regularization-based methods can directly benefit from our optimizer by plugging the original task loss in the curvature-regularized objective:

$$L(w^T) = L^{R_\rho^0}(w^T) + \gamma \cdot L^{R_\rho^1}(w^T), \quad (13)$$

where the zeroth- and first-order terms encourage flat minima aligned with generalization. These methods typically operate on data streams that includes exemplars, and our optimizer fits naturally into this training protocol using FLAD Algorithm. Expansion-based methods extend the model by allocating task-specific modules for new classes. In this case, the same objective in Eq. 13 can be applied to the added components during the learning phase, while shared components remain frozen. Final inference is performed after a post-processing step, such as classifier calibration or gating.




Figure 1: Empirical analysis of perturbation components in zeroth- and first-order sharpness minimization. (a) Average accuracy across four training settings using different perturbation directions. (b)  $\text{Tr}(\mathbf{H}\Sigma)$  during training for stochastic-noise and original variants. (c) 2D loss landscape visualizations around final model parameters and training trajectory.

This formulation demonstrates the flexibility of our method in adapting to various CL strategies, while preserving its benefits centered on flatness-induced generalization.

### Empirical Analysis

To validate the effectiveness behind our unified framework, we conduct an empirical investigation into the underlying mechanisms of zeroth- and first-order sharpness minimization by 1) identifying the effective components in the adversarial perturbation that contributes to improved generalization, and 2) quantitatively analyze the gradient-aligned and stochastic-noise components of the perturbations, providing deeper insights into the role of stochasticity in optimization.

**Effect of adversarial perturbation.** Previous work (Andriushchenko and Flammarion 2022; Li et al. 2024) has demonstrated that the direction of adversarial perturbation plays a critical role in zeroth-order sharpness minimization. In particular, decomposing the perturbation into a gradient-aligned component and a stochastic-noise component has been empirically shown to improve generalization.

We extend this analysis to the first-order setting by examining how different perturbation directions affect the performance. Specifically, we consider the following variants: GAM, Pre: perturbations computed using the previous minibatch, Random: perturbation directions are randomly sampled, Full: perturbation direction aligned with the full gradient component in Eq. 6, Noise: perturbation direction restricted to the stochastic-noise component in Eq. 6, corresponding to the first-order part of FLAD. We replace the perturbation in Eq. 9 with each component separately and evaluate their effects.

We conduct experiments using four training settings, which use ResNet-18 and Resnet-32 on CIFAR-10/100 with standard data augmentation and batchsize 128. As shown in Fig. 1a, GAM consistently improves over SGD across all settings. Importantly, Noise outperforms all other variants, indicating that the stochastic-noise component alone captures the most beneficial aspect of perturbation. Conversely, Full degrades performance, likely due to introducing conflict with learning dynamics.

This suggests that suppressing sample-induced gradient noise improves generalization without interfering with the main optimization trajectory. In contrast, penalizing the full gradient direction introduces conflict with learning dynamics and may degrade convergence. Together, these findings provide strong evidence for the effectiveness of our orthogonal decomposition strategy.

**Analysis of noise components.** To further verify the effectiveness of the stochastic-noise component, we analyze its interaction with the loss curvature using the metric  $\text{Tr}(\mathbf{H}\Sigma)$ , which captures the alignment between the Hessian  $\mathbf{H}$  and the gradient noise covariance  $\Sigma$  (Zhu et al. 2019). A higher value of it suggests a stronger tendency to escape narrow minima and explore flatter regions. We estimate  $\text{Tr}(\mathbf{H}\Sigma)$  using the top Hessian eigenpairs and project update gradients along those directions.

As shown in Fig. 1b, in early training, both the noise variants, which corresponding to the zeroth-order (FLAD-0th) and first-order (FLAD-1st) parts of FLAD, almost generally exhibits higher  $\text{Tr}(\mathbf{H}\Sigma)$  than their standard counterparts, with stronger fluctuations. This implies that the perturbation is more anisotropic and better aligned with high-curvature directions, facilitating escape from sharp minima. In the later training phase, the trend reverses, which indicates a reduction in curvature-aligned noise, promoting stable convergence. Consistently, the final loss landscape in Fig. 1c reveals that solutions obtained by noise variants lie in significantly flatter regions, confirming their improved generalization. These results reinforce that the stochastic-noise component enhances generalization, offering a form of implicit regularization shared by both zeroth- and first-order sharpness minimization.

