SLEEP-INSPIRED TWO-STAGE CONTINUAL LEARNING FOR SPIKING NEURAL NETWORKS

ABSTRACT

Continual learning systems are often plagued by catastrophic forgetting across sequential tasks, whereas the human brain retains prior knowledge through neural plasticity and memory consolidation. Spiking neural networks (SNNs), which emulate the brain's communication and neuronal dynamics, provide a promising substrate for brain-inspired continual learning. However, a pronounced stability and plasticity dilemma is still encountered in continual learning with SNNs. Neuronal dynamics are altered by incoming tasks, while current and historical samples are usually treated symmetrically by conventional replay. To overcome these limitations, a sleep inspired framework is proposed. Rapid daytime adaptation is separated from nighttime knowledge consolidation. Three complementary mechanisms are integrated, including sample-wise coefficients derived from membrane potential and spike statistics, source-dependent optimization for current and replay samples, and class-balanced consolidation via old class response distillation. Experiments on MNIST, DVS128 Gesture, CIFAR-10, and CIFAR-100 under class incremental and task incremental settings show consistent improvements in final accuracy over experience replay. These results demonstrate that temporally separating acquisition and consolidation can effectively coordinate stability and plasticity in SNN continual learning.

1 INTRODUCTION

Continual learning is adopted when complete retraining is rendered impractical by limited storage and computational resources, so that new knowledge can be acquired while earlier capabilities are retained. This process is primarily hindered by catastrophic forgetting, whereby representations learned from previous data are overwritten during subsequent optimization (Mccloskey & Cohen, 1989). Continual learning is commonly evaluated under class-incremental and task-incremental settings. In the former, predictions must be made over all observed classes, whereas task identity is provided in the latter (van de Ven et al., 2022).

Spiking neural networks (SNNs) are regarded as a suitable computational framework for continual learning under such resource constraints. Information is encoded through membrane integration and discrete spike emission across multiple time steps (Minhas et al., 2025a). The representation of spatial intensity and temporal dynamics through sparse spiking computation has also been demonstrated by recent neuron designs (Fan et al., 2025). During continual learning, these spatiotemporal representations must be retained across successive tasks, yet their stability cannot be guaranteed by computational efficiency alone. Membrane trajectories, firing patterns, and class responses established for earlier tasks can still be altered by sequential parameter updates, so previously learned behavior must be protected against catastrophic forgetting during continued adaptation.

Preservation of learned behavior is pursued through several complementary strategies. Important parameters can be constrained, historical samples can be replayed, previous responses can be distilled, and task-specific capacity can be allocated. Direct access to earlier experience is provided by replay, but its effectiveness can be limited by the imbalance between abundant current samples and a small historical memory. In class-incremental learning, the effective training distribution can become dominated by current classes, and the decision regions of earlier classes can be compressed by this imbalance. When a common optimization schedule is applied to both sources, the plasticity needed to acquire new knowledge and the stability needed to preserve historical knowledge are not explicitly distinguished.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

In SNNs, this coordination problem also involves neuronal dynamics because knowledge is encoded in both synaptic parameters and neuronal trajectories. The response of a spiking classifier is characterized by neuronal-state statistics; earlier classes are directly anchored through replay; and functional targets over the old output space are supplied by a frozen model. These signals serve different purposes and need not be assigned identical roles throughout training. A staged schedule in which rapid adaptation is followed by balanced consolidation is motivated by their complementary functions.

A sleep-inspired two-stage framework is proposed to implement this separation. During daytime adaptation, sample-wise coefficients are derived by a lightweight controller from neuronal-state statistics and a surprise score. The time-averaged membrane-potential mean, membrane-potential standard deviation, and spike rate are recorded for each sample over a 12-step forward pass. These statistics are detached from the model computation graph, and surprise is measured through cosine distance relative to an exponential moving average reference. The resulting coefficients are bounded between 0.1 and 0.9. Their minibatch mean is used only to balance a coefficient budget against preservation of old classifier rows during daytime optimization, without rescaling forward spike outputs.

The separation between adaptation and consolidation also governs the optimization of current and historical examples. During early daytime epochs, Mixup is restricted to current samples and combined with label smoothing, whereas replay samples remain unmixed and are optimized with ordinary cross-entropy. During later daytime epochs, the two sources are concatenated for shared label-smoothed refinement. Old-class response distillation from the frozen model is retained throughout daytime learning whenever a teacher is available. After daytime adaptation, current examples are used to update the class-indexed memory. Approximately class-balanced hard-label replay and old-class response distillation are then used during nighttime consolidation. Controller-dependent penalties are disabled during this stage, and the consolidated model, updated memory, and neuronal-state reference are transferred to the next task.

The main contributions are summarized below.

1. A sleep-inspired two-stage schedule is formulated for continual learning in SNNs. Rapid daytime adaptation and nighttime class-balanced consolidation are temporally separated.
2. A surprise-modulated controller is introduced to derive sample-wise coefficients from detached neuronal statistics. Their minibatch mean is used to balance a coefficient budget against preservation of old classifier weights.
3. A phase-dependent replay objective is developed. Current-only Mixup and unmixed replay are used during early daytime training, followed by shared refinement during later epochs. Old-class response distillation is retained during both daytime adaptation and nighttime consolidation.

