

# Distillation-Guided Structural Transfer for Continual Learning Beyond Sparse Distributed Memory

Huiyan Xue<sup>1\*</sup>, Xuming Ran<sup>2\*</sup>, Yaxin Li<sup>1</sup>, Qi Xu<sup>1†</sup>, Enhui Li<sup>1</sup>, Yi Xu<sup>1</sup>, Qiang Zhang<sup>1†</sup>

<sup>1</sup>School of Computer Science and Technology, Dalian University of Technology, China

<sup>2</sup>National University of Singapore, Singapore

xuqi@dlut.edu.cn, zhangq@dlut.edu.cn

## Abstract

Sparse neural systems are gaining traction for efficient continual learning due to their modularity and low interference. Architectures like Sparse Distributed Memory Multi-Layer Perceptrons (SDMLP) construct task-specific subnetworks via Top-K activation and have shown resilience against catastrophic forgetting. However, their rigid modularity poses two fundamental challenges: (1) the isolation of sparse subnetworks severely limits cross-task knowledge reuse; and (2) increased sparsity reduces interference but often degrades performance due to constrained feature sharing. We propose **Selective Subnetwork Distillation (SSD)**, a structurally guided continual learning framework that treats distillation not as a regularizer, but as a topology-aligned information conduit. By identifying neurons with high activation frequency, SSD selectively distills knowledge within previous Top-K subnetworks and output logits—without requiring replay or task labels—preserving both sparsity and functional specialization. Unlike conventional distillation, SSD operates under hard modular constraints and enables structural realignment without altering the sparse architecture. While our method is validated on SDMLP, its structure-aligned mechanism has the potential to generalize to other sparse networks as a plug-in module for promoting representation sharing. Comprehensive experiments on Split CIFAR-10, CIFAR-100, and MNIST demonstrate that SSD improves accuracy, retention, and manifold coverage, offering a structurally grounded solution to sparse continual learning.

## Introduction

In Continual Learning (CL), models are required to sequentially acquire new tasks while avoiding the forgetting of previously learned knowledge (French 1999). However, artificial neural networks often suffer from catastrophic forgetting during sequential learning, where representations of earlier tasks are overwritten by new ones, severely degrading long-term performance (De Lange et al. 2021; Aljundi et al. 2019).

Inspired by the sparse activation mechanisms in biological neural systems (Kanerva 1988), Sparse Distributed

Memory (SDM)-based improvements to Multi-Layer Perceptrons (MLPs) have been proposed (Li, Smith, and Doe 2023; Sung, Nair, and Raffel 2021). These methods leverage the Top-K activation function to retain only the most active neurons in each layer, dynamically forming input-dependent subnetworks (Frankle and Carbin 2019; Evci et al. 2022). This sparsity reduces task interference at the parameter level, enabling continual learning without replay (Mai et al. 2022). Nevertheless, the subnetworks activated for different tasks in SDMLP are typically disjoint and structurally isolated, which restricts cross-task knowledge transfer and limits the reuse of prior representations. This independence degrades generalization in dynamic environments, particularly under high sparsity, where the model struggles with representation bottlenecks and a trade-off emerges between accuracy and sparsity.

To address this structural bottleneck, we propose **Selective Subnetwork Distillation (SSD)**, which enhances continual learning through structure-aware knowledge transfer. Instead of relying on global distillation or task labels, SSD selects Top- $n$  neurons based on their historical activation frequency to perform layer-wise distillation and construct hierarchical connection paths between corresponding subnetworks. This structural routing mechanism, inspired by selective activation in neurophysiology (Olshausen and Field 2023), effectively establishes inter-subnetwork knowledge channels, enabling structured sharing and transfer (Ran et al. 2024). SSD surpasses SDMLP in both structural design and functional capacity. It addresses SDMLP’s core limitations: disconnected subnetworks and insufficient knowledge reuse, and the trade-off between accuracy and sparsity, thereby providing a structural enhancement (Wang et al. 2022) strategy tailored for sparse continual learning.

In contrast to conventional knowledge distillation techniques that operate in the output space via loss functions (Hinton, Vinyals, and Dean 2015; Gou et al. 2021), SSD functions not as a loss adjustment but as an architectural mechanism. It explicitly connects sparse task-specific subnetworks through inter-layer paths, enabling transferable representation sharing at the architectural level. This structural advantage allows SSD to function without sample replay or task labels, making it well-suited for more complex and dynamic continual learning scenarios (Bhat, Zonooz, and Arani 2022; Hafez and Erekmen 2024).

\*These authors contributed equally.

†Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

The primary contributions of this work are summarized as follows:

- A systematic analysis of structural bottlenecks in sparse continual learning frameworks such as SDMLP, particularly in terms of knowledge transfer and representational capacity (De Lange et al. 2021; Li et al. 2024);
- The proposal of **SSD**, a structure-enhanced method that performs layer-wise selective distillation based on neuron activation and subnetwork connectivity, achieving cross-task knowledge transfer without relying on replay or task labels;
- Extensive experiments on standard continual learning benchmarks including Split CIFAR-10, CIFAR-100, and MNIST, where SSD achieves significant improvements in classification accuracy and knowledge retention, validating its effectiveness and generality within sparse continual learning settings.

The remainder of this paper is organized as follows: Section 2 reviews related work on sparse continual learning and knowledge distillation; Section 3 introduces the background and mechanisms of Sparse Distributed Memory and SDMLP; Section 4 presents the design and training strategy of the proposed Selective Subnetwork Distillation; Section 5 reports experimental results and ablation studies; and Section 6 discusses the method’s advantages, limitations, and potential extensions, followed by the conclusion.

