DEL: BIOLOGICALLY PLAUSIBLE DENDRITIC LEARNING ENABLES CLASS-INCREMENTAL LEARNING

Anonymous authorsPaper under double-blind review

000

001

002 003 004

010 011

012

013

014

015

016

017

018

019

021

023

025

026027028

029

031

033

034

035

037

040

041

042

043

044

045

046

047

048

051

052

ABSTRACT

Class-incremental learning (CIL) enables models to acquire new knowledge while retaining prior knowledge, thereby adapting to continuous data streams. Because parameter drift and distribution shifts are inevitable, it suffers from catastrophic forgetting and the stability-plasticity dilemma. There are various strategies to address these challenges. Nevertheless, they still remain limited by the homogeneous representations, which reduce inter-class diversity and exacerbate forgetting. To overcome this bottleneck, we introduce dendritic learning (DeL), a biologically inspired framework that reduces homogeneous representations and thereby mitigates catastrophic forgetting. DeL leverages synaptic plasticity and multi-branch dendrites to extract diverse, discriminative features, fostering heterogeneous representation learning. A membrane layer integrates these features, and a subsequent somatic layer adapts them for downstream classification. By strengthening classspecific features, DeL also promotes robust memory consolidation. Experiments show that augmenting state-of-the-art CIL methods with DeL consistently boosts accuracy. Furthermore, DeL encourages more efficient representation learning, allowing the model to rely on fewer discriminative features. Code is available at https://github.com/anonymous/DeL.

1 Introduction

Artificial intelligence has achieved significant success across a wide range of real-world applications. Numerous specialized models have been developed to address specific tasks. However, a model that performs well on one class often struggles with another because of feature and parameter shifts Zhou et al. (2024b). To maintain performance, such models typically require retraining or fine-tuning through transfer learning, which incurs additional computational cost. Real-world scenarios frequently produce non-stationary data streams that demand a model capable of continuously learning novel classes, called as class-incremental learning (CIL). However, CIL faces the chronic problem of catastrophic forgetting, whereby the model gradually loses previously learned knowledge while assimilating new information. In contrast, humans exhibit a remarkable ability to preserve and integrate knowledge over time.

To achieve human-level learning, researchers have proposed various strategies to mitigate catastrophic forgetting. Because the loss of earlier knowledge is its principal cause Robins (1995), rehearsal-based CIL methods allocate a buffer that stores salient examples of past data, drawing inspiration from human memory consolidation Zhuang et al. (2022). The model periodically revisits these samples while acquiring new information. Nevertheless, in conditions where historical data are scarce or privacy constraints preclude storage, rehearsal buffers cannot sufficiently preserve old knowledge, thereby degrading performance. Besides, parameter drift also contributes to forgetting. Expansion-based CIL approaches Zhou et al. (2024a); Wu et al. (2022) therefore instantiate multiple lightweight subnetworks, each dedicated to a subset of classes, to minimize inter-class parameter drift. Although effective, both rehearsal- and expansion-based methods entail substantial computational overhead, the former is limited by buffer capacity, whereas the latter suffers from unbounded model growth. Parameter-regularization approaches Kirkpatrick et al. (2017) first estimate how much each parameter contributes to previously learned tasks, then actively constrain the parameters they judge critical, thereby preserving prior knowledge. Although conceptually appealing, these methods perform poorly in CIL because importance estimates obtained at one incremental step often conflict with those required later Van de Ven et al. (2022). Currently, knowledge-distillation

techniques rank among the most attention research directions for mitigating forgetting in CIL Li et al. (2025). In this paradigm, a model trained on earlier tasks acts as a teacher that transfers its learned representations to a student model as the latter acquires new classes. CIL typically adopts a self-distillation variant, the teacher and student share an identical backbone while maintaining separate classification heads, enabling the teacher's predictions to guide the student and incorporate new knowledge without erasing the old. Although numerous strategies have been proposed to alleviate catastrophic forgetting by balancing the acquisition of new information with the retention of old knowledge, they still struggle to resolve the stability–plasticity dilemma, resulting in a persistent performance gap between artificial models and humans.

