arXiv:2608.19514v1 [cs.LG] 20 Aug 2026

Graphical Abstract

In Two Minds about Lifelong Learning: Exploring Hemispheric Redundancy and Specialisation in Neural Models

Benjamin Smith, Levin Kuhlmann, Kaushik Roy, Gideon Kowadlo

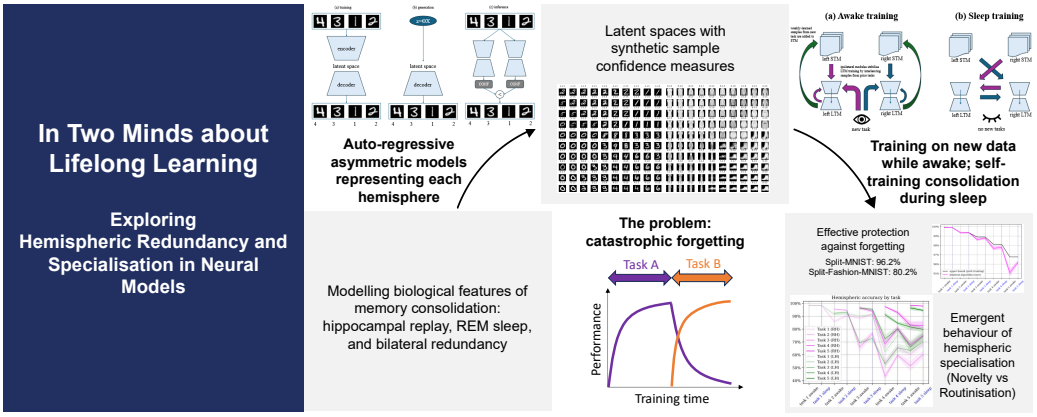

In Two Minds about Lifelong Learning: Exploring Hemispheric Redundancy and Specialisation in Neural Models

Benjamin Smith[a,*], Levin Kuhlmann[a], Kaushik Roy[c], Gideon Kowadlo[1,b]

[a]*Data Science and AI, Monash University, Wellington Road, Clayton, 3800, Victoria, Australia*
[b]*Cerenaut, Victoria, Australia*
[c]*CSIRO Robotics, CSIRO, Australia*

Abstract

Persistent intelligent systems require the ability to learn continually, but current machine learning approaches face significant challenges in this area compared to biological learning systems. Machine learning algorithms typically trade off retention of previously learned information and adaptation to new or changing data patterns. When continual learning capabilities are absent, algorithms must undergo retraining using the entire data set, an approach that becomes impractical when original training data are unavailable due to storage constraints, financial or computational costs, or privacy restrictions. However, biological animals can learn continually, without experiencing catastrophic forgetting. This paper attempts to build a high-level framework for how animals learn and preserve knowledge by modelling neural components and states that are known to be related to memory consolidation. We focus on three concepts: experience replay, REM sleep, and bilaterality. We propose 4MAS (4 Module Awake/Sleep), a novel macroarchitecture demonstrating how machine learning models might benefit from asymmetric hemispheres, each with their own long- and short-term memory mechanisms, and how a period of sleep between incremental learning tasks might benefit memory consolidation. Finally, we present results showing that our architecture achieves competitive results on the Split-MNIST, Split-Fashion-MNIST and Split-CIFAR-100 datasets, with 98.3%, 84.9%, and 29.29% accuracy re-

*Corresponding author

spectively.

Keywords: Continual Learning, Lifelong Learning, Experience Replay, Bilateral Deep Learning

---

1. Introduction

Contemporary artificial neural networks suffer from catastrophic forgetting (or catastrophic interference) [1, 2], where learning new tasks rapidly erodes previously acquired skills. While standard training assumes independent, identically distributed (i.i.d.) data, real-world applications present sequentially correlated distributions where class presentation frequencies vary arbitrarily.

This forgetting stems from the stability-plasticity dilemma: adapting weights to new data modifies parameters that are critical for earlier tasks [3, 4]. Joint training across all data streams is fundamentally impossible in continual learning settings where future data are unavailable and past data are restricted by storage or privacy constraints. Naive sequential fine-tuning, while feasible, suffers from severe catastrophic forgetting. Continual learning algorithms seek to resolve this dilemma without full dataset retraining.

In contrast, biological brains mitigate forgetting through specialised neural architectures and consolidation states. The mammalian hippocampus uses experience replay to weave together recent and distant memories, transferring knowledge to the neocortex during sleep [5, 6, 7]. Furthermore, sleep phases like rapid-eye-movement (REM) sleep exhibit highly coordinated bilateral (inter-hemispheric) activity, which contrasts with the unihemispheric states of slow-wave sleep [8, 9]. Finally, there is evidence that biological brains leverage two lateralised cortical hemispheres that acquire distinct representations, with the right hemisphere handling novel information and the left hemisphere optimised for stable routine tasks [10, 11].