## Related Work

### Sparse Subnetwork-Based Continual Learning Methods

Continual learning techniques can generally be categorized into three major paradigms: architectural approaches (Goodfellow et al. 2013; Springenberg et al. 2015), regularization-based approaches (Smith et al. 2022; Kirkpatrick et al. 2017; Zenke, Poole, and Ganguli 2017; Aljundi, Chakravarty, and Tuytelaars 2018), and replay-based methods (De Lange et al. 2021; Hsu et al. 2018). Among these, architectural approaches, which design task-specific network structures, often leverage sparse subnetwork strategies to minimize inter-task interference. (Aljundi et al. 2019; Xu et al. 2018; Mallya and Lazebnik 2018b; Schwarz et al. 2021b; Smith et al. 2022).

Specifically, some methods use weight masking or dynamic routing to partition pre-trained networks into task-specific subnetworks. For instance, Piggyback (Mallya and Lazebnik 2018b) learns binary masks to select substructures on fixed weights for new tasks, Supermasks in Superposition (Wortsman, Farajtabar et al. 2020) identifies sparse subnetworks without altering weights, and PackNet (Mallya and Lazebnik 2018a) prunes redundant weights to allocate capacity for subsequent tasks. These approaches typically require explicit task boundaries and additional memory for masks or weight snapshots. Alternatively, sparse activation mechanisms limit active neurons per layer to dynamically form task-specific subnetworks, well-suited for task-free continual learning. However, these methods often lack effective knowledge transfer across tasks, limiting long-term generalization.

## Knowledge Distillation in Continual Learning

Knowledge distillation (KD) (Hinton, Vinyals, and Dean 2015) has been widely applied to continual learning. Following the taxonomy in (Li et al. 2024), KD-based continual learning approaches can be categorized into three paradigms: regularization-based, data replay-based, and feature replay-based distillation.

Regularization-based methods constrain the current model’s predictions to match those of the previous model, enabling knowledge retention without data replay, such as Learning without Forgetting (LwF) (Li and Hoiem 2016) using logits distillation and Learning without Memorizing (LwM) (Dhar et al. 2019) employing attention-based feature alignment. Data replay-based methods store or generate past task samples, like iCaRL (Rebuffi et al. 2017) combining exemplars with distillation loss, Deep Generative Replay (DGR) (Shin et al. 2017) synthesizing data, and SparCL (Wang et al. 2022) optimize replay efficiency for resource-limited devices. Feature replay-based methods reconstruct old task feature spaces to avoid raw data storage, as in GFR (Belouadah, Popescu, and Kanellos 2020) and PLOP (Douillard et al. 2021) extends this idea to dense prediction tasks such as semantic segmentation.

These methods significantly enhance continual learning but focus on dense architectures and often rely on replay, limiting sparse network applications. Our Sparse Subnetwork Distillation (SSD) balances sparsity and cross-task connectivity, unlike LwF, iCaRL, DGR, or mask-based (Mallya and Lazebnik 2018b) and pruning-based (Mostafa and Wang 2019) methods, offering a lightweight, scalable continual learning solution.

## Background on Sparse Distributed Memory

Sparse Distributed Memory (SDM) is an associative memory model inspired by neurobiological principles, using sparse activation patterns to store and retrieve information based on similarity-driven neuron selection. Bricken et al. (Kanerva 1988) extended this principle to neural networks, introducing the Sparse Distributed Multi-Layer Perceptron (SDMLP), a continual learning framework that forms sparse, task-specific subnetworks to mitigate catastrophic forgetting.

SDMLP integrates several biologically inspired mechanisms into a standard MLP architecture. It first applies a **Top-K activation function** (Cekic, Bakiskan, and Madhoo 2022), selecting the  $K$  most active neurons per layer to dynamically form sparse subnetworks. Both inputs and weights are **L2-normalized**, and bias terms are removed to ensure balanced neuron participation and prevent dominance by a few neurons. Inspired by Dale’s principle, **non-negative weight constraints** (Dale 1935) enhance biological plausibility and sparse connectivity. Finally, SDMLP employs **SGD without momentum** to avoid stale gradient updates that may deactivate neurons. Through these mechanisms, SDMLP effectively isolates task-relevant knowledge within dedicated subnetworks, mitigating catastrophic forgetting in continual learning.