### Convergence Analysis

We analyze the convergence properties of the our algorithm under non-convex setting in Theorem 1.

**Theorem 1.** Assume the loss function is twice differentiable, bounded by  $M$ , and obeys the triangle inequality. Both the loss function and its second-order gradient are  $\beta$ -Lipschitz smooth. FLAD converges in all tasks with learn-

Methods	CIFAR-10		CIFAR-100				Tiny-ImageNet		Average
	N=5		N=10		N=5		N=8		Return
	AAA	Acc	AAA	Acc	AAA	Acc	AAA	Acc	AAA
Replay	61.84 ±6.66	41.68 ±3.61	49.89 ±1.04	29.07 ±1.30	53.94 ±1.72	33.62 ±0.36	39.51 ±0.47	19.09 ±1.17	+2.18%
w/SAM	60.91 ±5.42	42.57 ±4.32	51.16 ±1.13	31.31 ±2.12	54.91 ±1.59	<b>34.97</b> ±0.49	41.19 ±0.27	21.22 ±0.60	
w/GAM	<u>61.92</u> ±5.52	<u>43.07</u> ±6.82	50.26 ±0.93	30.30 ±1.39	54.49 ±1.76	<b>34.97</b> ±1.39	40.10 ±0.51	20.23 ±0.39	
w/C-Flat	60.46 ±6.62	42.32 ±2.85	51.15 ±1.27	31.07 ±1.86	54.85 ±1.65	<u>34.40</u> ±0.38	<u>41.95</u> ±0.10	<u>22.29</u> ±0.20	
w/FLAD	<b>62.60</b> ±4.85	<b>43.13</b> ±4.32	<b>51.81</b> ±0.95	<b>31.74</b> ±1.69	<b>55.19</b> ±1.59	<u>34.40</u> ±0.10	<b>42.41</b> ±0.18	<b>22.91</b> ±0.75	
iCaRL	66.54 ±4.42	50.63 ±3.17	51.11 ±1.23	30.05 ±1.83	56.84 ±1.72	35.91 ±0.76	41.12 ±0.49	20.38 ±0.53	
w/SAM	66.85 ±4.57	50.93 ±4.86	<u>51.45</u> ±0.96	30.14 ±1.98	57.33 ±1.87	36.95 ±0.62	42.35 ±0.27	21.83 ±0.40	
w/GAM	66.46 ±4.82	49.86 ±3.94	51.19 ±1.06	30.05 ±2.20	56.84 ±1.52	36.07 ±0.59	42.00 ±0.48	21.23 ±0.68	
w/C-Flat	<u>67.89</u> ±4.10	<u>51.62</u> ±6.38	51.09 ±1.21	30.05 ±1.83	57.21 ±1.90	36.75 ±0.10	<u>42.81</u> ±0.41	<u>22.17</u> ±0.35	
w/FLAD	<b>68.32</b> ±4.54	<b>52.53</b> ±6.02	<b>51.75</b> ±1.21	<b>30.36</b> ±2.06	<b>57.59</b> ±1.60	<b>37.13</b> ±0.97	<b>43.29</b> ±0.57	<b>23.04</b> ±0.86	
WA	72.03 ±3.58	61.95 ±4.33	62.56 ±1.26	47.29 ±0.25	66.89 ±2.00	53.59 ±0.21	51.92 ±0.38	36.40 ±0.23	+1.90%
w/SAM	<u>72.27</u> ±3.94	61.26 ±5.77	63.04 ±1.45	47.35 ±0.98	67.36 ±2.17	54.55 ±0.73	52.39 ±0.07	<u>37.81</u> ±0.33	
w/GAM	69.50 ±4.70	<u>62.18</u> ±2.78	<u>63.13</u> ±1.16	47.79 ±0.69	<u>67.38</u> ±2.10	54.51 ±0.43	50.95 ±1.15	36.56 ±0.46	