2 Related Works

Different forms of historical information are preserved by existing continual learning methods during sequential optimization. Important parameters are constrained through a Fisher information approximation in Elastic Weight Consolidation (Kirkpatrick et al., 2017), while previous output responses are retained through distillation in Learning without Forgetting (Li & Hoiem, 2017). Parameter importance is estimated from contributions accumulated along the optimization trajectory in Synaptic Intelligence (Zenke et al., 2017) and from output sensitivity in Memory Aware Synapses (Aljundi et al., 2018). Without direct access to historical inputs, old-class decision regions are protected only indirectly, and their separation from newly introduced classes cannot be guaranteed by parameter or response constraints alone. More direct access to earlier experience is provided by replay-based methods. Update directions are constrained using episodic samples in Gradient Episodic Memory (Lopez-Paz & Ranzato, 2017), and exemplar rehearsal is combined with response preservation in iCaRL (Rebuffi et al., 2017). Incoming and replayed samples are treated asymmetrically in ER-ACE to reduce abrupt representation changes (Caccia et al., 2022), while historical inputs are reconstructed in Deep Generative Replay (Shin et al., 2017). Even with replay, historical coverage is restricted by memory capacity or reconstruction quality, and balanced optimization across old and new classes is not ensured merely by the availability of past experience. The retention of historical information must therefore be complemented by explicit coordination of how current and replayed samples are optimized throughout training. In the proposed framework, this coordination is addressed through source-dependent daytime adaptation followed by a separate phase of approximately class-balanced consolidation.

Within SNNs, knowledge preservation has also been pursued through modifications to network structure and neuronal activity. Structural growth, pruning, and reuse are supported by DSD-SNN (Han et al., 2023), while activity subspaces are estimated and update-related traces are projected onto an orthogonal complement in HLOP-SNN (Xiao et al., 2024). Sparse spatial and temporal activation is encouraged through trace-based selection and adaptive thresholds in SA-SNN (Shen et al., 2024). Structural adaptation is further supported through adaptive logit alignment in ALADE-SNN (Ni et al., 2025) and context-guided neuronal reuse and expansion in SCA-SNN (Han et al., 2025). Short-term Hebbian adaptation and long-term meta-learning are coordinated in HLML-SNN (Xu et al., 2026), while replay storage and temporal costs are reduced through latent compression and timestep reduction in Replay4NCL (Minhas et al., 2025b). In the structural and activity-control approaches above, preservation is coupled to capacity management, subspace estimation, or explicit regulation of neuronal participation. Additional structural management or direct intervention in neuronal computation is consequently required by the corresponding preservation mechanisms. A diagnostic role is instead assigned to neuronal state in the proposed framework, so preservation strength can be adjusted without these interventions. Sample-wise coefficients are computed from detached membrane-potential means, membrane-potential dispersion, and spike rates, together with surprise measured relative to an exponential moving average reference. Only their minibatch mean is used to balance a coefficient penalty against preservation of old classifier weights. The remaining predictive parameters are not directly constrained by this regularizer, and no task-specific structural allocation, update projection, or forward-output rescaling is introduced by the controller.

Neuronal-state regularization is complemented by replay and response distillation because historical inputs and old-class response relationships cannot be reconstructed from compact state statistics alone. Input and label anchors are supplied by stored exemplars, while response distributions over previously learned classes are supplied by a frozen teacher. The use of these complementary signals must also be coordinated across learning stages. Improved retention has been reported when optimal stopping was combined with sleep replay consolidation (Vins et al., 2025). Alongside this temporal separation, the data distribution and preservation objective within each stage must be explicitly specified. In the proposed framework, distinct supervised objectives are assigned to daytime adaptation and nighttime consolidation. During early daytime adaptation, Mix-up and label smoothing are applied only to current samples, while unmixed replay samples are optimized with hard labels; shared label-smoothed refinement is subsequently performed. During nighttime consolidation, approximately class-balanced hard-label replay is combined with old-class response distillation, and controller-dependent penalties are disabled. Old-class distillation is retained in both stages, with teacher responses computed from the frozen model consolidated after the preceding task. A common consolidated reference is thus provided for all distillation inputs, whereas logits collected at different points along the optimization trajectory are retained in Dark Experience Replay (Buzzega et al., 2020). Distillation is restricted to the old-class output subspace, so teacher targets are not imposed on new-class outputs. The neuronal-state reference is updated only after consolidation, so subsequent deviations are measured against a consolidated activity state.

3 METHODOLOGY

The sleep-inspired two-stage framework is presented in this section. It is composed of a daytime adaptation stage, a neuronal-state-conditioned controller, and a nighttime class-balanced consolidation stage. Rapid acquisition and the preservation of historical knowledge are coordinated through the interaction of these components. The overall workflow is illustrated in Figure 1.

3.1 SLEEP-INSPIRED TWO-STAGE CONTINUAL LEARNING FRAMEWORK

Class-incremental learning is considered, in which a labeled set of new classes is introduced at task k and the classifier is expanded to cover all classes observed to date. The trainable model is composed of a feature extractor, a spiking classification head, and a lightweight controller. Real examples are retained in a class-indexed episodic memory with a fixed total capacity of 500 samples, consistent with experience replay (Rolnick et al., 2019). Before new output rows are added, the model consolidated after task k-1 is copied and frozen to serve as the teacher. Previous responses are supplied by this model rather than stored with each exemplar, as in Dark Experience Replay (Buzzega et al., 2020). Its outputs are consequently aligned with the old-class subspace of the expanded student.