Existing generative replay methods such as Generative Replay struggle to scale because a single generator must capture all historical distributions, leading to compounding representational drift [12]. To address these limitations, we propose 4MAS (4 Module Awake/Sleep), a novel continual learning macroarchitecture that implements two asymmetric hemispheres—each with dedicated short- and long-term memory modules—and an explicit offline sleep phase for cross-hemispheric consolidation; representing hippocampus and neocortex respectively. By distributing replay across two specialised

generators and anchoring the latent space during sleep, 4MAS minimises representational drift.

Our main contributions are:

- We propose a dual-hemisphere continual learning architecture (4MAS) that splits generative replay across specialised exploratory and conservative models, mimicking biological lateralisation.
- We introduce an explicit wake-sleep training cycle that uses cross-hemispheric consolidation (mutual fine-tuning during a simulated sleep phase) to stabilise latent spaces and prevent representational drift.
- We demonstrate that 4MAS achieves competitive Class-IL accuracies on Split-MNIST (98.3%), Split-Fashion-MNIST (84.9%), and Split-CIFAR-100 (29.29%); while displaying extremely low representational drift across tasks.

1.1. Biological motivation

Inspired by this neurobiological template, 4MAS employs two generative long-term memories as an asymmetrical ensemble: one retains plasticity for novel task acquisition, while the other stabilises historical patterns. Generative memories are paired with short-term memory buffers, allowing for rehearsal akin to hippocampal replay. Training follows an ultradian-inspired cycle, alternating between a unihemispheric awake phase for task learning and a bilateral sleep phase for cross-hemispheric memory harmonisation.

2. Background and related work

Continual learning is typically evaluated under three scenarios [13] of increasing difficulty:

- **Task-IL (Task Incremental Learning):** The model learns a sequence of tasks with distinct data distributions and separate output spaces. The task identity is explicitly provided during both training and inference, allowing the model to select task-specific parameters or heads.
- **Domain-IL (Domain Incremental Learning):** The model is exposed to a sequence of tasks with varying input distributions that share a common output space. The task identity is never provided, requiring the model to adapt to changing domains without contextual routing.

- **Class-IL (Class Incremental Learning):** The most challenging paradigm, where both input distributions and output spaces vary across tasks. No task identifiers are available at training or inference. The model must classify inputs across all classes seen so far, making it highly susceptible to catastrophic forgetting.

When evaluating these models, performance is typically benchmarked against fine-tuning (sequential training without forgetting mitigation; lower bound) and joint-training (simultaneous training on all data; upper bound). To bridge the gap to joint-training under Class-IL constraints, algorithms are commonly grouped into regularisation, parameter isolation, and rehearsal.

2.1. Regularisation

Regularisation methods restrict gradient updates to protect parameters that are critical for prior tasks. Early approaches froze lower layers of a network after training on a task [14]. Modern algorithms estimate parameter importance by calculating contribution to loss reduction (e.g., Synaptic Intelligence [15]), approximating Bayesian inference (e.g., Elastic Weight Consolidation [4]), or evaluating output function sensitivity (e.g., Memory Aware Synapses [16]). Recent methods perform geometric analysis to locate stable flat minima in the distribution manifold [17, 18] or approximate prior losses using the Hessian matrix eigenvalues [19].

While regularisation is highly effective in Task-IL, it struggles in Class-IL [20]. The strict constraints designed to protect existing knowledge prevent the model from adapting to novel classes, leading to representational paralysis as constraints accumulate over sequential tasks [18].

2.2. Parameter isolation

Parameter isolation limits training to a subset of parameters or dynamically expands the network structure. A common design adds task-specific output heads to a pretrained feature-extractor backbone [21, 22]. Under network expansion, the network freezes historical weights and adds new neurons or columns for each task [23, 24]. Another branch of research leverages Adaptive Resonance Theory to dynamically create category nodes as new distributions appear [25].

However, these approaches struggle in Class-IL settings because they require task identity during inference to route data to the appropriate subnetwork. Without explicit task identifiers, the model lacks an intrinsic mech-

anism for choosing the correct task pathway. Although sparse neural activation in large networks can mitigate routing issues [26], parameter isolation remains difficult to scale without introducing capacity exhaustion or inference-routing failures [27, 28].