w/C-Flat	71.59 ±1.95	61.45 ±4.26	62.24 ±1.57	47.04 ±1.02	67.24 ±1.98	<u>54.66</u> ±0.34	<u>52.71</u> ±0.84	35.61 ±0.08	
w/FLAD	<b>72.43</b> ±1.80	<b>62.46</b> ±4.26	<b>63.73</b> ±1.19	<b>48.35</b> ±0.74	<b>67.63</b> ±2.02	<b>54.69</b> ±0.47	<b>53.53</b> ±0.18	<b>39.00</b> ±0.15	
FOSTER	65.15 ±3.86	56.43 ±4.55	52.71 ±1.25	<b>39.71</b> ±1.68	56.32 ±1.74	37.91 ±0.55	48.32 ±0.66	38.10 ±0.55	
w/SAM	<u>74.80</u> ±3.26	68.52 ±2.78	52.40 ±1.89	37.81 ±1.80	<u>56.47</u> ±2.30	37.92 ±0.14	48.49 ±1.15	<u>38.42</u> ±1.79	
w/GAM	63.86 ±5.31	54.04 ±3.30	<u>52.93</u> ±1.09	39.41 ±1.95	56.19 ±1.75	38.09 ±0.52	48.49 ±1.03	38.27 ±1.00	
w/C-Flat	73.68 ±3.62	<u>68.81</u> ±4.65	52.05 ±1.84	38.71 ±1.85	56.35 ±2.00	37.81 ±0.66	<u>48.56</u> ±0.96	38.25 ±0.55	
w/FLAD	<b>75.08</b> ±2.17	<b>68.98</b> ±4.41	<b>53.00</b> ±2.01	39.13 ±1.80	<b>56.51</b> ±2.01	<b>38.82</b> ±0.85	<b>48.66</b> ±0.88	<b>38.54</b> ±1.15	
MEMO	<u>67.86</u> ±6.63	57.80 ±4.55	64.38 ±2.83	53.27 ±4.16	68.05 ±1.94	<b>61.58</b> ±8.40	56.53 ±0.55	44.29 ±0.18	+1.83%
w/SAM	65.99 ±6.51	54.96 ±2.96	<u>65.41</u> ±3.00	53.67 ±0.67	69.10 ±1.69	59.40 ±0.25	<u>57.53</u> ±0.66	<u>44.64</u> ±0.27	
w/GAM	66.92 ±6.69	55.34 ±5.56	64.60 ±2.72	53.22 ±0.82	68.36 ±1.85	57.78 ±1.23	57.22 ±0.84	44.43 ±0.27	
w/C-Flat	66.29 ±7.34	58.04 ±4.80	64.96 ±2.63	<b>54.07</b> ±0.74	<u>68.95</u> ±1.41	59.70 ±0.40	57.41 ±1.32	44.38 ±0.23	
w/FLAD	<b>68.51</b> ±6.07	<b>61.40</b> ±4.92	<b>65.47</b> ±2.75	<u>53.92</u> ±0.48	<b>69.30</b> ±1.50	<u>60.09</u> ±0.06	<b>57.91</b> ±0.95	<b>45.06</b> ±0.12	
PODNet	72.96 ±1.48	57.03 ±1.50	51.65 ±1.06	31.87 ±0.72	61.52 ±1.24	45.46 ±0.51	49.35 ±0.30	33.13 ±0.09	
w/SAM	74.13 ±1.51	57.17 ±3.09	52.55 ±1.20	33.03 ±0.89	61.99 ±1.28	45.97 ±0.23	50.25 ±0.64	33.90 ±0.31	
w/GAM	73.08 ±1.24	55.79 ±2.73	51.87 ±1.28	31.82 ±0.14	61.68 ±1.28	45.49 ±0.73	49.58 ±0.75	33.18 ±0.52	
w/C-Flat	<u>74.19</u> ±1.22	<u>57.20</u> ±2.72	<u>53.15</u> ±1.37	33.84 ±0.83	<u>62.21</u> ±1.14	<u>46.20</u> ±0.44	<u>50.31</u> ±0.69	<u>33.93</u> ±0.14	
w/FLAD	<b>74.64</b> ±1.58	<b>57.62</b> ±2.55	<b>53.74</b> ±1.26	<b>34.85</b> ±0.42	<b>62.47</b> ±1.19	<b>46.92</b> ±0.66	<b>50.70</b> ±0.45	<b>35.04</b> ±0.39	