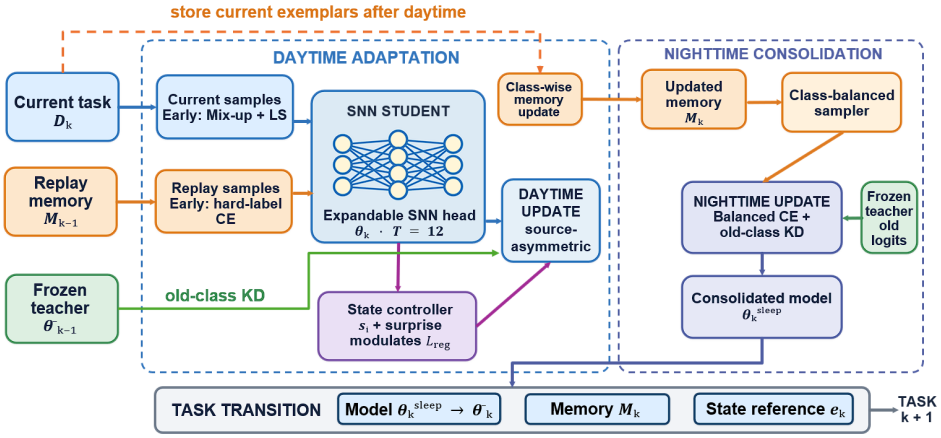

Figure 1: Overall workflow of the proposed sleep-inspired two-stage continual learning framework.

For each input, the feature extractor is evaluated once, and the resulting feature vector is supplied as a constant input current to the leaky integrate-and-fire classification head. Twelve simulation steps are used. The accumulated classification score is calculated as follows:

$$z_i = \sum_{\tau=1}^{12} o_{\tau,i} \tag{1}$$

where $o_{\tau,i}$ denotes the output spike vector of sample i at simulation step τ, with τ ranging from 1 to 12, and z_i denotes the corresponding accumulated classification score. The nondifferentiable spike function is optimized using an ATan surrogate derivative in accordance with surrogate-gradient learning (Neftci et al., 2019). The same temporally accumulated score is used for classification and distillation so that the two objectives are not optimized over inconsistent output representations. Because the feature extractor is evaluated once and only the spiking head is unfolded, the temporal overhead is confined to the component in which neuronal dynamics are explicitly modeled.

Learning is performed in two stages. This temporal separation is motivated by sleep-based consolidation in spiking networks (Golden et al., 2022) and complementary learning systems (Mcclelland et al., 1995). During daytime adaptation, current-task minibatches and approximately class-balanced replay samples are used. Source asymmetry is retained during the early epochs: Mixup and label smoothing are applied only to current examples, whereas replay examples remain unmixed and are optimized using ordinary cross-entropy loss. During the later epochs, current and replay examples are concatenated and optimized using a shared label-smoothed objective. After current examples have been used to update the memory, only the updated buffer is revisited during nighttime consolidation, with class participation kept approximately balanced. The consolidated network and neuronal-state reference are then transferred to the next task.

3.2 Neuron-State-Conditioned Adaptive Gating

Each sample is characterized by the membrane-potential mean, membrane-potential dispersion, and spike rate measured from the spiking classifier. Each statistic is first evaluated across classifier units at every simulation step and then averaged over the 12 steps. The detached neuronal-state vector is calculated as follows:

$$s_i = \mathrm{sg}\left([\mu_i^u, \sigma_i^u, r_i]^\top\right) \tag{2}$$

where s_i denotes the neuronal-state vector for sample i; μ_i^u and σ_i^u denote the time-averaged membrane-potential mean and the across-unit standard deviation, respectively; r_i denotes the time-averaged active-output ratio; the superscript u indicates aggregation over classifier units; and $\mathrm{sg}(\cdot)$ denotes the stop-gradient operator. The complete 12-step membrane trajectory is not retained; instead, the three complementary statistics are used to provide compact descriptions of response level, population dispersion, and spiking activity. The state is thereby kept diagnostic rather than

interventionist, and preservation strength can be adjusted by the controller without creating an auxiliary gradient path through the membrane trajectory, spiking classifier, or feature extractor.

The deviation of a current response from the consolidated activity state is measured relative to the exponential moving-average reference from the preceding task. The normalized surprise score is calculated as follows:

$$\hat{\delta}_i = \mathrm{Norm}_B(1 - \cos(s_i, e_{k-1})) \tag{3}$$

where $\hat{\delta}_i$ denotes the normalized surprise score, e_{k-1} denotes the neuronal-state reference associated with the preceding consolidated model, $\cos(\cdot, \cdot)$ denotes cosine similarity, and Norm_B denotes minibatch min-max normalization with a numerical stabilizer. The unnormalized distance is retained for a single-sample batch, and a neutral value of 0.5 is used before the first reference has been initialized. Cosine distance is used to emphasize changes in the relative neuronal-state pattern, whereas task-dependent response scales are mapped to a common gating range by minibatch normalization.

The detached state vector is mapped to a base logit by a multilayer perceptron with layer dimensions of 3, 16, 8, and 1. Layer normalization and ReLU activation are applied after the first layer, and ReLU activation is applied after the second layer. The sample-wise adaptive coefficient is calculated as follows:

$$g_i = 0.1 + 0.8 \ \mathrm{sigmoid}\Big(f_\phi(s_i) + 0.5\hat{\delta}_i\Big) \tag{4}$$

where g_i denotes the adaptive coefficient, f_ϕ denotes the controller, $\mathrm{sigmoid}(\cdot)$ denotes the sigmoid function, and $\hat{\delta}_i$ denotes the normalized surprise score. The affine sigmoid constrains g_i to [0.1, 0.9], thereby preserving nonzero capacity for both plasticity and retention rather than allowing the regularizer to be fully disabled or become fully dominant. Neuronal novelty is allowed to exert a direct influence through the explicit surprise path even before the lightweight controller has been fully calibrated. Because only the minibatch mean enters the regularizer, optimization pressure is changed without rescaling spike outputs or introducing a test-time gating branch. The regularization loss is calculated as follows:

$$L_{\mathrm{reg}} = 0.005\bar{g} + 0.05(1 - \bar{g}) \left\| W_k^{\mathrm{old}} - \bar{W}_{k-1} \right\|_F^2 \tag{5}$$

where $\bar{g}$ denotes the minibatch mean of g_i, W_k^{old} contains the current classifier rows associated with previously learned classes, and $\bar{W}_{k-1}$ denotes the corresponding frozen post-consolidation snapshot. The subscript F denotes the Frobenius norm, and the outer exponent 2 indicates that the norm is squared. In this formulation, $\bar{g}$ is used as an adaptive plasticity budget: larger values relax anchoring of old-class rows but incur a linear coefficient cost, whereas smaller values strengthen retention. Because the penalty is restricted to old-class decision weights, the classifier bias, new output rows, feature extractor, and remaining spiking parameters remain available for adaptation, thereby avoiding the rigidity of a global parameter constraint. Both regularization terms are disabled during nighttime consolidation.

3.3 Source-Aware Asymmetric Replay and Class-Balanced Consolidation

During the first half of daytime training, current examples are processed using Mixup with a Beta-distribution parameter of 0.2 and label smoothing with a coefficient of 0.1 (Zhang et al., 2018; Szegedy et al., 2016). Replay examples remain unchanged and are optimized using ordinary hard-label cross-entropy loss. After the source-specific losses have been averaged within their respective minibatches, the early-stage classification loss is calculated as follows:

$$L_{\mathrm{cls}}^{\mathrm{early}} = \frac{1}{2} L_{\mathrm{mix}}^{\mathrm{cur}} + \frac{1}{2} L_{\mathrm{CE}}^{\mathrm{rep}} \tag{6}$$

where L_{mix}^{cur} denotes the current-sample Mixup loss and L_{CE}^{rep} denotes the replay hard-label cross-entropy loss. When replay is unavailable during the first task, only the current-sample Mixup loss is used. The early-stage asymmetry is intentional: the support of newly introduced classes is broadened

by augmentation, whereas each scarce replay exemplar is kept unmixed as a stable anchor for an old-class decision region. During the second half of daytime training, Mixup is disabled and the two sources are concatenated. The late-stage classification loss is calculated as follows:

$$L_{\text{cls}}^{\text{late}} = L_{\text{LS}}(B^{\text{cur}} \cup B^{\text{rep}}) \tag{7}$$

where L_{LS} denotes the label-smoothed cross-entropy loss, B^{cur} denotes the current-task minibatch, and B^{rep} denotes the replay minibatch. Once initial acquisition has been completed, the source distinction is removed by this shared late-stage objective, and the expanded classifier is recalibrated over the combined visible class space without modifying the output space or the frozen teacher.

For each task after the first, old-class responses on the combined current-task and replay inputs are constrained by the frozen teacher. The teacher and student distributions are obtained by applying softmax with a temperature of 2 to the corresponding old-class outputs. The distillation loss is calculated as follows:

$$L_{\text{KD}} = \frac{T_d^2}{C_{k-1}} H_B(q, p) \tag{8}$$

where L_{KD} denotes the temperature-scaled distillation loss, T_d denotes the distillation temperature, C_{k-1} denotes the number of previously learned classes, and H_B denotes the batch-averaged soft cross-entropy between the teacher distribution q and the student distribution p. The gradient attenuation introduced by temperature softening is compensated for by the factor T_d^2, whereas division by C_{k-1} prevents the distillation term from increasing solely because additional old classes have accumulated. Alignment is preserved by restricting both distributions to old-class outputs, and teacher targets are not imposed on unseen new classes. The loss is omitted for the first task because a frozen teacher is unavailable. The complete daytime loss is calculated as follows:

$$L_{\text{day}} = L_{\text{cls}} + L_{\text{KD}} + L_{\text{reg}} \tag{9}$$

where L_{cls} is selected from Eq. (6) or Eq. (7) according to the daytime phase, L_{KD} preserves old-class responses in function space, and L_{reg} anchors old decision weights in parameter space. Complementary information is preserved rather than duplicated by these terms: the current task is acquired through label supervision, historical behavior on observed inputs is maintained through distillation, and old-class decision geometry is stabilized by selective weight anchoring. After daytime adaptation, current examples are used to update the class-indexed episodic memory. A subset of current and historical examples is retained in the class-indexed episodic memory, whose fixed total capacity is 500 samples shared across all observed classes. An approximately equal replay quota is assigned to each represented class, and sampling is performed without replacement within each class.

During nighttime consolidation, only the updated, approximately class-balanced memory is used. Mixup, label smoothing, and controller-dependent penalties are removed. Hard-label classification and old-class distillation are retained. The nighttime loss is calculated as follows:

$$L_{\text{night}} = L_{\text{CE}}^{\text{bal}} + L_{\text{KD}}^{\text{bal}} \tag{10}$$

where L_{CE}^{bal} denotes the hard-label classification loss on the class-balanced memory, and L_{KD}^{bal} denotes the corresponding old-class distillation loss. Nighttime consolidation is deliberately performed without Mixup, label smoothing, or controller-dependent penalties: once acquisition is complete, class geometry is restored by approximately class-balanced hard-label replay, while historical responses are stabilized by old-class distillation. After consolidation, the mean detached state vector is computed from a class-balanced memory sample. The reference is initialized with this mean and subsequently updated using momentum 0.9. Before the classifier is expanded for the next task, the post-nighttime model is copied and frozen as the teacher. The predictive parameters, class-indexed memory, and neuronal-state reference are therefore transferred together.