2.3. Rehearsal: replay and generative replay

Rehearsal techniques interleave historical data with new task inputs. Experience replay stores a memory buffer of real samples from earlier tasks, and Gradient Episodic Memory (GEM) uses those stored examples as inequality constraints on each update so the loss on previous tasks does not increase, which reduces forgetting while still allowing positive backward transfer [29]. Maximally Interfered Retrieval (MIR) instead keeps a finite replay memory, estimates the parameter update from the current batch, and then replays the buffered samples whose losses would increase the most under that update, making it distinct from GEM because it prioritises most-interfered samples rather than enforcing explicit gradient constraints [30]. AdaER further adapts replay by using Contextually-Cued Memory Recall to select memories based on both data-conflicting and task-conflicting cues, and it also updates the buffer with Entropy-Balanced Reservoir Sampling to keep a more balanced, informative memory, distinguishing it from MIR's interference-only retrieval and GEM's constraint-based updates [31].

While buffer-based replay is effective, it scales poorly because representing complex, high-dimensional distributions requires a prohibitive number of stored samples [32].

Generative replay avoids storing real data by training a generative model to synthesise historical inputs [33]. To scale beyond simple datasets, models reconstruct latent feature representations rather than raw training samples. For example, Generative Feature Replay (GFR) uses a feature extractor to train a generator on latent distributions [34], mimicking biological systems where memory replay occurs at representational rather than raw sensory levels [35]. Brain-Inspired Replay (B-IR) combines generative replay with parameter isolation to achieve state-of-the-art results [12]. Nonetheless, generative models suffer from representational drift as the feature extractor updates over time, requiring distillation constraints or frozen features to anchor the latent space [36].

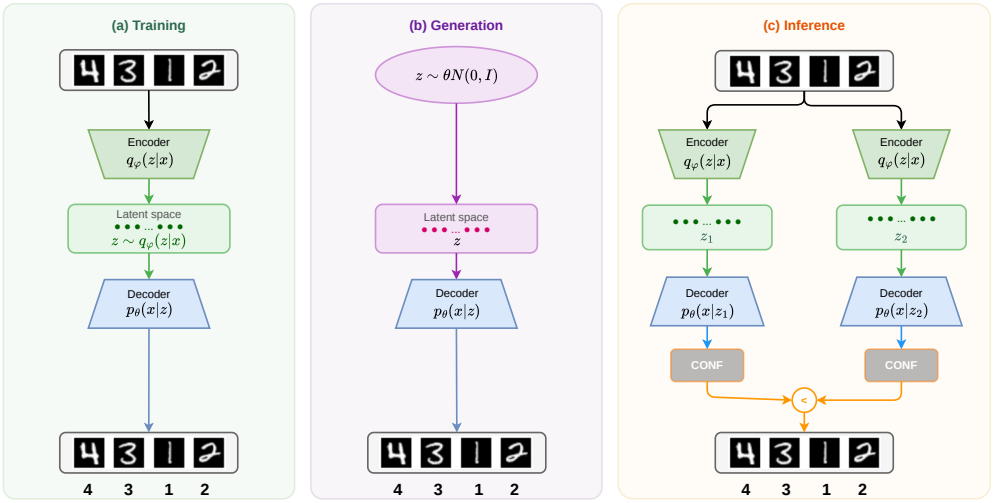

Figure 1: Long-term memory inputs and outputs: (a) LTM training on image-label pairs. (b) Synthetic sample generation with temperature τ. (c) Inference confidence-based selection.

3. Methods

We present 4MAS (4 Module Awake/Sleep), a macroarchitecture for continual learning designed to model biological memory consolidation processes. As illustrated in Figure 1, the system consists of two lateralised hemispheres. Each hemisphere contains:

1. A generative Long-Term Memory (LTM) representing the neocortex, which learns task distributions and classifies incoming data.
2. A Short-Term Memory (STM) buffer representing the hippocampus, which stores a small set of episodic exemplars.

The architecture restricts data flow to a biologically inspired model where task acquisition is unihemispheric and offline consolidation is bilateral, facilitating knowledge transfer and specialisation.

3.1. Benchmarks and experimental design

To evaluate 4MAS under the Class-IL constraint, we employ Split-MNIST [37, 38] and Split-Fashion-MNIST [39, 40], shown in Figure 2. The standard 10-class datasets are split into 5 sequential tasks of 2 classes each, presented