Table 1: Average accuracy (%) across all phases using 6 advanced methods, which span three categories in CL: Memory-based methods, Regularization-based methods and Expansion-based methods (w/ and w/o four optimizer plugged in). The mean and standard deviation was estimated over 3 runs. The best results are in bold and the second-best results are underlined. Average Return in the last column represents the average boost of FLAD towards C-flat in each row.

ing rate  $\eta \leq 1/\beta$ , the perturbation radius  $\rho \leq 1/4\beta$ , and  $\eta_i^T = \eta/\sqrt{i}$ ,  $\rho_i^T = \rho/\sqrt{i}$  for each epoch  $i$  in any task  $T$ .

$$\frac{1}{n^T} \sum_{i=1}^{n^T} \mathbb{E}[\|\nabla L(w_i)\|^2] \leq \frac{C_1 + \log n^T}{8\sqrt{n^T}} + \frac{64\gamma^2\sqrt{n^T} - C_2}{\beta^2 n^T}, \quad (14)$$

where  $n^T$  is the total iteration numbers of task  $T$ ,  $C_1 = 32M(\beta - 1)$ ,  $C_2 = 32\gamma^2$  only depends on  $\gamma, M, \beta$ .

See the proof in Appendix A. For non-convex stochastic optimization, Theorem 1 shows that FLAD has the convergence rate  $\mathcal{O}(\log n^T / \sqrt{n^T})$ .

## Experiments

### Experimental Setup

**Datasets.** We evaluate the performance of our method on CIFAR-10, CIFAR-100 (Keskar et al. 2017) and Tiny-

ImageNet (Deng et al. 2009). Specifically, we split CIFAR-10 into 5 tasks (2 classes per task), CIFAR-100 into 5/10 tasks (20/10 classes per task), Tiny-ImageNet into 8 tasks (25 classes per task). We adopt two standard and widely used CL metrics: Average Accuracy (Acc), which measures the final performance by averaging classification accuracy over all  $N$  tasks at the end of training; and Average Anytime Accuracy (AAA), which averages the accuracy over all learned tasks after training on each new task.

**Baselines.** To evaluate the efficacy of our method, we plug it into 6 leading baselines spanning the three major CL category, Replay-based methods: Replay (Rolnick et al. 2019) and iCaRL (Rebuffi et al. 2017), which store raw exemplars from previous tasks, and PODNet (Douillard et al. 2020) is akin to iCaRL, incorporating knowledge distillation to preserve past knowledge; Regularization-based methods: WA (Zhao et al. 2020), which mitigates prediction bias

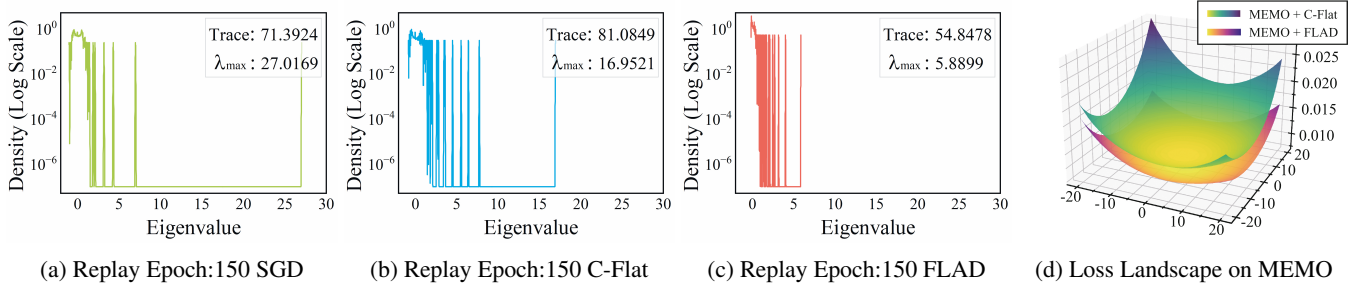


Figure 2: Analysis of generalization. In (a)(b)(c), we report Hessian eigenvalue distributions and the trace of SGD, C-Flat, and FLAD on replay  $N=5$ . In (d), we show the loss landscape around final model parameters on MEMO  $N=10$ .