These limitations arise primarily from the biologically implausible architectures of conventional artificial neural networks Bellitto et al. (2024). Contemporary models employ oversimplified Mc-Culloch–Pitts (MCP) neurons that interact solely through weighted summations. Consequently, their plasticity and stability depend almost entirely on architectural design and parameter values, especially in classifiers based on fully connected networks. This homogeneous processing drives successive layers to generate highly correlated activations, hindering any principled trade-off between stability and plasticity. In contrast, the human brain consolidates memories through precise synaptic and dendritic connectivity patterns, thereby forming stable long-term engrams Ryan et al. (2015). Synaptic plasticity, which modulates the strength of these connections, provides the physiological substrate that enables humans to assimilate new information while preserving prior knowledge Mansvelder et al. (2019). Motivated by this evidence, we argue that embedding biologically inspired mechanisms, specifically, synaptic plasticity for rapid adaptation and structured connectivity for durable memory, could markedly narrow the stability–plasticity gap in class-incremental learning.

To this end, we propose a biologically plausible dendritic learning (DeL) for CIL. DeL models neurons with a more realistic structure than the traditional MCP abstraction, incorporating a synaptic layer, a dendritic layer, a membrane layer, and a somatic layer. The synaptic layer that receives features from the backbone and promotes diverse activations to highlight discriminative features. The dendritic layer is composed of multiple branches that repeatedly process synaptic signals with distinct plasticity profiles, thereby sustaining long-term potentiation, which is an essential mechanism for memory formation Toni et al. (1999). Finally, the membrane and the somatic layers that integrate the dendritic responses to produce class-specific outputs. DeL is model-agnostic and can be attached to any CIL method. We integrate DeL into several representative CIL paradigms and observe consistent accuracy improvements, confirming its ability to mitigate catastrophic forgetting. We further evaluate DeL on multiple benchmarks to characterize its stability–plasticity trade-off and conduct extensive ablation studies to quantify the contribution of each architectural component. Collectively, these results demonstrate that a biologically plausible dendritic learner can extract discriminative features from streaming data and perform robust incremental classification.

2 Related Work

Class-incremental Learning (CIL) seeks to prevent catastrophic forgetting while balancing model stability and plasticity. Prior work can be summarized into four main categories. Rehearsal-based methods mimic human complementary learning by periodically revisiting a small buffer of past samples. The key challenge is selecting representative exemplars that preserve the discriminative structure of previous tasks. Recent work adopts multi-criteria selection to identify such exemplars Rolnick et al. (2019); Zhuang et al. (2022). Parameter-regularization methods estimate the importance of each parameter to previously learned classes and constrain important ones during subsequent training. EWC Kirkpatrick et al. (2017) introduces this paradigm, storing a Fisher-based importance matrix that matches the network's dimensionality. Expansion-based methods dynamically grow their architectures, allocating new capacity for incoming tasks to mitigate parameter drift. DER Yan et al. (2021) appends an additional backbone for each new task and aggregates features through an enlarged fully connected layer. Although such methods often achieve state-of-the-art accuracy, they demand substantial memory and computation. Knowledge-distillation methods transfer information from a frozen "teacher" to a "student", learning new classes. LwF Li & Hoiem (2018) first applied logits-based distillation to align outputs of model across tasks. iCaRL Rebuffi et al. (2017) augments it with prototype rehearsal exemplars. Subsequent work incorporates feature-level distillation to pre-




Figure 1: Overview of class-incremental learning with DeL. *Left*: The CIL process incorporating DeL. *Right*: The detailed architecture of the DeL module.

serve intermediate representations. TagFex Zheng et al. (2025) combines task-agnostic expansion with simultaneous feature and logits distillation to encourage diverse representation learning.

Collectively, these four research streams provide complementary perspectives on mitigating forgetting, yet each still faces trade-offs among accuracy, efficiency, and biological plausibility. By contrast, our method introduces a fully articulated dendritic architecture that leverages synaptic plasticity to boost feature diversity and safeguard long-term knowledge. Furthermore, its dendritic branches adaptively modulate the connection strengths between synapses and dendrites, enabling the model to extract discriminative features more effectively during incremental learning.