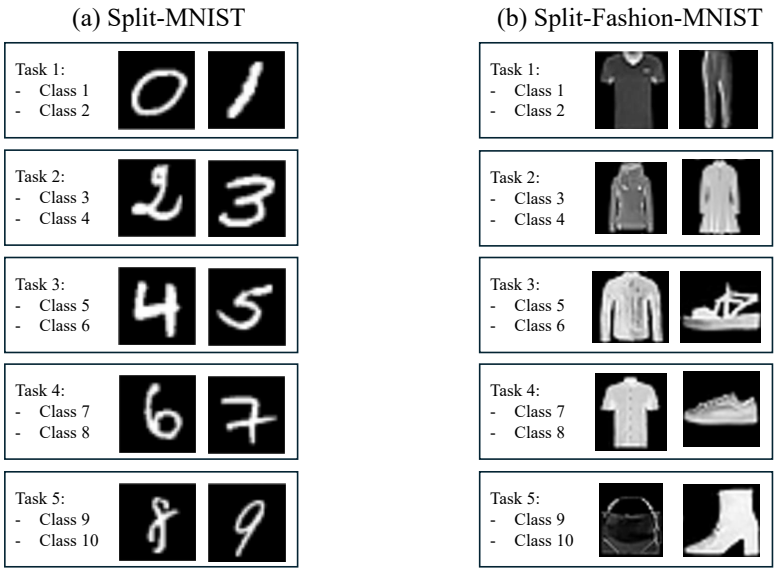

Figure 2: Class-IL datasets: (a) Split-MNIST, (b) Split-Fashion-MNIST.

without task identifiers at both training and inference. Hyperparameter tuning and initial method development were conducted primarily on these two MNIST variants. To test scalability and out-of-the-box generalizability, we then expanded evaluation to the more challenging Split-CIFAR-100 dataset [38]. Unless otherwise stated, each configuration is run for 10 random seeds, and we report the mean and standard error of the mean of each metric across seeds. The model trains for 10 epochs per class, with a synthetic epoch length of 10,000 samples. Ablation studies disentangle the contributions of the STM buffers, the dual-hemisphere ensemble, and the sleep phase.

3.2. Sleep/wake phases

The training cycle alternates between a unihemispheric "awake" phase and a bilateral "sleep" phase for each task, Figure 3. During the awake phase, Algorithm 1, each LTM hemisphere is trained on three randomly interleaved sources: new task data, exemplars stored in the ipsilateral STM from prior tasks, and self-generated synthetic samples. After training, the LTM evaluates the training data and updates its STM buffer with samples that meet the selection criteria (Section 3.4).

During the sleep phase, Algorithm 2, LTMs undergo fine-tuning on contralateral representations (generations from the opposite LTM and exemplars

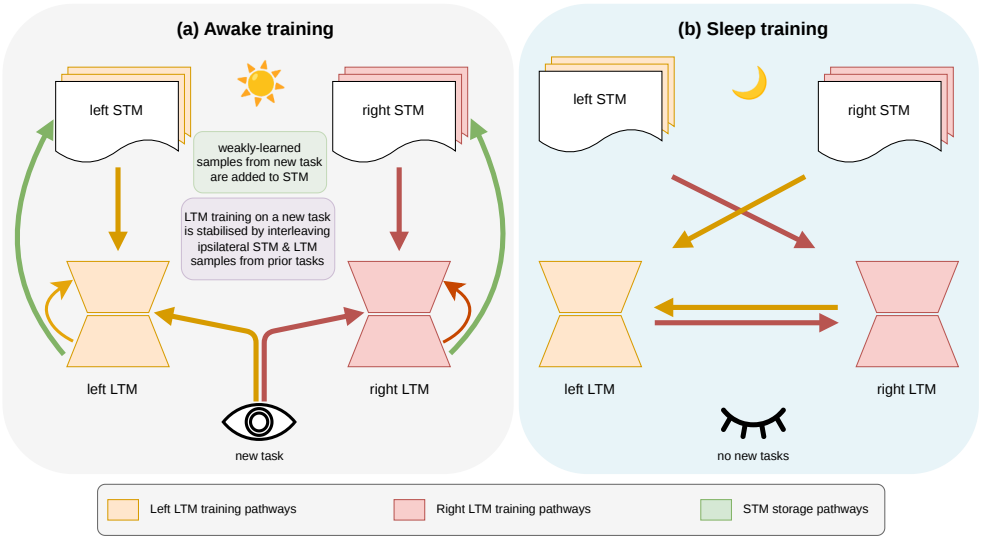

Figure 3: Awake/sleep training phases: (a) Awake training: LTMs learn new task data interleaved with STM samples. (b) Sleep training: LTMs are fine-tuned on contralateral representations.

from the opposite STM) during a period where no new training data are available. To allow subtle adjustments to internal representations without destroying learned task-specific parameters, the sleep learning rate is scaled down by a multiplier $\lambda = 0.1$, mimicking the lower firing rates observed across brain regions during REM sleep [41].