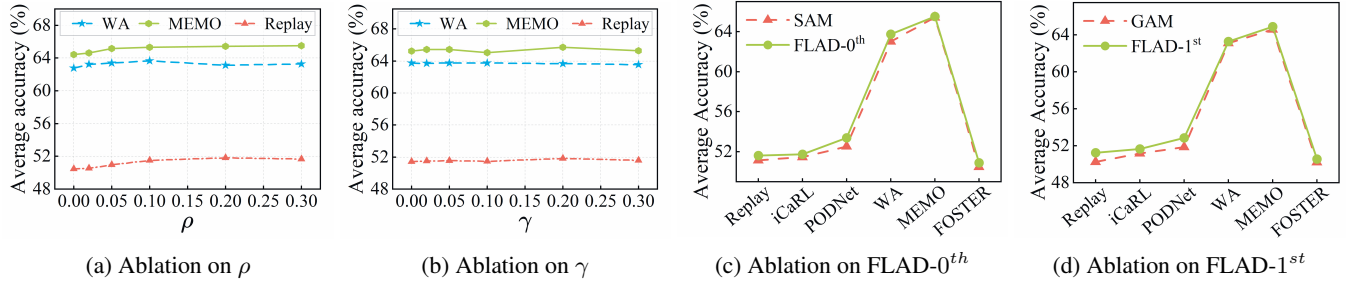


Figure 3: Ablation experiments. In (a)(b), we conduct ablation experiments on parameter  $\rho$  and  $\gamma$ . In (c)(d), we investigate the impact of our decomposition strategy across 6 CL methods.

via fairness-aware regularization; Expansion-based methods: FOSTER (Wang et al. 2022) and MEMO (Zhou et al. 2023b), which dynamically expand the network using modular substructures or parameter freezing.

**Network and training details.** All experiments are conducted on RTX 4090Ti GPUs with 96GB of RAM. For each dataset, we evaluate all methods using the same network architecture following repo (Zhou et al. 2023a, 2024). Details are provided in Appendix A. All methods are trained using vanilla SGD (Zinkevich 2003), into which FLAD and other methods are seamlessly plugged. For fair comparison, hyperparameters  $\rho, \gamma, \lambda_0$  and  $\lambda_1$  are initialized identically across tasks, while  $\lambda \in \{0.5, 0.7, 0.9\}$  is tuned based on validation performance.

### Performance in Continual Learning

In this experiment, we plug our method into 6 advanced methods that cover the full range of CL methods. Table 1 demonstrates that when applied across multiple benchmarks and continual settings, our method yields consistent and significant performance improvements. These enhancements manifest 1) Universality: Our method benefits all baselines, indicating its broad compatibility and minimal assumptions on the underlying CL mechanism, 2) Generalization across datasets: On each dataset, our approach consistently achieves higher final and anytime accuracies, suggesting robustness to scale and domain shift, 3) Scenario resilience: Whether under coarse- or fine-grained task splits, the gains remain stable, highlighting its adaptability across different incremental conditions. These results demonstrate

that our framework serves as a versatile and effective augmentation to existing CL frameworks, offering generalization gains with negligible integration cost.

### Analysis of Generalization

Hessian eigenvalues are widely used to quantify the sharpness of loss minima, with smaller maximum eigenvalues indicating flatter optima. Since flatness-aware minimization aim to reduce curvature by penalizing sharp directions, we assess whether our method indeed leads to flatter solutions by analyzing the Hessian spectrum during training.

We report the Hessian eigenvalue distributions and trace throughout CL in Fig. 2. Compared to vanilla SGD and C-Flat, our method yields a pronounced reduction in both the top eigenvalue and the trace, significantly suppresses high-curvature directions, indicating the discovery of flatter regions in the loss landscape. These results empirically validate that noise-aware perturbations promote flatness.

To provide a more intuitive understanding, we visualize the loss surface by PyHessian (Yao et al. 2020) in Fig. 2d. The landscape of MEMO becomes noticeably smoother and wider when integrated with our framework. These visual and spectral results jointly support the effectiveness of our framework in guiding the model toward flatter minima, which contributes to improved continual generalization.

### Ablation Experiments

We conduct ablation experiments to evaluate the effect of key hyperparameters and the impact of our proposed decomposition strategy. We first investigate the sensitivity of our

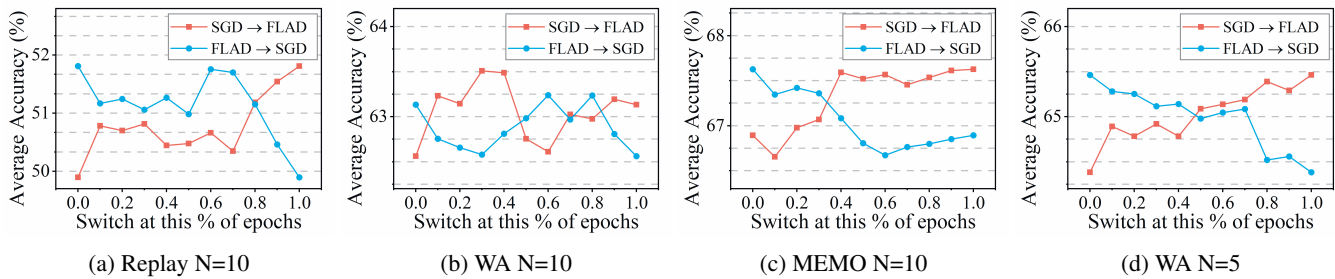


Figure 4: To further reduce overhead, we apply FLAD in a limited number of epochs within each task on different method and different settings, the horizontal axis is the transition point, before and after which we use different optimizers (SGD or FLAD).