3 Method

3.1 Problem Definition

CIL seeks to build a unified classifier that can assimilate new classes from a continuously evolving data stream while preserving knowledge of previously learned classes. Let the data stream be $\mathbf{D} = \{\mathfrak{D}_1, \mathfrak{D}_2, \dots, \mathfrak{D}_L\}$, where the lth incremental batch is $\mathfrak{D}_l = \{(\mathbf{x}_i^l, y_i^l)\}_{i=1}^{n_l}$, containing n_l samples. Each instance $\mathbf{x}_i^l \in \mathbb{R}^D$ is paired with a label $y_i^l \in Y_l$, in which Y_l is the label sets of incremental class l. Notably, $\forall i \neq l, Y_i \cap Y_l = \emptyset$. A corresponding test set is $\mathbf{D}^t = \{\mathfrak{D}_1^t, \mathfrak{D}_2^t, \dots, \mathfrak{D}_L^t\}$. Under the non-rehearsal assumption, samples from earlier increments are not reused when learning subsequent ones. After the lth increment, the model is evaluated on the cumulative label set $\mathcal{Y}_l = \bigcup_{k=1}^l Y_k$, thereby measuring its ability both to acquire new knowledge and to retain what it has previously learned.

3.2 DEL: BIOLOGICALLY PLAUSIBLE DENDRITIC LEARNING

Our method is grounded in neurophysiological evidence showing that behavioral learning leads to changes in synaptic strength, and that manipulating synaptic strength can alter the information stored in memory Whitlock et al. (2006); Nabavi et al. (2014). Synaptic plasticity mechanisms that enhance synaptic strength contribute not only to the formation of memory traces but also to their reactivation during memory recall Ryan et al. (2015); Nonaka et al. (2014). We embed these mechanisms in a dendritic learning (DeL) network to emulate human learning and memory formation, as well as to confront the challenges of CIL. DeL comprises four functional layers, i.e., synaptic, dendritic, membrane, and somatic, working in concert to balance plasticity and stability throughout incremental learning, as shown in Figure 1. Moreover, conventional models process information homogeneously, therefore, they must learn many redundant features to separate classes, which strains memory capacity Jaini et al. (2018). Once parameter or distribution shifts, such models struggle to generate reliable features, leading to rapid forgetting. Mitigating these shifts typically requires either adopting architectures with additional parameters or using more distinctive features, both of




Figure 2: Grad-CAM Selvaraju et al. (2017) visualizations for the fully-connection (FC) and DeL networks. The upper image corresponds to the FC model, and the lower image to DeL. DeL concentrates on highly discriminative regions, thereby reducing memory pressure.

which lessen memory pressure and diminish catastrophic forgetting. DeL's biologically inspired architecture and adaptive synaptic scaling extract discriminative features with far less redundancy. Figure 2 highlights this advantage. DeL is attention to highly distinctive features than conventional ones, and utilize them to make accurate classification. In addition, DeL is model-agnostic and can be seamlessly integrated into any CIL method by replacing the original logit classifier with DeL. In the case of knowledge distillation–based models, only logit distillation through DeL is required, without the need for any additional modifications.

3.2.1 SYNAPTIC LAYER

The layer ingests the feature stream produced by the backbone at every incremental step. To overcome the limitation of homogeneous representations, it must discriminate and refine task-specific features. Because class distributions differ, treating all inputs uniformly degrades performance. Biological neurons modulate synaptic strength in response to each input, amplifying salient signals and attenuating less informative ones, before forwarding them to the dendrites. We emulate this synapse plasticity by applying a learnable sigmoid function that maps each feature to the interval of (0,1). Meanwhile, considering the distribution difference of activations from backbone, we utilize the learnable function to evaluate their synaptic strength, i.e.,

$$S_{i,j} = \frac{1}{1 + e^{-(w_{i,j}x_i - \theta_{i,j})}} \tag{1}$$

where $w_{i,j}$ and $\theta_{i,j}$ denote the weight and threshold parameters of the learnable sigmoid that modulates the synaptic strength of dendrite j, thereby producing diverse connection patterns. The details of the synaptic layer are given in the supplementary file.