Algorithm 1 Incremental task – awake phase

Input: Task $t \in [1, \infty)$, learning rate α, model parameters θ_{left}, θ_{right}, memories $\mathcal{M}_{left}$, $\mathcal{M}_{right}$, classes seen C, memory capacity M.
for $h \in \{left, right\}$ **do**
    $w \leftarrow \{w_i\}_{i=1}^{C}$ where $w_i = 1/C$
    for $e = 1 \dots E_{awake}$ **do**
        Draw batch β_t from $\mathcal{D}_t$.
        if $t > 1$ **then** generate replay $\mathcal{G}_h \sim p_{\theta_h}(x|z)$ and set $\beta \leftarrow \beta_t \cup \mathcal{G}_h \cup \mathcal{M}_h$ **else** $\beta \leftarrow \beta_t$
        $\beta \leftarrow \text{class_weighted_rebalance}(\beta, w)$
        $\theta_h \leftarrow \theta_h - \alpha \nabla_\theta \mathcal{L}(\theta_h; \beta)$
        $g \leftarrow \text{bincount}\big(\arg\max_c \hat{y}(\mathcal{G}_h)\big)$
        $w \leftarrow \{1/(g_i + \varepsilon)\}_{i=1}^{C}$
    end for
    for $i = 1 \dots C$ **do**
        $\mathcal{P}_{h,i} \leftarrow \text{Oversample}(\mathcal{D}_{t,i} \cup \mathcal{M}_{h,i})$ ▷ Expand candidate pool for class i
        $\mathcal{M}_{h,i} \leftarrow \text{TopK}_{\lfloor M/C \rfloor}\Big(\mathcal{P}_{h,i},\, CONF(\hat{y})\Big)$ ▷ Select top-k highest confidence
    end for
    $\mathcal{M}_h \leftarrow \bigcup_{i=1}^{C} \mathcal{M}_{h,i}$
end for

Algorithm 2 Incremental task – sleep phase

Input: Task $t \in [1, \infty)$, learning rate α, models θ_{left}, θ_{right}, memories $\mathcal{M}_{left}$, $\mathcal{M}_{right}$, sleep learning multiplier $\lambda = 0.1$.
for h in $\{left, right\}$ **do**
    $w \leftarrow \{w_i\}_{i=1}^{C}$ where $w_i = 1/C$
    for $e = 1 : E_{sleep}$ **do**
        $\beta \leftarrow \theta_h(z) \cup \mathcal{M}_h$
        $\tau_h \leftarrow \theta_h - \lambda \alpha \nabla_\tau \mathcal{L}(\theta_h; \beta)$.
    end for
end for

3.3. Long-term memories

Each LTM is both a generator and a classifier within a single variational auto-encoder (VAE). The classification vector is appended to the input image, and the model is trained to reconstruct the joint vector $x \frown y$. This joint parameter space ensures that synthetic images and labels are tightly coupled. Let $\hat{y}$ denote the reconstructed label channel, normalised over the C classes seen so far. We define prediction confidence as the negative cross-entropy of $\hat{y}$ against its own arg-max class:

$$CONF(\hat{y}) = -\mathcal{L}_{CE}\big(\hat{y}, \text{onehot}(\arg\max_c \hat{y}_c)\big) = \log \max_c \hat{y}_c \tag{1}$$

$CONF \in (-\infty, 0]$, with larger values indicating higher confidence. Since $\mathcal{L}_{CE}$ is the same term the VAE minimises on the label channel, Equation 1 doubles as a measure of reconstruction quality and requires no separate discriminative head, providing stability when learning new tasks [42]. Confidence is used in three places: rejecting low-confidence generations during replay, ranking candidates for STM storage (Section 3.4), and arbitrating between hemispheres at inference.

Standard VAE latent sampling $z \sim \mathcal{N}(0, I)$ can suffer from posterior collapse, producing low-variance synthetic samples that degrade replay quality over sequential tasks. To resolve this, we apply a temperature gain τ to the latent coordinates during sampling: $z = \tau X, X \sim \mathcal{N}(0, I)$ [43]. Setting $\tau > 1$ expands the explored region of the latent space, generating sharper and more diverse samples, Figure 4. During replay, low-confidence generations near class boundaries are rejected to avoid interpolating between classes. At inference, inputs are processed by both LTMs, and the class prediction from the hemisphere with higher confidence is selected. To maintain class balance during replay, generation frequencies are weighted inversely to generation counts Algorithm 1, line 10, biased linearly by task age.

3.4. Short-term memories and lateralisation

Each STM acts as a memory buffer storing 50 exemplars. We evaluate two buffer storage selection mechanisms:

1. **Comparative Confidence Selection (CCS):** Samples from the task and prior STM are passed through the post-training LTM and ranked by confidence. In symmetric configurations, moderate-confidence quantile range samples ($\sim 50\%$) are selected in both hemispheres to represent distribution boundaries while retaining discernibility. Under