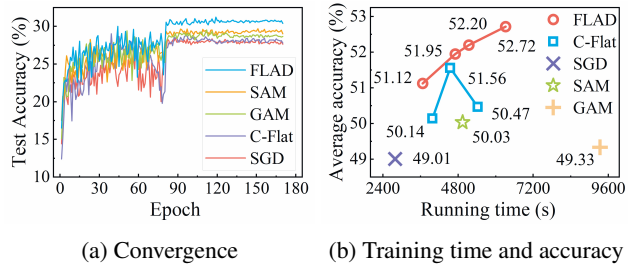


Figure 5: Convergence and computation overhead. In (a), we compare test accuracy of different optimizers during training process. In (b), we compare accuracy and training time of different optimizers.

method to the weight of first-order term  $\gamma$  in Eq. 13 and the perturbation radius  $\rho$  as defined in Eq. 4 and Eq. 9,  $\gamma = 0$  reduces the method to its zeroth-order variant. As shown in Fig. 3a, our method consistently outperforms the vanilla optimizer across a wide range of  $\gamma$  values, confirming the utility of incorporating first-order curvature information. Similarly, Fig. 3b illustrates the influence of  $\rho$ , it demonstrate that models trained with  $\rho > 0$  uniformly outperform those without gradient ascent.

We further assess the contribution of our decomposition by comparing the performance of the decomposed and original variants across all 6 CL methods. In each case, we replace the gradient term with its decomposed variant, which isolates stochastic-noise directions. As summarized in Fig. 3c and Fig. 3d, the decomposed version consistently outperforms the original across methods. This highlights the importance of separating stochastic-noise from global trends, and affirms the value of our decomposition design in promoting generalization under continual learning.

### Convergence and Computation Overhead

To evaluate the efficiency of our method, we analyze both convergence speed and computational cost with the Replay baseline on CIFAR-100. As shown in Fig. 5a, our method achieves the fastest convergence among all optimizers and reaches the highest final accuracy. This demonstrates that our method not only improves generalization but also accelerates training dynamics.

To further reduce overhead, we explore applying our

method only during a subset of epochs. Surprisingly, we find that even limited use of our method leads to substantial performance gains across all cases. As summarized in Fig. 4, applying our optimizer in just 10–20% of epochs already yields significant improvement over vanilla SGD, and in many cases, it achieves performance comparable to or even better than using the optimizer throughout training. As a consequence, applying our method for only 30% of the epochs reduces the computational overhead by at least 50% compared to full-method training, highlighting its practicality and scalability for efficient continual learning.

We compare runtime and accuracy trade-offs across different optimizers in Fig. 5b. The results show that 1) Using our method for only 10% of the epochs (red line) already improves performance over SGD while consuming slightly less time, 2) at 20% usage, our method outperforms GAM, SAM, and C-Flat with comparable or lower computational cost, 3) training with only 20 epochs using our method achieves higher accuracy than training with 50 epochs using C-Flat, and even 200 epochs using other optimizers.

These results confirm that our method offers a favorable balance between efficiency and effectiveness. Without introducing additional modules or increasing model complexity, it can be seamlessly integrated into existing CL methods, enabling faster convergence and better performance with fewer optimization steps.

## Conclusion

In this paper, we have presented FLAD, a flatness decomposition framework for continual learning that isolates the stochastic-noise component of sharpness-aware perturbations. Empirical analyses reveal that this stochastic-noise direction plays a key role in escaping sharp minima effectively and enhancing generalization. By combining zeroth- and first-order decompositions with a partial application strategy, FLAD achieves strong performance gains with minimal computational overhead. Extensive experiments demonstrate that FLAD is effective across a diverse set of continual learning paradigms, highlighting its plug-and-play capability and adaptability. Its ability to improve learning efficiency with fewer optimization steps highlights the potential of curvature-guided methods for scalable and effective continual learning.

## Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grants 62192781, 62576268, 62137002, the Key Research and Development Project in Shaanxi Province No. 2023GXLH-024, and the Project of China Knowledge Centre for Engineering Science and Technology.