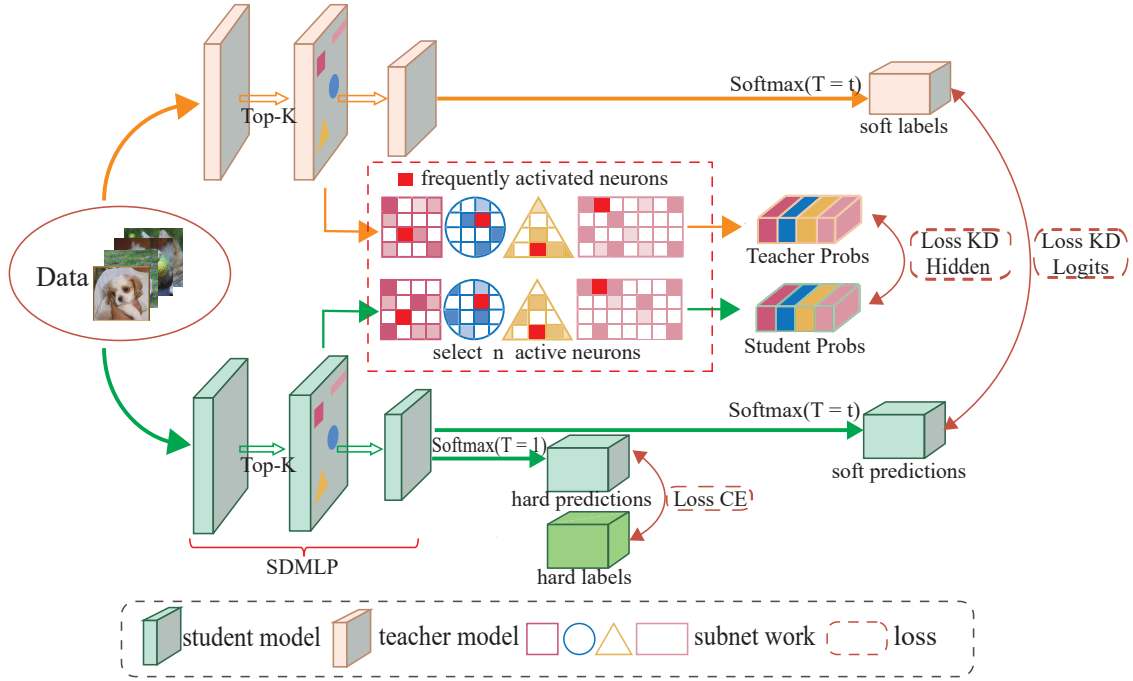


Figure 1: Overview of the Selective Subnetwork Distillation (SSD) framework. Given the same input, both teacher and student SDMLPs activate task-specific sparse subnetworks via Top-K selection. SSD identifies the Top- $n$  most frequently activated neurons across training and selectively distills their hidden activations, preserving structural sparsity while aligning task-relevant representations. Shared neurons across tasks (red blocks) promote subnetwork connectivity and enhance representation reuse. SSD further distills the output logits for stable decision boundaries.

## Selective Subnetwork Distillation

### Overview

We propose Selective Subnetwork Distillation (SSD), a structural framework that enhances sparse subnetwork connectivity through selective neuron routing to address the limitations of SDMLP in cross-task knowledge transfer. SSD employs a teacher-student architecture, where the student model is the SDMLP for the current task, and the teacher model, starting from task 1, is the SDMLP trained on the previous task to retain prior knowledge (Li, Smith, and Doe 2023; French 1999). Due to the sequential nature of continual learning, task 0 lacks prior knowledge for distillation, so SSD activates the teacher model from task 1 onward.

Unlike conventional self-distillation, SSD selectively distills the Top- $n$  active neurons from the previous task, enhancing inter-task representation transfer while preserving inactive neurons for dynamic subnetwork construction in new tasks (Olshausen and Field 2023; Frankle and Carbin 2019). The SSD framework is illustrated in Figure 1, with its workflow comprising three steps: neuron selection, layer-wise distillation, and loss function optimization.

### Selective Neuron Identification

The core of SSD lies in identifying neurons essential for retaining prior knowledge within sparse subnetworks. Here, **Top-K** denotes the sparse activation with a fixed number  $K$

of active neurons per layer, whereas **Top- $n$**  refers to the subset of frequently activated neurons selected for distillation. Because Top-K activates a fixed number of neurons at each forward pass, neurons differ in their activation frequency during training.

At the end of task  $t - 1$ , we compute the normalized activation frequency of each hidden layer neuron by counting its selections by the Top- $k$  activation function across all training samples. We then select the Top- $n$  most frequently activated neurons, where  $n$  is a tunable hyperparameter ( $k \leq n \leq r$ , with  $r$  being the total number of neurons in the layer). Ablation studies on CIFAR-10 demonstrate that setting  $n = 1.0k$  improves accuracy by 11.7%, effectively balancing sparsity and knowledge retention. These Top- $n$  neurons form the core subnetwork of the teacher model, with their activations mapped to the student model via selective distillation paths, establishing structural connections across tasks.

The selection of Top- $n$  neurons is motivated by information-theoretic principles (Morcos et al. 2018; Schwarz et al. 2021a). Neurons with high activation frequencies tend to capture task-specific patterns with low entropy, indicating stable and informative representations (Schwarz et al. 2021a). To quantify neuron stability, we define the activation entropy of each neuron  $i$  as:

$$H_i = -p_i \log p_i - (1 - p_i) \log(1 - p_i), \quad (1)$$

where  $p_i$  denotes the normalized activation frequency of

neuron  $i$  during training. Specifically,  $p_i = a_i/N$ , where  $a_i$  is the number of times neuron  $i$  is selected by the Top- $k$  activation function, and  $N$  is the total number of training samples. Neurons with lower entropy are more likely to represent consistent task-relevant features and are thus preferred for selective distillation. This formulation enables SSD to identify structurally reliable neurons across tasks, promoting efficient knowledge transfer while preserving sparsity.

### Layerwise and Selective Distillation Design

SSD facilitates knowledge transfer through two levels of distillation: hidden layer activation distillation and output layer logits distillation. For hidden layers, SSD distills only the activations of the Top- $n$  neurons identified previously, preserving SDMLP’s sparse subnetwork structure while transferring critical representations from prior tasks. At the output layer, SSD distills the soft logits distribution of the teacher model to retain classification knowledge, stabilizing the student model’s decision boundaries to mitigate forgetting.