3.2.2 DENDRITIC LAYER

The dendritic layer receives heterogeneous synaptic signals and executes complementary functions. The synaptic layer forces discriminative features and can down-weight less informative connections, thereby reducing the feature set. However, over-pruning still risks degrading classification accuracy. To balance selectivity and coverage, we introduce a dendritic layer with multi-branch that mirrors the anatomical complexity of biological dendrites. Given the same input, each branch applies a distinct pattern of synaptic plasticity, enabling the network to capture diverse functional combinations and lessening its reliance on any single feature subset. Dendritic layer integrates synaptic outputs as $D_j = \sum_{i=1}^N S_{i,j}$, where $S_{i,j}$ is the ith synaptic signal on dendrite j. N is the input size. This repeated aggregation strengthens plasticity events, promoting robust learning and memory consolidation.

Table 1: Average and last performance comparison on six datasets with **ResNet-18** as the backbone. We report all compared methods with their source-code provided by PyCIL library and trained from scratch. The best performance is shown in bold.

Methods	CIFAR100		CUB200		VTAB		ImageNet-A		ImageNet-R		OmniBench	
	$ar{A}$	A_L	$ar{A}$	A_L	\bar{A}	A_L	\bar{A}	A_L	$ar{A}$	A_L	\bar{A}	A_L
Finetune	25.75	9.00	21.72	10.39	41.30	29.70	6.73	2.57	14.10	6.60	24.34	8.89
w/ DeL	26.78	9.20	25.88	11.32	42.81	26.99	7.18	2.11	14.96	7.35	24.61	8.96
Replay	60.60	43.66	38.30	29.81	53.72	47.74	6.06	2.30	23.23	15.67	51.54	34.32
w/ DeL	64.01	47.45	46.83	38.80	58.83	49.90	6.24	1.91	28.40	21.28	53.93	36.37
LwF	46.27	26.31	27.50	12.98	44.72	27.10	6.60	3.03	18.67	9.25	37.63	19.67
w/ DeL	52.70	30.32	33.34	18.62	48.49	31.27	8.36	3.23	23.27	12.52	40.90	20.50
iCarL	60.72	43.65	36.77	25.06	58.09	53.16	6.87	2.44	22.20	14.47	51.60	31.81
w/ DeL	64.06	46.59	45.33	33.29	62.49	54.61	6.83	2.57	27.56	18.63	53.42	33.78
DER	72.92	62.77	45.72	40.50	58.03	51.07	7.10	2.50	33.06	30.88	65.74	55.99
w/ DeL	74.68	64.31	50.48	43.98	64.51	56.22	6.30	2.50	37.74	34.50	66.69	58.45
TagFex	73.69	62.55	51.08	44.27	61.79	48.54	5.89	0.07	41.66	38.08	65.45	56.78
w/ DeL	75.25	65.68	55.43	48.52	65.41	56.19	6.81	2.30	45.15	40.83	66.07	57.76

3.2.3 MEMBRANE LAYER

The membrane layer aggregates the outputs of the dendritic branches. Because each dendrite performs a distinct function, the membrane applies a weighted fusion, i.e., $I = \sum_{j=1}^{M} g_j D_j$, where g_j is a learnable, branch-specific gain. The adaptive gain further strengthens the model's capacity to learn novel information while preserving previously acquired knowledge.