## References

- Andriushchenko, M.; and Flammarion, N. 2022. Towards Understanding Sharpness-Aware Minimization. In *Proceedings of the 39th International Conference on Machine Learning*, 639–668. PMLR.
- Bian, A.; Li, W.; Yuan, H.; Yu, C.; Wang, M.; Zhao, Z.; Lu, A.; Ji, P.; and Feng, T. 2024. Make Continual Learning Stronger via C-Flat. In Globerson, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J.; and Zhang, C., eds., *Advances in Neural Information Processing Systems*, volume 37, 7608–7630.
- Chaudhry, A.; Ranzato, M.; Rohrbach, M.; and Elhoseiny, M. 2019. Efficient Lifelong Learning with A-GEM. In *International Conference on Learning Representations*.
- Chen, R.; Jing, X.-Y.; Wu, F.; and Chen, H. 2024. Sharpness-aware gradient guidance for few-shot class-incremental learning. *Knowledge-Based Systems*, 299: 112030.
- Chen, Z.; Zhang, J.; Kou, Y.; Chen, X.; Hsieh, C.-J.; and Gu, Q. 2023. Why Does Sharpness-Aware Minimization Generalize Better Than SGD? In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 72325–72376.
- De Lange, M.; Aljundi, R.; Masana, M.; Parisot, S.; Jia, X.; Leonardis, A.; Slabaugh, G.; and Tuytelaars, T. 2022. A Continual Learning Survey: Defying Forgetting in Classification Tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(7): 3366–3385.
- Deng, D.; Chen, G.; Hao, J.; Wang, Q.; and Heng, P.-A. 2021. Flattening Sharpness for Dynamic Gradient Projection Memory Benefits Continual Learning. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P. S.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 18710–18721.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255.
- Douillard, A.; Cord, M.; Ollion, C.; Robert, T.; and Valle, E. 2020. PODNet: Pooled Outputs Distillation for Small-Tasks Incremental Learning. In Vedaldi, A.; Bischof, H.; Brox, T.; and Frahm, J.-M., eds., *ECCV*, 86–102.
- Foret, P.; Kleiner, A.; Mobahi, H.; and Neyshabur, B. 2021. Sharpness-aware Minimization for Efficiently Improving Generalization. In *International Conference on Learning Representations*.
- Hadsell, R.; Rao, D.; Rusu, A. A.; and Pascanu, R. 2020. Embracing Change: Continual Learning in Deep Neural Networks. *Trends in Cognitive Sciences*, 24(12): 1028–1040.
- He, H.; Huang, G.; and Yuan, Y. 2019. Asymmetric Valleys: Beyond Sharp and Flat Local Minima. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; Alché-Buc, F. d.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32.
- Hu, Z.; Li, Y.; Lyu, J.; Gao, D.; and Vasconcelos, N. 2023. Dense Network Expansion for Class Incremental Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11858–11867.
- Keskar, N. S.; Mudigere, D.; Nocedal, J.; Smelyanskiy, M.; and Tang, P. T. P. 2017. On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima. In *International Conference on Learning Representations*.
- Lee, S.; Jeon, H.; Son, J.; and Kim, G. 2024. Learning to Continually Learn with the Bayesian Principle. In Salakhutdinov, R.; Kolter, Z.; Heller, K.; Weller, A.; Oliver, N.; Scarlett, J.; and Berkenkamp, F., eds., *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, 26621–26639. PMLR.
- Li, T.; Zhou, P.; He, Z.; Cheng, X.; and Huang, X. 2024. Friendly Sharpness-Aware Minimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5631–5640.
- Lin, H.; Zhang, B.; Feng, S.; Li, X.; and Ye, Y. 2023. PCR: Proxy-Based Contrastive Replay for Online Class-Incremental Continual Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 24246–24255.
- Lopez-Paz, D.; and Ranzato, M. A. 2017. Gradient Episodic Memory for Continual Learning. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30.
- Lu, A.; Feng, T.; Yuan, H.; Song, X.; and Sun, Y. 2024. Revisiting Neural Networks for Continual Learning: An Architectural Perspective. In *IJCAI*, 4651–4659.
- Mehta, S. V.; Patil, D.; Chandar, S.; and Strubell, E. 2023. An Empirical Investigation of the Role of Pre-training in Lifelong Learning. *Journal of Machine Learning Research*, 24(214): 1–50.
- Parisi, G. I.; Kemker, R.; Part, J. L.; Kanan, C.; and Wermter, S. 2019. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113: 54–71.
- Rebuffi, S.-A.; Kolesnikov, A.; Sperl, G.; and Lampert, C. H. 2017. iCaRL: Incremental Classifier and Representation Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2001–2010.
- Rolnick, D.; Ahuja, A.; Schwarz, J.; Lillicrap, T.; and Wayne, G. 2019. Experience Replay for Continual Learning. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; Alché-Buc, F. d.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32.