4 EXPERIMENT

4.1 EXPERIMENTAL SETTINGS

Evaluation is performed on MNIST, DVS128 Gesture, CIFAR-10, and CIFAR-100. MNIST and CIFAR-10 are divided into five sequential tasks, with two classes assigned to each task. CIFAR-100 is divided into ten tasks, with ten classes assigned to each task. DVS128 Gesture is divided into three tasks containing four, four, and three gesture classes, respectively. Under Class-IL, predictions are made over all classes observed up to the current task without access to task identity at inference time. Under Task-IL, task identity is provided, and the output corresponding to the identified task is used.

Dataset-dependent architectures are used for a practical comparison. For fine-tuning, joint training, and experience replay, a fully connected artificial neural network is used for MNIST, a ResNet-18-based artificial neural network is used for CIFAR-10 and CIFAR-100, and a convolutional SNN is used for DVS128 Gesture. For the proposed method, an improved fully connected SNN is used for MNIST. For CIFAR-10 and CIFAR-100, a ResNet-18-based spiking architecture incorporating the proposed daytime and nighttime mechanisms is used. An improved convolutional SNN is used for DVS128 Gesture. Twelve simulation steps are used by all proposed spiking models.

For Class-IL, 10, 40, 50, and 45 daytime epochs are used for each task on MNIST, DVS128 Gesture, CIFAR-10, and CIFAR-100, respectively. For Task-IL, 12, 25, 45, and 45 epochs are used in the same dataset order. A batch size of 128 is used in the CIFAR implementations. Adam optimization is performed with an initial learning rate of 0.001 and a weight decay of 0.0002. A class-indexed memory is used to store at most 500 exemplars in total across all observed classes. On CIFAR-10 and CIFAR-100, ten nighttime consolidation epochs are performed after each daytime phase using approximately class-balanced memory replay. The same task order and preprocessing protocol are used for all methods within each dataset.

4.2 CLASS-INCREMENTAL LEARNING RESULTS

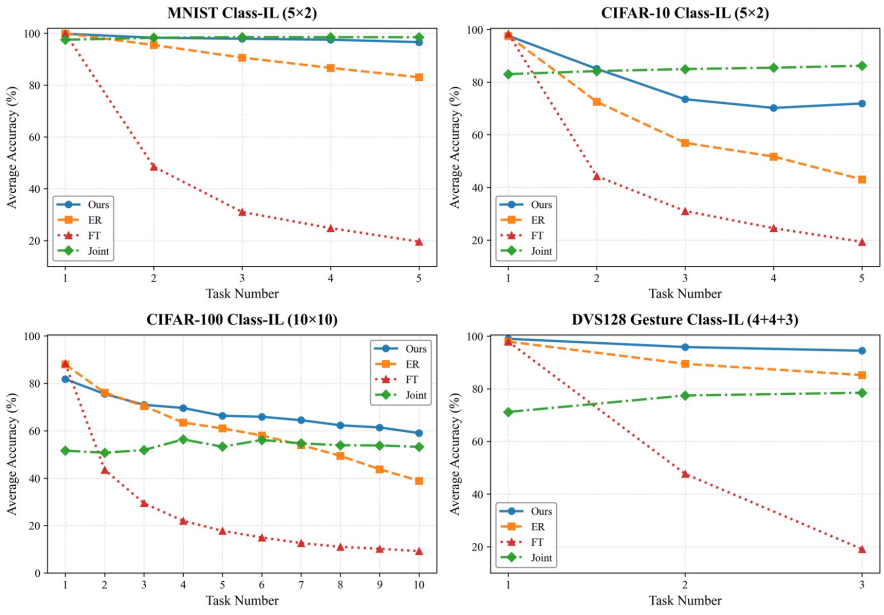

Figure 2: Average accuracy under Class-IL on four benchmarks.

The evolution of average accuracy is shown in Figure 2 as the visible class space is expanded. Joint training is retained only as an offline reference. All classes are simultaneously accessible to this model, and differences in architecture and optimization are also present. Joint training is therefore not treated as a strict upper bound. The primary evidence is provided by the matched sequential comparison among fine-tuning, experience replay, and the proposed method.

Two consistent patterns are observed. First, the advantage over experience replay is enlarged as more tasks are introduced. On MNIST, a reduction of only 3.21 percentage points is observed for the proposed method from the first task to the final task, whereas a reduction of 16.92 points is observed for experience replay. A final margin of 13.57 points is consequently obtained. On DVS128 Gesture, the margin is increased from 1.04 points to 9.26 points. The largest separation is obtained on CIFAR-10. Nearly identical first-task accuracies of 97.60 percent and 97.55 percent are obtained, whereas final accuracies of 71.88 percent and 43.04 percent are obtained. A final advantage of 28.84 points is therefore achieved. Second, a delayed benefit is observed on CIFAR-100. The proposed method's accuracy is 6.20 points lower after the first task and 0.55 points lower after the second task. Experience replay is surpassed from the third task onward, and a final advantage of 20.12 points is obtained. This pattern is interpreted as a limited sacrifice in early-task fitting in exchange for stronger long-term retention. Sequential performance above the offline reference on DVS128 Gesture and CIFAR-100 is attributed to differences in architecture and optimization rather than to superiority over a matched offline upper bound.

Forgetting is evaluated directly in Figure 4. For each previously learned task, forgetting is defined as the difference between its best accuracy during the stream and its final accuracy. The mean across previously learned tasks is represented by the dashed line. Loss of acquired knowledge can thereby be distinguished from the absolute difficulty of each dataset.