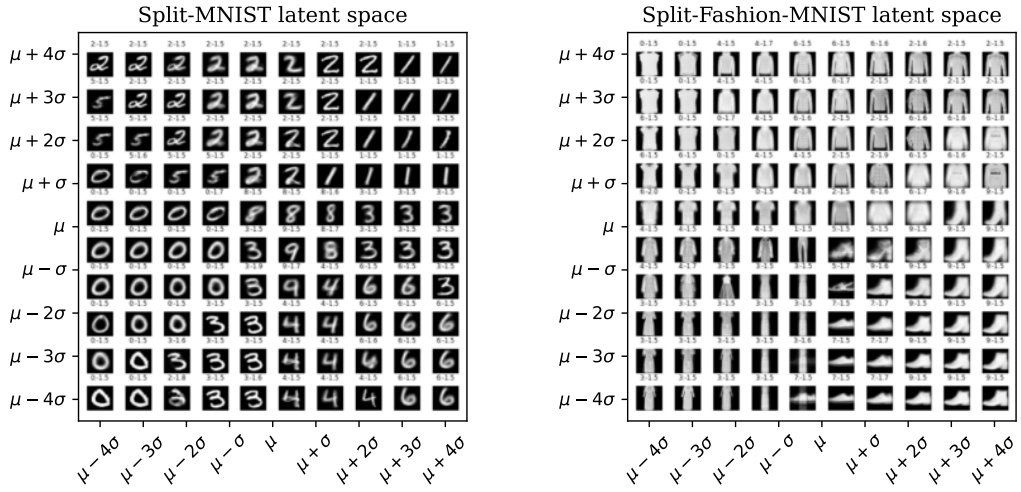

Figure 4: 2-dimensional mapping of the LTM's latent space after training. τ scales the sampling radius in units of the prior standard deviation. High temperature $\tau > 1$ increases the diversity of output representations (16 latent dimensions used in final configurations).

asymmetric configurations, STM storage selection is also lateralised: the Left Hemisphere selects high-confidence anchor samples to reinforce representations against drift, while the Right Hemisphere selects moderate-confidence samples to explore decision boundaries. Specifically, we filter replay samples using the raw probability $\max_c \hat{y}_c = \exp(CONF(\hat{y})) \in [0, 1]$, retaining candidates within the moderate-confidence quantile range $\{L, R\} = \{0.1, 0.5\}$. This design helps avoid VAE posterior collapse and representational drift.

2. **Latent Space Cluster Centroids (LSCC):** We apply K-Means clustering in the VAE latent space across task and STM samples, retaining exemplars closest to the centroids. This grounds representations and further prevents representational drift.

To model biological hemispheric lateralisation (stability vs. plasticity), we configure hyperparameter asymmetry. The Left Hemisphere (LH) is configured for stability and routine processing (lower generator temperature τ and task-age bias favouring older tasks). The Right Hemisphere (RH) is configured for plasticity and exploration (higher generator temperature τ and no age bias), enabling rapid adaptation to novel distributions.

4. Results and discussion

Our results show that 4MAS achieves competitive Class-IL accuracies across all benchmarks: $98.3 \pm 0.0\%$ on Split-MNIST, $84.9 \pm 0.3\%$ on Split-Fashion-MNIST, and $29.29 \pm 0.29\%$ on Split-CIFAR-100 (Table 1). These results are competitive with other recent Class-IL methods, such as B-IR (93.5%, 74.6%, and 27.85% accuracy) [38] and AdaER (89.6% and 74.0% accuracy) [31]. As shown in Table 1, 4MAS demonstrates strong resilience to catastrophic forgetting. On Split-MNIST, 4MAS achieves 98.3% accuracy, closely matching the Joint training ceiling of 98.0%. On Split-Fashion-MNIST, our method achieves $84.9 \pm 0.3\%$, representing a significant advancement over B-IR ($74.6 \pm 0.4\%$) and GR ($71.3 \pm 0.3\%$). This performance boost is directly tied to our low Forgetting scores (1.0% on Split-MNIST, 10.6% on Split-Fashion-MNIST) and exceptionally low representational drift (1.2 and 1.0, respectively, compared to GR's 215.0 and 48.5). This demonstrates that the dual-hemisphere and sleep-phase architecture stabilises internal latent representations, preventing the drift that typically destabilises single-generator networks.

On the challenging Split-CIFAR-100 dataset and when scaling each architecture to 140M trainable parameters, 4MAS improves on the unilateral B-IR baseline ($27.85 \pm 0.55\%$) and significantly outperforms other generative replay methods. The performance gap relative to Joint training ($51.9 \pm 0.4\%$) is primarily due to the expressive capacity of the standard flat VAE decoder, which struggles to reconstruct high-frequency details for 100 complex classes.

4.1. Backward and Forward Transfer Analysis

To evaluate how sequence learning affects historical and future task performance, we analyze Backward Transfer (BWT) and Forward Transfer (FWT), as defined in [29]:

$$\text{BWT} = \frac{1}{T-1}\sum_{i=1}^{T-1}(R_{T,i} - R_{i,i}), \quad \text{FWT} = \frac{1}{T-1}\sum_{i=2}^{T}(R_{i-1,i} - \bar{b}_i) \tag{2}$$

where $R_{T,i}$ represents test accuracy on task i after training on task T, and $\bar{b}_i$ denotes random baseline performance for task i.