- Saha, G.; Garg, I.; and Roy, K. 2021. Gradient Projection Memory for Continual Learning. In *International Conference on Learning Representations*.
- Shi, G.; CHEN, J.; Zhang, W.; Zhan, L.-M.; and Wu, X.-M. 2021. Overcoming Catastrophic Forgetting in Incremental Few-Shot Learning by Finding Flat Minima. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P. S.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 6747–6761.
- Sun, W.; Li, Q.; Zhang, J.; Wang, W.; and Geng, Y.-A. 2023. Decoupling Learning and Remembering: a Bilevel Memory Framework with Knowledge Projection for Task-Incremental Learning. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 20186–20195.
- Sun, Z.; Mu, Y.; and Hua, G. 2023. Regularizing Second-Order Influences for Continual Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 20166–20175.
- Tran Tung, L.; Nguyen Van, V.; Nguyen Hoang, P.; and Than, K. 2023. Sharpness and Gradient Aware Minimization for Memory-based Continual Learning. In *Proceedings of the 12th International Symposium on Information and Communication Technology*, 189–196.
- van de Ven, G. M.; Tuytelaars, T.; and Tolias, A. S. 2022. Three types of incremental learning. *Nature Machine Intelligence*, 4(12): 1185–1197.
- Wang, F.-Y.; Zhou, D.-W.; Ye, H.-J.; and Zhan, D.-C. 2022. FOSTER: Feature Boosting and Compression for Class-Incremental Learning. In Avidan, S.; Brostow, G.; Cissé, M.; Farinella, G. M.; and Hassner, T., eds., *ECCV*, 398–414.
- Wang, L.; Zhang, X.; Su, H.; and Zhu, J. 2024. A Comprehensive Survey of Continual Learning: Theory, Method and Application. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(8): 5362–5383.
- Wu, Y.; Huang, L.-K.; Wang, R.; Meng, D.; and Wei, Y. 2024. Meta Continual Learning Revisited: Implicitly Enhancing Online Hessian Approximation via Variance Reduction. In *The Twelfth International Conference on Learning Representations*.
- Yang, E.; Shen, L.; Wang, Z.; Liu, S.; Guo, G.; and Wang, X. 2023. Data Augmented Flatness-aware Gradient Projection for Continual Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 5630–5639.
- Yao, Z.; Gholami, A.; Keutzer, K.; and Mahoney, M. W. 2020. PyHessian: Neural Networks Through the Lens of the Hessian. In *2020 IEEE International Conference on Big Data (Big Data)*, 581–590.
- Zhang, X.; Xu, R.; Yu, H.; Zou, H.; and Cui, P. 2023. Gradient Norm Aware Minimization Seeks First-Order Flatness and Improves Generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 20247–20257.
- Zhao, B.; Xiao, X.; Gan, G.; Zhang, B.; and Xia, S.-T. 2020. Maintaining Discrimination and Fairness in Class Incremental Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 13205–13214.
- Zhou, D.-W.; Wang, F.-Y.; Ye, H.-J.; and Zhan, D.-C. 2023a. PyCIL: a Python toolbox for class-incremental learning. *Science China Information Sciences*, 66(9): 197101.
- Zhou, D.-W.; Wang, Q.-W.; Qi, Z.-H.; Ye, H.-J.; Zhan, D.-C.; and Liu, Z. 2024. Class-Incremental Learning: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12): 9851–9873.
- Zhou, D.-W.; Wang, Q.-W.; Ye, H.-J.; and Zhan, D.-C. 2023b. A Model or 603 Exemplars: Towards Memory-Efficient Class-Incremental Learning. In *International Conference on Learning Representations*.
- Zhu, Z.; Wu, J.; Yu, B.; Wu, L.; and Ma, J. 2019. The Anisotropic Noise in Stochastic Gradient Descent: Its Behavior of Escaping from Sharp Minima and Regularization Effects. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97, 7654–7663. PMLR.
- Zinkevich, M. 2003. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning*, 928–935.