Substantial dataset-dependent differences are revealed by the forgetting profiles. An average forgetting rate of 0.26 percent is obtained on MNIST, and measurable loss is concentrated in the first task. An average rate of 2.60 percent is obtained on DVS128 Gesture, and most of the reduction is assigned to the earliest gesture group. A markedly larger rate of 19.57 percent is obtained on CIFAR-10. Losses of 26.00 and 29.95 points are observed for Tasks 0 and 1, whereas only 7.60 points are lost on Task 3. An average rate of 10.72 percent is obtained on CIFAR-100 despite the ten-task sequence. Losses between 9 and 18.5 points are repeatedly observed for early and middle tasks, whereas only 1.50 points are lost on Task 8. The remaining interference is therefore concentrated in knowledge exposed to the largest number of later updates. Lower forgetting is measured on CIFAR-100 than on CIFAR-10, but lower final accuracy is still obtained. Absolute performance is consequently limited by both retention and representational difficulty.

A cross-method comparison is provided in Table 1. Accuracies between 16.18 percent and 25.57 percent are obtained by regularization-based methods across the four Class-IL benchmarks. Historical input distributions cannot be reconstructed by parameter constraints alone. Accuracies of 97.03 percent on Split-MNIST and 94.10 percent on DVS128 Gesture are obtained by the proposed method, thereby yielding the highest listed values. On CIFAR-10, an accuracy of 71.28 percent is obtained. This result is close to the 70.51 percent obtained by Dark Experience Replay and is 1.42 points lower than the result obtained by the supervised-replay variant of Dark Experience Replay. On CIFAR-100, an accuracy of 60.13 percent is obtained. Differences of 0.34, 0.26, and 1.17 points are observed relative to DSD-SNN, Dark Experience Replay, and its supervised-replay variant, respectively. The table is used to establish competitiveness rather than a universal ranking because different architectures and protocols are used across the listed studies.

Detailed forgetting-rate profiles and the complete Class-IL cross-method comparison are provided in Appendix A.

4.3 Task-Incremental Learning Results

Task-IL performance is presented in Figure 3. Prediction is restricted to the active output space by the supplied task identity. Inter-task class competition is therefore removed, and the available margin for improvement is reduced. The remaining differences are mainly associated with the resistance of the shared representation to sequential drift.

Different regimes are observed across the four datasets. Near saturation is reached on MNIST. Final accuracies of 99.41 percent and 99.04 percent are obtained by the proposed method and experience replay, respectively. The gain of 0.37 points is therefore interpreted as improved stability rather than increased classification capacity. On DVS128 Gesture, an accuracy of 100 percent is obtained by the proposed method after the second task, and a final advantage of 2.94 points is retained after the third task. Continued perturbation of the shared temporal representation by the last gesture group is indicated by the final decline of both curves. A less uniform pattern is observed on CIFAR-10. The proposed method's accuracy is 2.35 points lower after the third task, but a final advantage of

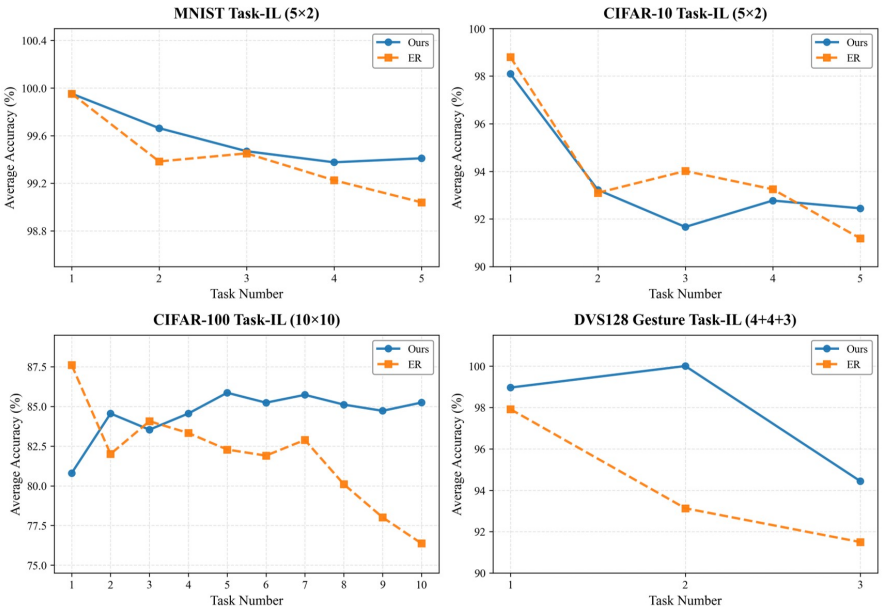

Figure 3: Average accuracy under Task-IL on four benchmarks.

1.26 points is recovered after further consolidation. A clear long-sequence advantage is obtained on CIFAR-100. An initial disadvantage of 6.80 points is obtained by the proposed method, after which its accuracy is maintained between 83.53 percent and 85.86 percent from the third task onward. A decline to 76.37 percent is observed for experience replay, and the final margin is expanded to 8.87 points. The importance of competition in a shared output space for Class-IL difficulty is indicated by the smaller gains on MNIST and CIFAR-10. The continued importance of feature preservation when task identity is available is indicated by the persistent CIFAR-100 gain.

The curve-based interpretation is supported by the reference comparison in Table 2. An accuracy of 99.65 percent is obtained on Split-MNIST, for which accuracies above 98 percent are already obtained by most methods. On DVS128 Gesture, an accuracy of 97.22 percent is obtained, which is 3.20 points higher than that of the supervised-replay variant of Dark Experience Replay listed in the table. On CIFAR-10, an accuracy of 93.34 percent is obtained. Differences of only 0.06 and 0.54 points are observed relative to Dark Experience Replay and its supervised-replay variant, respectively. On CIFAR-100, an accuracy of 85.32 percent is obtained. The strongest listed non-proposed result is exceeded by 6.73 points, and experience replay is exceeded by 8.89 points. A larger benefit is therefore observed when a longer task sequence is processed by a shared representation. Because different protocols are aggregated in the table, the principal evidence for comparative claims is provided by the matched curves.