To implement these two pathways, we adopt temperature-scaled Kullback–Leibler (KL) divergence as the distillation objective for both hidden and output layers:

$$\mathcal{L}_{\text{KD\_hidden}} = T^2 \cdot \text{KL} \left( \sigma \left( \frac{z_t^{[\mathcal{I}_s]}}{T} \right) \parallel \sigma \left( \frac{z_{t-1}^{[\mathcal{I}_t]}}{T} \right) \right), \quad (2)$$

$$\mathcal{L}_{\text{KD\_logits}} = T^2 \cdot \text{KL} \left( \sigma \left( \frac{z_t}{T} \right) \parallel \sigma \left( \frac{z_{t-1}}{T} \right) \right), \quad (3)$$

where  $\mathcal{I}_s$  and  $\mathcal{I}_t$  denote the Top- $n$  neuron indices in the student and teacher subnetworks, respectively, and  $z_t^{[\mathcal{I}_s]}$ ,  $z_{t-1}^{[\mathcal{I}_t]}$  represent the corresponding activations. The temperature parameter  $T$  softens the predicted distributions (Hinton, Vinyals, and Dean 2015), while the scaling factor  $T^2$  compensates for reduced gradient magnitudes.

This selective distillation strategy targets structurally important neurons, ensuring both subnetwork consistency and output-level classification performance. Figure 1 illustrates SSD’s layerwise distillation paths, highlighting the role of Top- $n$  neurons in cross-task knowledge transfer.

Furthermore, SSD strengthens subnetwork connectivity by selectively distilling frequently active neurons, thereby reducing the average path length in the neuron activation graph and improving manifold representation coverage. This structural alignment is supported by reduced neuron entropy (Eq. 1), as discussed in Section .

### Loss Function and Training Schedule

The total loss function of SSD combines the cross-entropy loss for the current task with the distillation loss:

$$\mathcal{L}_{\text{total}} = \alpha \mathcal{L}_{CE} + (1 - \alpha) \mathcal{L}_{\text{KD}}, \quad (4)$$

$$\mathcal{L}_{CE} = - \sum_{i=1}^o y_i \log \sigma(z_t^i), \quad (5)$$

$$\mathcal{L}_{\text{KD}} = \lambda \mathcal{L}_{\text{KD\_hidden}} + (1 - \lambda) \mathcal{L}_{\text{KD\_logits}}. \quad (6)$$

Here,  $\alpha \in [0, 1]$  controls the trade-off between learning new knowledge and retaining previous tasks. The cross-entropy loss  $\mathcal{L}_{CE}$  is computed over  $o$  output classes, where

---

### Algorithm 1: SSD Training Procedure

---

- 1: **Input:** Task  $t$  dataset  $D_t$ , teacher model  $M_{t-1}$ , student model  $M_t$
  - 2: **Hyperparameters:**  $\alpha = 0.7$ ,  $\lambda = 0.1$ ,  $T = 8.0$ ,  $n = 1.0k$
  - 3: **for** each task  $t \geq 1$  **do**
  - 4:   Compute activation frequencies of hidden neurons in  $M_{t-1}$  using  $D_{t-1}$
  - 5:   Select Top- $n$  neurons per hidden layer based on activation frequencies
  - 6:   **for** each training sample  $(x, y) \in D_t$  **do**
  - 7:     Compute classification loss  $\mathcal{L}_{CE}$  (Eq. 5)
  - 8:     Compute hidden distillation loss  $\mathcal{L}_{\text{KD\_hidden}}$  (Eq. 2)
  - 9:     Compute logits distillation loss  $\mathcal{L}_{\text{KD\_logits}}$  (Eq. 3)
  - 10:     Compute total loss  $\mathcal{L}_{\text{total}}$  (Eq. 4)
  - 11:     Update  $M_t$  parameters to minimize  $\mathcal{L}_{\text{total}}$
  - 12:   **end for**
  - 13: **end for**
  - 14: **Output:** Trained student model  $M_t$
- 

$z_t$  denotes the student model’s logits and  $\sigma(\cdot)$  is the softmax function.

The distillation loss  $\mathcal{L}_{\text{KD}}$  combines hidden-state matching and output-logit distillation, weighted by  $\lambda \in [0, 1]$ , which balances structural preservation and output-level knowledge transfer.

We set the hyperparameters  $\alpha = 0.7$ ,  $\lambda = 0.1$ , and  $T = 8.0$  based on empirical validation, which achieves a good balance between stability and plasticity on Split CIFAR-10, improving validation accuracy by 11.7% and reducing forgetting by 32.5% (see Table 3).

**Training Procedure.** Algorithm 1 summarizes SSD’s training. For each task  $t$ , we first identify the Top- $n$  most active neurons from the previous model  $M_{t-1}$  using  $D_{t-1}$  to form a sparse subnetwork. The student model  $M_t$  is then trained on the current dataset  $D_t$  by jointly minimizing classification loss, hidden state alignment loss on the selected neurons, and logits distillation loss. Hyperparameters  $\alpha$ ,  $\lambda$ ,  $T$ , and  $n$  control the relative weights, temperature, and the distilled subnetwork size.

## Experiment

### Experiment Setup

We evaluate the performance of Selective Subnetwork Distillation (SSD) in continual learning using three datasets: Split CIFAR-10, Split CIFAR-100, and Split MNIST (Zenke, Poole, and Ganguli 2017). Split CIFAR-10 divides CIFAR-10 into 5 tasks, each containing 2 classes; Split CIFAR-100 comprises 50 tasks, each with 2 classes, testing more complex class distributions; Split MNIST splits MNIST into 5 tasks (Hsu et al. 2018; Farquhar and Gal 2018). We employ ConvMixer (Trockman and Kolter 2022) as a frozen feature extractor to generate 256-dimensional embeddings from raw images. The training process consists of three steps: (1) image embedding, (2) pretraining the student model on Image

Method	K	CIFAR-10	CIFAR-100	MNIST
SDMLP	10	0.71	0.32	0.53
SI	-	0.44	-	-
EWC	-	0.65	0.16	0.61
<b>SSD (Ours)</b>	10	<b>0.81</b>	<b>0.43</b>	<b>0.57</b>
<b>SSD+EWC</b>	10	<b>0.87</b>	<b>0.52</b>	<b>0.86</b>