3.2.4 SOMATIC LAYER

The somatic layer performs the final transformation, adapting the aggregated signal to the requirements of the downstream task, $O=f(I-\Delta)$, where $f(\cdot)$ is the activation function and Δ is a learnable threshold that adjusts to different tasks. Typically, the somatic layer applies a sigmoid function to the outputs of the membrane layer, producing the logits for each class. In classification problems, these logits reflect the model's confidence in each class, with larger values indicating higher certainty.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

4.1.1 Datasets

Follow the work Sun et al. (2025), we use six datasets to verify the performance, including CI-FAR100 Krizhevsky et al. (2009), CUB200 Wah et al. (2011), ImageNet-A Hendrycks et al. (2021b), ImageNet-R Hendrycks et al. (2021a), VTAB Zhai et al. (2019), and OmniBenchmark Zhang et al. (2022). These datasets contain 100, 200, 200, 200, 50, and 300 classes, respectively, and we construct incremental streams of 10, 10, 20, 20, 10, and 30 new classes per step, starting from an initial zero-class state. Their details are summarized in the supplementary file.

4.1.2 Baseline Methods

To evaluate DeL's performance, we integrated it into six representative CIL algorithms. Fine-tune serves as a plain baseline that employs no forgetting-mitigation strategy. Replay represents rehearsal-based methods, whereas DER Yan et al. (2021) follows an expansion strategy. LwF Li & Hoiem (2018) and iCaRL Rebuffi et al. (2017) both rely on knowledge distillation, and TagFex Zheng et al. (2025) combines knowledge distillation with rehearsal and dynamic expansion. More details of all models are provided in the supplementary file.




Figure 3: Performance curve of different methods with and without DeL. The solid and dot-line denote the variant with and without DeL, respectively.

Table 2: The ablation study of Synaptic Plasticity on CIFAR 100 with **DER**. The best performance is shown in bold.

SA	SN	CIFA	R100	VTAB		
		$ar{A}$	A_L	$ar{A}$	A_L	
X	X	71.93	61.18	58.03	49.30	
\checkmark	X	74.68	64.31	64.51	56.22	
\checkmark	\checkmark	73.69	63.58	63.85	56.19	

4.1.3 IMPLEMENTATION DETAILS

For a fair comparison, we keep the training and evaluation protocols identical across all methods. Every model is implemented in PyTorch using the PyCIL Sun et al. (2023) and trained from scratch on an NVIDIA RTX 3090 GPU with ResNet18 backbone. Optimization is performed with SGD (initial learning rate =0.1, weight decay $=5\times10^{-4}$). The initial task is trained for 200 epochs. We use the random seed of 1993 to guarantee the consistency of data partition.

4.1.4 MEASUREMENT METRIC

We assess model performance with two complementary metrics. One of them is the top-1 accuracy after the final increment (A_L) , another one is the average incremental accuracy \bar{A} . Let $A_l(l=1,\ldots,L)$ denote the top-1 accuracy measured after the lth increment on the cumulative test set that includes every class learned thus far. The average incremental accuracy is calculated as $\bar{A}=\frac{1}{L}\sum_{l=1}^{L}A_l$.

4.2 Performance Results

For each model, we compute both A and A_L , then compare those with the DeL-enhanced variants. The experimental results are summarized in Table 1. DeL boosts replay methods by 4.1% in \bar{A} and 3.7% in A_L , and raises LwF, a knowledge-distillation approach, by 4.2% and 3.0%, respectively. It achieves improvements of 3.9% and 3.1% in iCaRL. As baselines, the expansion-based methods, DER and TagFex, outperform other types of methods, which indicates that larger parameters efficiently inhibit parameter shift and store more features to prevent forgetting. DeL further improves their performance. The experimental results strongly demonstrate that synaptic plasticity in DeL produces heterogeneous and highly discriminative representations, enabling the network to rely on a smaller yet more informative feature set and thus reducing memory overhead while mitigating catas-

Table 3: The parameter analysis of DeL on VTAB with **DER**. The best performance is shown in bold.

M	1	2	4	8
\bar{A}	63.73	64.51	64.43	62.54
A_L	54.91	56.22	56.09	55.80




Figure 4: Grad-CAM visualizations across incremental tasks T_1 to T_5 . Top: Grad-CAM results from iCaRL. Bottom: Grad-CAM results from iCaRL with DeL.




Figure 5: The Grad-CAM of activation maximization maps, which are prone to catastrophic forgetting due to their dependence on the classifier.

trophic forgetting. Furthermore, we give the performance curve in Figure 3, showing that methods with DeL achieve a large performance gap compared to methods without it. Moreover, the additional computation and train time are reported in supplementary file.