Table 1: Class-IL evaluation across Split-MNIST, Split-Fashion-MNIST, and Split-CIFAR-100 benchmarks. Accuracy metrics report final task performance after learning all tasks.

Dataset	Method	Final Acc (%)	Forgetting (%)	BWT (%)	FWT (%)	Compute (x)
Split-MNIST	Joint	98.5 ± 0.0	1.0	-1.0	0.0	1.0
	4MAS (Our method)	98.3 ± 0.0	1.2	-1.2	0.0	1.0
	B-IR	93.5 ± 0.2	6.4 ± 0.2	-6.4 ± 0.2	0.2 ± 0.0	0.2 ± 0.0
	GR	91.2 ± 0.4	8.9 ± 0.4	-8.9 ± 0.4	0.3 ± 0.0	42.1 ± 1.8
	LwF	24.2 ± 0.4	0.1	0.1	0.0	15.2
	EWC	19.9 ± 0.0	79.2	-79.2	0.0	61.3
	oEWC	19.9 ± 0.0	79.2	-79.2	0.0	58.1
	SI	19.9 ± 0.0	79.3	-79.3	0.0	33.5
	Fine-Tuning	19.7 ± 0.1	79.5 ± 0.1	-79.5 ± 0.1	0.3 ± 0.0	84.2 ± 2.8
Split-Fashion-MNIST	Joint	88.3 ± 0.2	7.5 ± 0.2	-7.5 ± 0.2	0.0	1.0
	4MAS (Our method)	84.9 ± 0.3	10.6	-10.6	0.0	1.0
	B-IR	74.6 ± 0.4	30.2 ± 0.3	-30.2 ± 0.3	0.5 ± 0.0	0.2 ± 0.0
	GR	71.3 ± 0.3	25.4 ± 0.3	-25.4 ± 0.3	0.6 ± 0.1	48.5 ± 2.3
	LwF	20.1 ± 0.2	0.3	0.3	0.0	12.9
	EWC	19.9 ± 0.0	79.2	-79.2	0.0	55.5
	oEWC	19.8 ± 0.4	79.2	-79.2	0.0	50.3
	SI	19.9 ± 0.0	79.3	-79.3	0.0	30.4
	Fine-Tuning	20.0 ± 0.0	79.4 ± 0.0	-79.4 ± 0.0	0.5 ± 0.1	77.3 ± 2.3
Split-CIFAR-100	Joint	51.9 ± 0.4	16.2 ± 0.4	-16.2 ± 0.4	0.0	2.4
	4MAS (Our method, 140M)	29.29 ± 0.29	25.8 ± 1.2	-25.8 ± 1.2	0.0	0.43 ± 0.07
	GR	7.9 ± 0.1	66.1 ± 0.5	-66.1 ± 0.5	0.0	10.3
	B-IR (140M)	27.85 ± 0.55	52.9	-52.9	0.0	5.1
	LwF	10.8 ± 0.2	4.6 ± 0.3	-3.9 ± 0.3	0.0	1.9
	EWC	8.1 ± 0.1	81.6 ± 0.4	-81.6 ± 0.4	0.0	4.2
	SI	9.2 ± 0.2	79.8 ± 0.4	-79.8 ± 0.4	0.0	3.1
	Fine-Tuning	8.1 ± 0.1	81.3 ± 0.4	-81.3 ± 0.4	0.0	3.6

Backward Transfer Dynamics.. Table 1 demonstrates that 4MAS consistently minimises negative BWT compared to existing baselines. On Split-MNIST, Split-Fashion-MNIST, and Split-CIFAR-100, 4MAS achieves BWT scores of -1.2%, -10.6%, and -25.8% respectively, significantly outperforming competitive memory and replay methods (B-IR and GR). 4MAS tracks closely with the offline Joint Training baseline (e.g., -25.8% vs. -16.2% on CIFAR-100), confirming that our approach effectively freezes and preserves past decision boundaries during new class assimilation.

Forward Transfer Limitations in Class-IL.. Across all methods, FWT remains near 0.0%. This behaviour is characteristic of Class-IL benchmarks evaluated from scratch: without a shared pre-trained feature extractor, feature representations learned on early tasks do not inherently transfer zero-shot accuracy to orthogonal class boundaries in subsequent tasks. Thus, performance superiority in 4MAS is driven almost entirely by backward stability rather than forward inductive bias.