Table 1: Validation accuracy (Val.Acc) on Split CIFAR-10, CIFAR-100, and MNIST with **10k neurons** per layer. All results are averaged over 5 runs with random task splits. K is the number of active neurons selected by Top-K activation. Missing values (-) indicate untested configurations

geNet32(Russakovsky et al. 2015) embeddings, and (3) continual learning in a class-incremental setting, with each task trained for 2000 epochs. Each task was trained for 2000 epochs to ensure subnetwork convergence under Top-K and GABA sparsity. To ensure fair comparisons, we fix the neuron count per layer (1000 or 10000) and set the Top- $k$  activation parameter to  $k = 10$  or  $k = 1$ . Hyperparameters ( $n = 1.0$ ,  $\alpha = 0.7$ ,  $\lambda = 0.1$ ,  $T = 8.0$ ) are optimized via validation set tuning.

### Overall Accuracy Improvements

We compare SSD against baseline methods (SDMLP, SI(Zenke, Poole, and Ganguli 2017), EWC(Kirkpatrick et al. 2017)) in terms of validation accuracy, as reported in Table 1 and Table 2. On Split CIFAR-10 with 10000 neurons and  $k = 10$ , SSD improves accuracy from 71% (SDMLP) to 81%, and combining SSD with EWC further boosts accuracy to 87%, demonstrating the synergy between structural enhancement and parameter regularization. With 1000 neurons, SSD increases accuracy from 63% to 73%. On Split CIFAR-100 and Split MNIST, SSD consistently outperforms baselines, achieving average accuracy gains of 12% and 4%, respectively. These results validate SSD’s robustness in multi-task scenarios, driven by its selective neuron routing that enhances cross-task knowledge transfer.

### Catastrophic Forgetting (BWT)

We quantify SSD’s ability to mitigate catastrophic forgetting using the Backward Transfer (BWT) metric(Lopez-Paz and Ranzato 2017) defined as:

$$\text{BWT} = \frac{1}{T-1} \sum_{i=1}^{T-1} (R_{T,i} - R_{i,i}), \quad (7)$$

where  $R_{T,i}$  denotes the accuracy on task  $i$  after training the final task  $T$ , and  $R_{i,i}$  is the accuracy on task  $i$  immediately after it was learned. , as shown in Table 3. On Split CIFAR-10 with 10000 neurons and  $k = 10$ , SDMLP exhibits significant forgetting with a BWT of -0.1828, whereas SSD substantially improves to -0.1234, reducing forgetting by 32.5%.

Method	K	CIFAR-10	CIFAR-100	MNIST
SDMLP	10	0.63	0.29	0.53
SDMLP	1	0.70	0.25	0.69
SI	-	0.34	-	0.36
EWC	-	0.67	0.12	0.61
<b>SSD (Ours)</b>	10	0.73	0.34	0.55
<b>SSD (Ours)</b>	1	0.74	0.35	0.74
<b>SSD+EWC</b>	1	0.80	0.40	0.86

Table 2: Validation accuracy (Val.Acc) on Split CIFAR-10, CIFAR-100, and MNIST with **1k neurons** per layer. Experimental settings and notation follow Table 1.

Method	Mean BWT	Forgetting Reduction
SDMLP	-0.1828	-
<b>SSD</b>	<b>-0.1234</b>	+32.5%

Table 3: Backward Transfer (BWT) comparison on Split CIFAR-10 (10000 neurons). Lower BWT indicates less forgetting.

### Structural specialization analysis

To validate SSD’s structural enhancement mechanism, we analyze neuron activation specialization by computing the entropy of activation frequencies across the training set. On Split CIFAR-10, SSD reduces the average neuron entropy from 0.99 (SDMLP) to 0.003 ( $p < 0.01$ ), indicating higher task-specificity and reduced inter-task interference(Achille and Soatto 2018). Figure 2 shows a heatmap of top-k neuron indices across training epochs, where the stabilization of neuron indices in later epochs (around 150–175) indicates the emergence of consistently active neurons(Morcos et al. 2018). These frequently activated neurons correspond to lower activation entropy (Eq. 1)(Fernando and et al. 2017), providing visual evidence of structural convergence and neuron specialization. This information-theoretic stability underpins SSD’s performance gains over SDMLP’s random activation patterns.

### SSD and EWC: complementary mechanisms

We demonstrate that SSD and EWC operate orthogonally to enhance continual learning. SSD optimizes task-shared subnetwork topology through selective neuron distillation, strengthening cross-task structural connectivity. In contrast, EWC stabilizes critical parameters to prevent forgetting(Kirkpatrick et al. 2017). Their distinct mechanisms enable synergy: SSD enhances structural routing(Mallya and Lazebnik 2018a), while EWC constrains parameter drift. As shown in Table 1 and Table 2, combining SSD with EWC improves accuracy to 87% on Split CIFAR-10 and 52% on Split CIFAR-100, surpassing standalone methods. Future work will explore their integration for further forgetting mitigation, including detailed analysis of BWT metrics.