4.3 ABLATION STUDY

We conduct an ablation study on CIFAR-100 and VTAB to quantify the contribution of synaptic plasticity in DeL. The results are summarized in Table 2. **SA** denotes the learnable synaptic-plasticity function, whereas **SN** performs input normalization through layer normalization Ba et al. (2016). The results reveal that incorporating synaptic plasticity markedly improves performance of the model. Moreover, the modest drop when incorporating normalization is removed suggests that the synaptic layer accommodates substantial distribution shifts, allowing the model to learn effectively from streaming data. DeL therefore maintains strong performance even without explicit normalization.

Dendritic morphology endows the network with greater computational capacity and robust learning dynamics Chavlis & Poirazi (2025). It also modulates plasticity; therefore, we examine how the number of dendritic branches M affects performance, as shown in Table 3. In principle, increasing M enhances synaptic plasticity and nonlinear expressiveness, yet excessive plasticity undermines stability. Our results show that M=2 offers the best trade-off, whereas larger or smaller values lead to performance declines.

4.4 VISUALIZATIONS

We employ Grad-CAM Selvaraju et al. (2017) to visualize model's attention regions and assess the impact of homogeneous representations. As Figure 4 illustrates, baseline iCaRL attends dispropor-




Figure 6: The weight probability density functions after class-incremental learning for iCaRL with and without DeL.




Figure 7: Information entropy distributions for iCaRL with and without DeL. Entropies are computed from classifier activations on all test images of the CUB dataset during class-incremental learning.

tionately to non-target regions, generating spurious memories that accelerate subsequent forgetting, which evidences that networks built on MCP neurons struggle to isolate class-relevant features due to homogeneous representation. In contrast, iCaRL enhanced with DeL consistently focuses on truly discriminative cues, most notably the bird's body and head, thereby preserving salient attention patterns and maintaining robust feature representations. This advantage is particularly pronounced at step T_2 . These observations suggest that DeL effectively curbs representation drift and efficiently store memory, yielding a more favorable plasticity–stability trade-off.

Furthermore, Figure 5 shows that the baseline iCaRL model correctly attends to the bird's body in the early tasks, however, as class-incremental learning progresses its attention drifts toward background regions, indicating catastrophic forgetting. By the final tasks, Grad-CAM no longer highlights critical features such as the head, chest, and wings, and the representations of earlier classes have markedly decayed. In contrast, once we embed DeL, the model's attention remains stable across all ten tasks. High-response regions (in red) consistently cover key features, chiefly the chest and back. Therefore, even after the final task the model localizes these regions while relying on fewer overall attention areas. This behavior evidences that DeL enhances neural selectivity, preserves memory, and effectively resists catastrophic forgetting by suppressing homogeneous representations and promoting truly discriminative features.

4.5 Analysis of Effectiveness

To better understand how DeL mitigates catastrophic forgetting, we analyze its weight distribution after each incremental step on the CUB dataset in Figure 6. From this figure, the weight distribution of conventional iCaRL shifts at every stage, revealing parameter drift that drives forgetting. By contrast, incorporating DeL can suppress the distribution drift to a certain extent. Importantly, DeL exhibits a more concentrated distribution, showing that DeL can efficiently utilize weights to constrain drift and concentrate model capacity on the most salient features via synaptic plasticity. Thereby it can preserve earlier knowledge while assimilating new classes.