5 Conclusion

A sleep-inspired two-stage continual learning framework has been proposed to mitigate catastrophic forgetting in SNNs. Neuronal-state-conditioned regularization is combined with source-dependent replay during daytime adaptation, whereas approximately class-balanced memory replay is used during nighttime consolidation. Old-class response distillation is retained in both stages to preserve previously learned knowledge. Rapid acquisition of new knowledge is thereby supported while interference with earlier tasks is limited. Experimental results across multiple benchmarks demonstrated consistent improvements in final classification accuracy over experience replay under both Class-IL and Task-IL. These results support the effectiveness of separating adaptation from consolidation in balancing stability and plasticity. Adaptive memory allocation, more efficient consolidation schedules, and hardware-aware evaluation will be investigated to improve early-task retention and reduce consolidation costs.

AI USE STATEMENT

Generative AI assisted with LaTeX typesetting, table preparation, correction of research code, mathematical notation conversion, and formatting checks. All scientific ideas, experiments, results, and claims were provided by the authors, who reviewed and approved the final submission.

REFERENCES

Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: learning what (not) to forget. In *Proceedings of the European conference on computer vision*, pp. 139–154, 2018.

Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. In *Advances in neural information processing systems 33*, pp. 15920–15930, 2020.

Lucas Caccia, Rahaf Aljundi, Nader Asadi, Tinne Tuytelaars, Joelle Pineau, and Eugene Belilovsky. New insights on reducing abrupt representation change in online continual learning. In *10th International Conference on Learning Representations*, pp. 1–27, 2022.

Liangwei Fan, Hui Shen, Xiangkai Lian, Yulin Li, Man Yao, Guoqi Li, and Dewen Hu. A multisynaptic spiking neuron for simultaneously encoding spatiotemporal dynamics. *Nature Communications*, 16(1):7155–7173, 2025.

Ryan Golden, Jean Erik Delanois, Pavel Sanda, and Maxim Bazhenov. Sleep prevents catastrophic forgetting in spiking neural networks by forming a joint synaptic weight representation. *PLoS Computational Biology*, 18(11):1–31, 2022.

Bing Han, Feifei Zhao, Yi Zeng, Wenxuan Pan, and Guobin Shen. Enhancing efficient continual learning with dynamic structure development of spiking neural networks. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pp. 1–9, 2023.

Bing Han, Feifei Zhao, Yang Li, Qingqun Kong, Xianqi Li, and Yi Zeng. Similarity-based context aware continual learning for spiking neural networks. *Neural Networks*, 184:107037–107048, 2025.

James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences of the United States of America*, 114(13):3521–3526, 2017.

Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017.

David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, pp. 6470–6479, 2017.

James L Mcclelland, Bruce L Mcnaughton, Randall C O'reilly, Carol Barnes, Sue Becker, Marlene Behrmann, Loren Frank, Mark Gluck, Robert French, Patrick Lynn, Mort Mishkin, Catherine Myers, David Plaut, William Skaggs, Craig Stark, Robert Sutherland, Larry Squire, and Matthew Wilson. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102(3):419–466, 1995.

Michael Mccloskey and Neal J Cohen. Catastrophic interference in connectionist networks: the sequential learning problem. *Psychology of learning and motivation*, 24:109–165, 1989.

Mishal Fatima Minhas, Rachmad Vidya Wicaksana Putra, Falah Awwad, Osman Hasan, and Muhammad Shafique. Continual learning with neuromorphic computing: foundations, methods, and emerging applications. *IEEE Access*, 13:124824–124873, 2025a.

Mishal Fatima Minhas, Rachmad Vidya Wicaksana Putra, Falah Awwad, Osman Hasan, and Muhammad Shafique. Replay4NCL: An efficient memory replay-based methodology for neuromorphic continual learning in embedded AI systems. In *62nd ACM/IEEE Design Automation Conference (DAC)*, pp. 1–7, 2025b.

Emre O. Neftci, Hesham Mostafa, and Friedemann Zenke. Surrogate gradient learning in spiking neural networks. *IEEE Signal Processing Magazine*, 36(6):51–63, 2019.

Wenyao Ni, Jiangrong Shen, Qi Xu, and Huajin Tang. ALADE-SNN: Adaptive logit alignment in dynamically expandable spiking neural networks for class incremental learning. In *Proceedings of the AAAI conference on artificial intelligence*, pp. 19712–19720, 2025.

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*, pp. 2001–2010, 2017.

David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy P. Lillicrap, and Greg Wayne. Experience replay for continual learning. In *Advances in Neural Information Processing Systems*, pp. 350–360, 2019.

Jiangrong Shen, Wenyao Ni, Qi Xu, and Huajin Tang. Efficient spiking neural networks with sparse selective activation for continual learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 611–619, 2024.

Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *Advances in neural information processing systems 30*, pp. 2994–3003, 2017.

Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, and Jon Shlens. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.

Gido M. van de Ven, Tinne Tuytelaars, and Andreas S. Tolias. Three types of incremental learning. *Nature Machine Intelligence*, 4(12):1185–1197, 2022.

Diana Vins, Jean Erik Delanois, and Maxim Bazhenov. When to learn and when to stop: quitting at the optimal time. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 29517–29518, 2025.

Mingqing Xiao, Qingyan Meng, Zongpeng Zhang, Di He, and Zhouchen Lin. Hebbian learning based orthogonal projection for continual learning of spiking neural networks. In *12th International Conference on Learning Representations*, pp. 16271–16293, 2024.