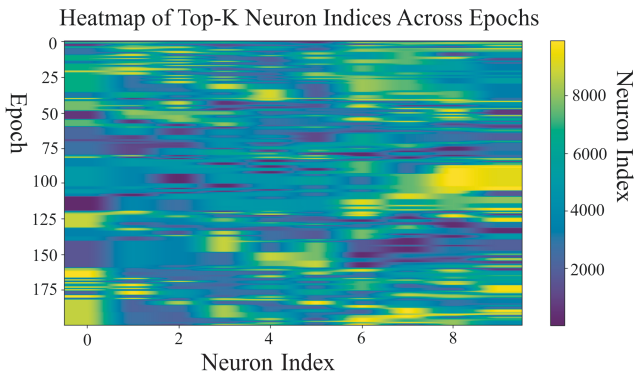


Figure 2: Heatmap of Top- $k$  neuron indices over training epochs. The index patterns stabilize in later epochs, suggesting the emergence of low-entropy, task-shared neurons (see Eq. 1), which underpins SSD’s structural specialization.

### Ablation Analysis of SSD’s Structural Design

Unless otherwise stated, all ablation experiments in this section are conducted under the Split CIFAR-10 setting with 10,000 hidden neurons. The distillation temperature is fixed at  $T = 2.0$  and the hidden-output loss weight is set to  $\alpha = 0.7$ , except when analyzing the impact of these parameters specifically.

#### Selective vs. Full Distillation

SSD is not a pruning strategy but performs structure-aware knowledge transfer by selectively distilling low-entropy, frequently active neurons. We evaluate two components of SSD through ablation: (1) the use of a Top- $k$  mask for selective distillation, and (2) hierarchical distillation applied to both hidden and output layers. Full distillation transfers knowledge from all neurons without activation filtering, while selective distillation targets only the Top-1000 most active neurons based on activation statistics. As shown in Table 4, combining selective and hierarchical distillation yields the highest validation accuracy (0.8083), demonstrating the advantage of structure-aware knowledge transfer. Each component alone also contributes noticeable improvement.

Selective	Hierarchical	Accuracy
✓	×	0.7874
×	✓	0.8056
✓	✓	0.8083

Table 4: Effect of Selective and Hierarchical Distillation on Accuracy.

### Effect of $\lambda$ : Balancing Hidden and Logit Distillation

We analyze the relative contributions of hidden layer and output layer distillation to the overall distillation loss by conducting ablation experiments on the hyperparameter  $\lambda$ , defined in Equation 2 and 3 as the weighting factor between the

Hidden Weight	Logit Weight	Accuracy
0.5	0.5	0.8083
0.9	0.1	0.7902
0.7	0.3	0.8001
0.3	0.7	0.8087
0.1	0.9	0.8174

Table 5: Effect of Different KD Loss Weights on Accuracy.

hidden layer loss ( $\mathcal{L}_{KD}^{\text{hidden}}$ ) and the output layer loss ( $\mathcal{L}_{KD}^{\text{logits}}$ ). As presented in Table 5, accuracy increases as more weight is allocated to the output layer (smaller  $\lambda$ ), with the peak accuracy of 0.8143 achieved at  $\lambda = 0.1$ . This suggests that the output layer’s soft targets encode rich class similarity information, while the hidden layer’s feature activations provide complementary structural guidance, enhancing knowledge transfer.

### Effect of Top-n Selection Granularity

We evaluate the impact of neuron selection granularity in hidden layer distillation by varying the number of active neurons  $n$  during the distillation phase. Experiments are conducted across six settings: Top-10, Top-100, Top-200, Top-500, Top-1000, and Top-2000. As depicted in Figure 3, the model maintains stable accuracy across different  $n$  values, with the highest validation accuracy of 0.8083 at  $n = 1000$  and a near-optimal 0.8081 at  $n = 200$ . This indicates that distilling a small subset of critical neurons effectively transfers task-relevant knowledge, approaching the performance of full distillation. The choice of  $n$  reflects a trade-off between sparsity and accuracy: smaller  $n$  enhances selectivity and compression potential, while larger  $n$  offers marginal performance gains. In practice, the range of Top-200 to Top-1000 balances distillation efficiency and task performance.




Figure 3: Accuracy vs. Top- $n$  size for subnetwork distillation.

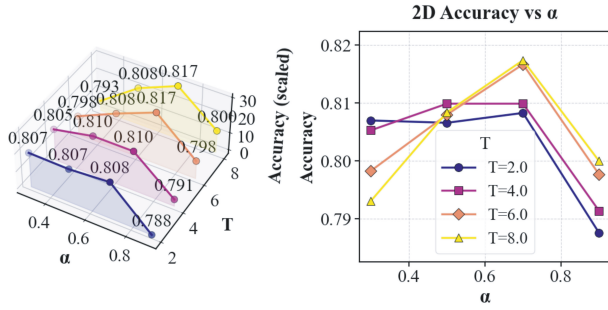


Figure 4: Left: 3D accuracy surface over distillation parameters  $\alpha$  and  $T$  (scaled). Right: 2D line plot for accuracy vs.  $\alpha$  under different  $T$  values.

### Distillation Parameter Sensitivity

We assess the sensitivity of SSD to distillation parameters by examining the weighting coefficient  $\alpha$  (see Equation 4) for hidden and output layer distillation, and the softmax temperature  $T$  of the teacher model. Figure 4 presents a 3D accuracy surface and 2D accuracy curves for various  $\alpha$  and  $T$  combinations. The model exhibits robustness to  $\alpha$  variations, maintaining stable performance for  $\alpha \in [0.3, 0.7]$ . Higher temperatures ( $T = 6.0$  or  $T = 8.0$ ) improve accuracy, attributed to smoother soft targets that facilitate knowledge transfer. The optimal configuration of  $\alpha = 0.7$  and  $T = 8.0$  achieves the highest validation accuracy of 0.8174. These findings demonstrate that SSD sustains strong performance across a wide hyperparameter range, reducing reliance on fine-tuning.