Furthermore, we compute information entropy for each neuron to quantify its class specificity and importance for classification results. High entropy indicates mixed selectivity, meaning a neuron responds to multiple classes, whereas low entropy reflects class-specific tuning. Figure 7 plots the




Figure 8: Probability density functions of class selectivity. Each histogram illustrates the distribution of neuron selectivity across classes at each incremental learning stage. The number of bins corresponds to the total number of classes in the CUB dataset.

entropy distribution after each incremental step. The baseline model shows uniformly high entropy, suggesting that nearly all neurons participate in every classification decision and therefore struggle to preserve the class-specific features needed to recall old knowledge. In contrast, the DeL-augmented model displays a broader and often bimodal entropy distribution. Many neurons shift toward lower entropy values, demonstrating distinct activations for different classes, while a smaller subset retains higher entropy, preserving flexibility for new tasks. This adaptive redistribution implies that DeL dynamically adjusts the proportion of class-specific neurons as new classes arrive, which mitigates catastrophic forgetting. Specifically, the entropy distribution becomes bimodal at certain increments, suggesting that the model adjusts the proportion of class-specific neurons to match the current class composition. Detailed calculation procedures appear in the supplementary file. Overall, the results confirm that DeL's synaptic-plasticity mechanism fosters a richer pool of class-specific neurons, enabling the network to remember discriminative features more effectively during incremental learning.

Besides, the entropy analysis shows that DeL yields a larger pool of class-specific neurons through synaptic plasticity, yet the exact specificity of these neurons remains uncertain. To clarify how many classes each neuron represents, we compute a selection index and apply a 0.5 activation threshold due to the use of sigmoid function. Figure 8 presents the resulting class-activation counts. At every incremental step, neurons in the baseline model activate predominantly for the most recently learned classes, with few responses to earlier classes. In contrast, neurons in the DeL-enhanced model respond to a broader range of classes; although the network still favors new information, it preserves considerably more activations for previously learned classes. These findings provide additional evidence that DeL effectively mitigates catastrophic forgetting.

5 CONCLUSION

We propose a biologically inspired dendritic learning module (DeL) designed to enhance class-incremental learning by alleviating the issues of catastrophic forgetting and homogeneous representation. DeL integrates synaptic plasticity and dendritic computation into a four-layer architecture that promotes diverse and discriminative feature extraction. Extensive experiments demonstrate that DeL improves the performance of representative CIL methods across multiple benchmarks. As model-agnostic, we discuss the performance of DeL by incorporating it into various types of CIL methods. Grad-CAM visualizations show that DeL captures discriminative features and directs more attention to task-relevant object regions. Because it requires fewer stored exemplars, DeL alleviates catastrophic forgetting. Furthermore, through ablation and analysis of effectiveness, we show that DeL reduces parameter drift, increases the proportion of class-specific neurons, and stabilizes attention to semantically meaningful image regions. These results confirm that embedding biologically plausible mechanisms into artificial neural networks can significantly enhance continual learning performance by improving feature selectivity and memory retention. Although DeL is advances CIL through biological mechanism, its present architecture lacks flexibility because its hyperparameters are fixed. We will explore dynamic DeL structures to enhance scalability in future work.

REFERENCES

- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint* arXiv:1607.06450, 2016.
- Giovanni Bellitto, Federica Proietto Salanitri, Matteo Pennisi, Matteo Boschini, Lorenzo Bonicelli, Angelo Porrello, Simone Calderara, Simone Palazzo, and Concetto Spampinato. Saliency-driven experience replay for continual learning. *Advances in Neural Information Processing Systems*, 37:103356–103383, 2024.
- Spyridon Chavlis and Panayiota Poirazi. Dendrites endow artificial neural networks with accurate, robust and parameter-efficient learning. *Nature Communications*, 16(1):943, 2025.
- Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 8340–8349, October 2021a.
- Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15262–15271, June 2021b.
- Priyank Jaini, Pascal Poupart, and Yaoliang Yu. Deep homogeneous mixture models: representation, separation, and approximation. *Advances in Neural Information Processing Systems*, 31, 2018.
- 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.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, 2009.
- Songze Li, Tonghua Su, Xu-Yao Zhang, and Zhongjie Wang. Continual learning with knowledge distillation: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 36(6): 9798–9818, 2025.
- Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2935–2947, 2018.
- Huibert D Mansvelder, Matthijs B Verhoog, and Natalia A Goriounova. Synaptic plasticity in human cortical circuits: cellular mechanisms of learning and memory in the human brain? *Current Opinion in Neurobiology*, 54:186–193, 2019.
- Sadegh Nabavi, Rocky Fox, Christophe D Proulx, John Y Lin, Roger Y Tsien, and Roberto Malinow. Engineering a memory with ltd and ltp. *Nature*, 511(7509):348–352, 2014.
- Ayako Nonaka, Takeshi Toyoda, Yuki Miura, Natsuko Hitora-Imamura, Masamitsu Naka, Megumi Eguchi, Shun Yamaguchi, Yuji Ikegaya, Norio Matsuki, and Hiroshi Nomura. Synaptic plasticity associated with a memory engram in the basolateral amygdala. *Journal of Neuroscience*, 34(28): 9305–9309, 2014.
- Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- Anthony Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 7(2): 123–146, 1995.
- David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. *Advances in Neural Information Processing Systems*, 32, 2019.