Jiangshuai Xu, Peiyun Xue, Jiacheng Song, Xuhui Huang, and Qingshan Hou. HLML-SNN: Fast continual learning in spiking neural networks achieved via hebbian learning-driven meta-learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 27260–27268, 2026.

Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 3987–3995, 2017.

Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. Mixup: beyond empirical risk minimization. In *6th International Conference on Learning Representations*, pp. 1–13, 2018.

A Additional Experimental Results

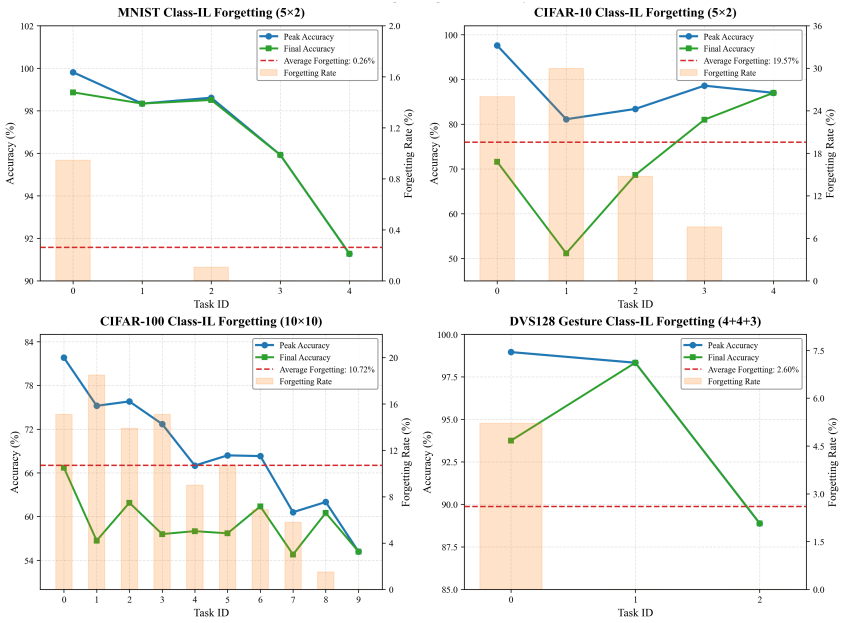

Figure 4: Forgetting-rate analysis of the proposed method under Class-IL.

Table 1: Class-IL comparison with representative continual-learning methods (%).

Method	Split-MNIST	DVS128 Gesture	CIFAR-10	CIFAR-100
EWC	20.01 ± 0.06	20.49 ± 0.19	16.18 ± 1.37	17.25 ± 0.09
SI	19.99 ± 0.06	21.53 ± 0.27	17.27 ± 0.87	17.26 ± 0.11
LwF	23.85 ± 0.44	24.93 ± 0.41	18.53 ± 0.12	25.57 ± 0.27
MAS	19.50 ± 0.30	21.60 ± 0.26	20.25 ± 1.54	17.07 ± 0.12
iCaRL	92.49 ± 0.12	67.76 ± 0.43	49.02 ± 3.20	37.83 ± 0.21
GEM	86.3 ± 1.4	–	25.54 ± 0.76	–
ER	82.1 ± 1.5	86.67 ± 0.39	44.79 ± 1.86	37.57 ± 0.21
ER-MIR	87.6 ± 0.7	87.81 ± 0.40	47.6 ± 1.1	–
ER-ACE	82.98 ± 1.79	85.94 ± 0.27	35.16 ± 1.34	38.75
A-GEM	29.02 ± 5.3	27.23 ± 0.28	20.04 ± 0.34	20.38 ± 1.45
DGR	90.35 ± 0.24	–	–	9.67 ± 0.22
DER	96.12	90.79 ± 0.53	70.51 ± 1.67	60.39 ± 0.27
DER++	96.37	91.48 ± 0.45	72.7 ± 1.36	61.30 ± 0.10
DSD-SNN	–	–	–	60.47 ± 0.72
Ours	97.03 ± 1.01	94.10 ± 0.48	71.28 ± 1.51	60.13 ± 0.96

Table 2: Task-IL comparison with representative continual-learning methods (%).

Method	Split-MNIST	DVS128 Gesture	CIFAR-10	CIFAR-100
EWC	98.64 ± 0.22	84.38 ± 0.43	80.89 ± 0.69	76.34 ± 0.29
SI	99.09 ± 0.15	73.41 ± 0.79	82.13 ± 0.58	74.84 ± 0.39
LwF	99.60 ± 0.03	72.92 ± 0.71	85.49 ± 1.32	78.59 ± 0.24
MAS	99.22 ± 0.21	75.29 ± 0.47	81.68 ± 1.05	–
iCaRL	–	88.54 ± 0.28	75.63 ± 1.42	–
GEM	98.42 ± 0.10	77.17 ± 0.45	90.44 ± 0.94	–
ER	98.98 ± 0.07	92.44 ± 0.40	91.19 ± 0.94	76.43 ± 0.24
ER-MIR	–	89.89 ± 0.24	–	–
ER-ACE	–	88.61 ± 0.37	–	–
A-GEM	98.54 ± 0.10	77.78 ± 0.53	83.88 ± 1.49	73.30 ± 0.39
DGR	99.50 ± 0.03	–	–	71.40 ± 0.32
DER	–	93.97 ± 0.35	93.40 ± 0.39	–
DER++	–	94.02 ± 0.49	93.88 ± 0.50	–
Ours	99.65 ± 0.10	97.22 ± 0.65	93.34 ± 1.34	85.32 ± 0.75