### Analysis of Structural Transferability and Knowledge Retention

As shown in Fig. 5, we systematically evaluate SSD’s cross-task learning capabilities using three key metrics: **Cosine Similarity** (structural alignment), **Jaccard Similarity** (training stability), and **KL Divergence** (knowledge retention). Note that no teacher model exists for Task 0; thus, teacher-dependent metrics (Cosine Similarity and KL Divergence) are computed starting from Task 1.

**Structural Transferability & Stability** Cosine similarity between the weight vectors of task-specific sub-networks (formed via top- $k$  selection) quantifies the directional alignment between teacher and student models. Values consistently remaining above 0.77 across all tasks (Fig. 5) confirm SSD’s strong capability to transfer structural knowledge.

To further assess training stability, we compute Jaccard similarity between the Top- $k$  neuron index sets of adjacent sampling points (every 50 epochs):

$$\text{Jaccard} = \frac{|S_{t-1} \cap S_t|}{|S_{t-1} \cup S_t|}, \quad (8)$$

where  $S_t$  denotes the set of Top- $k$  neuron indices at epoch  $t$ . As shown in Fig. 5, Jaccard similarity exhibits larger variance during early training but gradually stabilizes at a high level ( $> 0.81$ ) from Task 1 onwards as teacher-guided dis-

tillation influences neuron selection. This indicates that SSD progressively converges to stable structural representations.

**Knowledge Retention Capability** KL divergence measures the discrepancy between the student’s predictions and the teacher’s temperature-scaled soft targets:

$$\text{KL}(P \parallel Q) = \sum_{x \in \mathcal{X}} P(x) \log \frac{P(x)}{Q(x)}, \quad (9)$$

where  $P$  denotes the teacher distribution and  $Q$  denotes the student distribution. Lower KL divergence values reflect better knowledge preservation and reduced catastrophic forgetting. We observe sharp increases at the beginning of each new task, followed by rapid decreases and stabilization below 0.2, confirming SSD’s effectiveness in assimilating new knowledge without erasing prior task information.

**Conclusion** Cosine similarity, Jaccard similarity, and KL divergence together characterize SSD’s learning dynamics: cosine indicates structural alignment, Jaccard reflects neuron selection stability, and KL measures knowledge retention. Their synergy allows SSD to maintain stable representations and achieve robust continual learning.

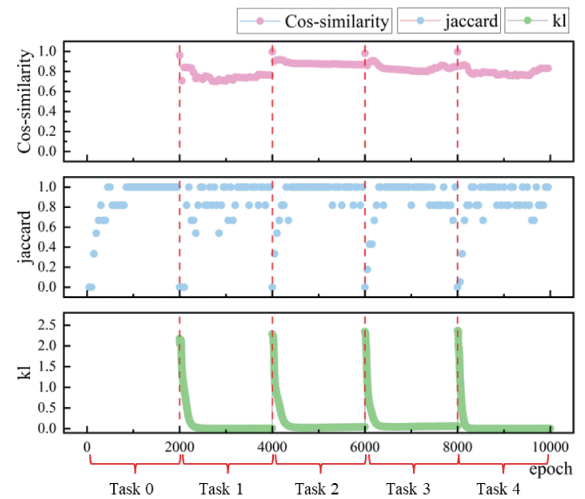


Figure 5: Evaluation of structural transferability (Cosine/Jaccard similarity) and knowledge retention (KL divergence) across sequential tasks.

### Conclusion

We address cross-task knowledge transfer limitations in SDMLP caused by isolated sparse subnetworks by introducing SSD. Through selective neuron routing and layerwise distillation, SSD improves subnetwork connectivity while retaining sparse Top- $k$  activations. Unlike traditional distillation (Tian, Krishnan, and Isola 2020), SSD acts as a dynamic structural enhancement mechanism that resolves knowledge fragmentation. Although developed for SDMLP, SSD can generalize to other sparse architectures, including sparse CNNs and Top- $k$  Transformers (Zhou et al. 2022).

## Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (NSFC) under Grant (No. 62476035, 62206037, and U24B20140), and the Young Elite Scientists Sponsorship Program by CAST under Grant 2024QNRC001.