- Tomás J Ryan, Dheeraj S Roy, Michele Pignatelli, Autumn Arons, and Susumu Tonegawa. Engram cells retain memory under retrograde amnesia. *Science*, 348(6238):1007–1013, 2015.
 - Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 618–626, 2017.
 - Hai-Long Sun, Da-Wei Zhou, Hanbin Zhao, Le Gan, De-Chuan Zhan, Da-Wei Ye, Han-Jia Zhou, Fu-Yun Wang, Han-Jia Ye, and De-Chuan Zhan. PyCIL: a python toolbox for class-incremental learning. *Science China Information Sciences*, 66(197101), 2023.
 - Hai-Long Sun, Da-Wei Zhou, Hanbin Zhao, Le Gan, De-Chuan Zhan, and Han-Jia Ye. Mos: Model surgery for pre-trained model-based class-incremental learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(19):20699–20707, Apr. 2025.
 - Nicolas Toni, P-A Buchs, Irina Nikonenko, CR Bron, and Dominique Muller. Ltp promotes formation of multiple spine synapses between a single axon terminal and a dendrite. *Nature*, 402 (6760):421–425, 1999.
 - Gido M Van de Ven, Tinne Tuytelaars, and Andreas S Tolias. Three types of incremental learning. *Nature Machine Intelligence*, 4(12):1185–1197, 2022.
 - C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.
 - Jonathan R Whitlock, Arnold J Heynen, Marshall G Shuler, and Mark F Bear. Learning induces long-term potentiation in the hippocampus. *Science*, 313(5790):1093–1097, 2006.
 - Tz-Ying Wu, Gurumurthy Swaminathan, Zhizhong Li, Avinash Ravichandran, Nuno Vasconcelos, Rahul Bhotika, and Stefano Soatto. Class-incremental learning with strong pre-trained models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9601–9610, 2022.
 - Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3014–3023, June 2021.
 - Xiaohua Zhai, Joan Puigcerver, Alexander Kolesnikov, Pierre Ruyssen, Carlos Riquelme, Mario Lucic, Josip Djolonga, Andre Susano Pinto, Maxim Neumann, Alexey Dosovitskiy, et al. A large-scale study of representation learning with the visual task adaptation benchmark. *arXiv* preprint arXiv:1910.04867, 2019.
 - Yuanhan Zhang, Zhenfei Yin, Jing Shao, and Ziwei Liu. Benchmarking omni-vision representation through the lens of visual realms. In *European Conference on Computer Vision*, pp. 594–611. Springer, 2022.
 - Bowen Zheng, Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Task-agnostic guided feature expansion for class-incremental learning. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 10099–10109, 2025.
 - Da-Wei Zhou, Hai-Long Sun, Han-Jia Ye, and De-Chuan Zhan. Expandable subspace ensemble for pre-trained model-based class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 23554–23564, 2024a.
 - Da-Wei Zhou, Qi-Wei Wang, Zhi-Hong Qi, Han-Jia Ye, De-Chuan Zhan, and Ziwei Liu. Class-incremental learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):9851–9873, 2024b.
 - Chen Zhuang, Shaoli Huang, Gong Cheng, and Jifeng Ning. Multi-criteria selection of rehearsal samples for continual learning. *Pattern Recognition*, 132:108907, 2022.