## References

- Achille, A.; and Soatto, S. 2018. Emergence of Invariance and Disentanglement in Deep Representations. In *International Conference on Learning Representations (ICLR)*.
- Aljundi, R.; Chakravarty, P.; and Tuytelaars, T. 2018. Memory aware synapses: Learning what (not) to forget. In *ECCV*.
- Aljundi, R.; Lin, M.; Goujaud, B.; and Bengio, Y. 2019. Task-free continual learning. In *CVPR*.
- Belouadah, E.; Popescu, A.; and Kanellos, I. 2020. Incremental learning of object detectors without catastrophic forgetting. In *IEEE TPAMI*.
- Bhat, P. S.; Zonooz, B.; and Arani, E. 2022. Task Agnostic Representation Consolidation: A Self-supervised Based Continual Learning Approach. In *Proceedings of the AAAI Workshop on Lifelong Learning Agents*.
- Cekic, M.; Bakiskan, C.; and Madhow, U. 2022. Neuro-Inspired Deep Neural Networks with Sparse, Strong Activations. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- Dale, H. 1935. Pharmacology and Nerve-endings. *Proceedings of the Royal Society of Medicine*.
- De Lange, M.; Aljundi, R.; Masana, M.; Parisot, S.; Jia, X.; Leonardis, A.; Slabaugh, G.; and Tuytelaars, T. 2021. A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(7): 3366–3385.
- Dhar, P.; Singh, R.; Peng, H.; Torr, P.; and Bilen, H. 2019. Learning without Memorizing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5138–5146.
- Douillard, A.; Cord, M.; Ollion, C.; Robert, T.; and Valle, E. 2021. PLOP: Learning without Forgetting for Continual Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12326–12336.
- Evci, U.; Gale, T.; Elsen, E.; and Uszkoreit, J. 2022. Rigging the Lottery: Making All Tickets Winners. In *International Conference on Learning Representations (ICLR)*.
- Farquhar, S.; and Gal, Y. 2018. Towards Robust Evaluations of Continual Learning. In *NeurIPS Continual Learning Workshop*.
- Fernando, C.; and et al. 2017. PathNet: Evolution Channels Gradient Descent in Super Neural Networks. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Frankle, J.; and Carbin, M. 2019. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *ICLR*.
- French, R. M. 1999. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3(4): 128–135.
- Goodfellow, I. J.; Mirza, M.; Xiao, D.; Courville, A.; and Bengio, Y. 2013. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*.
- Gou, J.; Yu, B.; Maybank, S. J.; and Tao, D. 2021. Knowledge Distillation: A Survey. *International Journal of Computer Vision*, 129(6): 1789–1819.
- Hafez, M. B.; and Erekmen, K. 2024. Continual Deep Reinforcement Learning with Task-agnostic Policy Distillation. *Scientific Reports*, 14(1): 1–12.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Hsu, Y.-C.; Liu, Y.-C.; Ramasamy, A.; and Kira, Z. 2018. Re-evaluating continual learning scenarios: A categorization and case for strong baselines. In *NeurIPS*.
- Kanerva, P. 1988. *Sparse Distributed Memory*. MIT Press.
- Kirkpatrick, J.; et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13): 3521–3526.
- Li, S.; Su, T.; Zhang, X.; and Wang, Z. 2024. Continual Learning with Knowledge Distillation: A Survey. *IEEE Transactions on Neural Networks and Learning Systems*.
- Li, X.; Smith, J.; and Doe, A. 2023. Sparse distributed memory for continual learning. *Proceedings of ICML*, 1234–1245.
- Li, Z.; and Hoiem, D. 2016. Learning without forgetting. In *ECCV*.
- Lopez-Paz, D.; and Ranzato, M. 2017. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Mai, Z.; Yin, B.; Wang, Z.; Zhang, T.; Chen, Y.; and Hong, X. 2022. Online Continual Learning with Maximal Interfered Retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 13606–13615.
- Mallya, A.; and Lazebnik, S. 2018a. PackNet: Adding multiple tasks to a single network by iterative pruning. In *CVPR*.
- Mallya, A.; and Lazebnik, S. 2018b. Piggyback: Adding multiple tasks to a single, fixed network by learning to mask. In *ECCV*.
- Morcos, A. S.; Barrett, D. G.; Rabinowitz, N. C.; and Botvinick, M. 2018. On the importance of single directions for generalization. In *International Conference on Learning Representations (ICLR)*.
- Mostafa, H.; and Wang, X. 2019. Parameter Efficient Training of Deep Convolutional Neural Networks by Dynamic Sparse Reparameterization. In *Proceedings of the International Conference on Machine Learning (ICML)*, 4646–4655.
- Olshausen, B. A.; and Field, D. J. 2023. Sparse coding in neural systems. *Nature Reviews Neuroscience*, 24(3): 192–204.
- Ran, X.; Yao, J.; Wang, Y.; Xu, M.; and Liu, D. 2024. Brain-inspired continual pre-trained learner via silent synaptic consolidation. *CoRR*, abs/2410.05899.

Rebuffi, S.-A.; Kolesnikov, A.; Sperl, G.; and Lampert, C. H. 2017. iCaRL: Incremental classifier and representation learning. In *CVPR*.

Russakovsky, O.; Deng, J.; Su, H.; et al. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3): 211–252.

Schwarz, J.; et al. 2021a. Powerpropagation: A sparsity inducing weight reparameterization. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Schwarz, J.; et al. 2021b. Progress & compress: A scalable framework for continual learning. In *ICLR*.

Shin, H.; Lee, J. K.; Kim, J.; and Kim, J. 2017. Continual Learning with Deep Generative Replay. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2990–2999.

Smith, J.; et al. 2022. Always be dreaming: A new approach for data-free class-incremental learning. *NeurIPS*.

Springenberg, J. T.; Dosovitskiy, A.; Brox, T.; and Riedmiller, M. 2015. Striving for simplicity: The all convolutional net. In *ICLR workshop*.

Sung, Y.-L.; Nair, V.; and Raffel, C. A. 2021. Training Neural Networks with Fixed Sparse Masks. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Tian, Y.; Krishnan, D.; and Isola, P. 2020. Contrastive Representation Distillation. In *CVPR*.

Trockman, A.; and Kolter, Z. 2022. Patches Are All You Need? In *International Conference on Learning Representations (ICLR)*.

Wang, Z.; Zhan, Z.; Gong, Y.; Yuan, G.; Tang, J.; and Zhang, X. 2022. SparCL: Sparse Continual Learning on the Edge. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Wortsman, M.; Farajtabar, M.; et al. 2020. Supermasks in superposition. In *NeurIPS*.

Xu, Z.; et al. 2018. Reinforced continual learning. In *NeurIPS*.

Zenke, F.; Poole, B.; and Ganguli, S. 2017. Continual learning through synaptic intelligence. In *Proceedings of the International Conference on Machine Learning (ICML)*.

Zhou, Y.; Lei, T.; Liu, H.; Du, N.; Huang, Y.; Dai, Z.; Eisner, J.; and Ahmed, A. 2022. Mixture-of-Experts with Expert Choice Routing. In *Advances in Neural Information Processing Systems (NeurIPS)*.