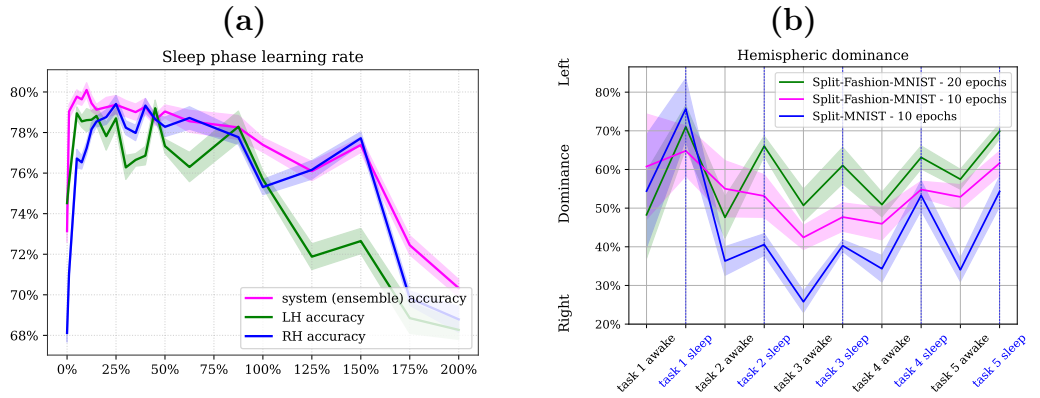

Figure 5: (a) Performance variation with sleep phase learning rates (LR) as a percentage of awake phase LR, Split-Fashion-MNIST. (b) Memory consolidation during sleep phase increased LH dominance, while new task learning during awake phase increased RH dominance; produced by lateralisation techniques (class representation bias by age in LH and asymmetric generator temperatures)

4.2. Sleep tuning

We evaluated sleep phase learning rates (LR) ranging from 0% to 200% of the awake phase rate (Figure 5a). On Split-Fashion-MNIST a reduced sleep rate of 5–10% ($\approx 2 \times 10^{-5}$) performed best ($84.9 \pm 0.3\%$), while removing sleep ($LR = 0\%$) yielded $74.0 \pm 0.59\%$. Larger learning rates degraded performance, confirming that low sleep learning rates enable fine-tuning on the contralateral representation without overwriting specialised features.

Furthermore, average ensemble accuracy exceeded either isolated hemisphere. This gain stems from effective confidence-based routing between specialised hemispheres rather than standard ensemble variance reduction; indeed, without sleep-phase consolidation, the ensemble fails to outperform the strongest individual hemisphere (Section 4.6). Hemispheric dominance shifted dynamically: wake training on new tasks increased RH dominance, while sleep consolidation restored LH dominance (Figure 5b). This shift aligns with Goldberg's Novelty-Routine hypothesis, reflecting a transition from initial RH-driven processing of novel representations to consolidated, routinised LH schemas.

4.3. Lateralisation

Figure 6a shows that the LH accuracy decays slowly, maintaining stability, while the RH learns new tasks quickly but forgets faster, supporting the

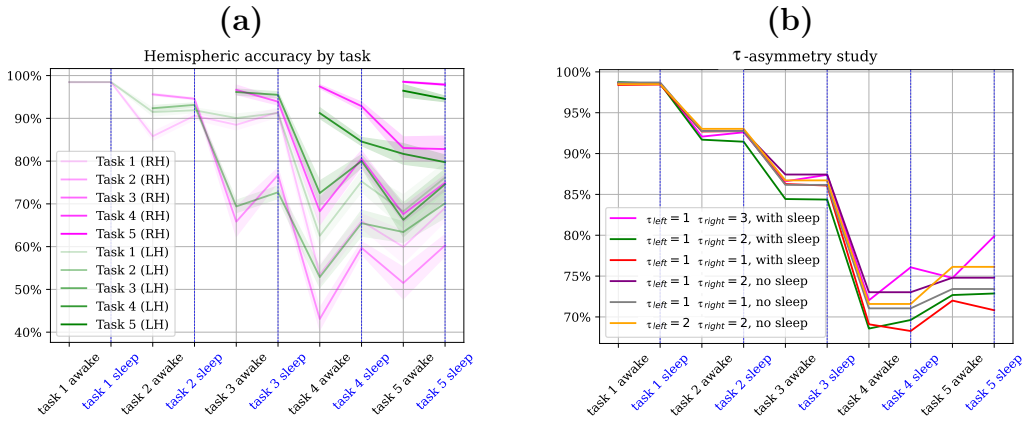

Figure 6: Fashion-MNIST: (a) Hemispheric accuracy for each task. (b) τ-asymmetry results, with and without sleep training phase.

		τ - right				
		1	**2**	**3**	**4**	**5**
τ - left	**1**	0.6301	0.7122	0.7445	0.7132	0.6984
	2	0.7021	0.6895	0.6662	0.6342	0.6178
	3	0.6768	0.6116	0.5907	0.5595	0.5504
	4	0.6593	0.6128	0.5561	0.5551	0.5233
	5	0.6453	0.576	0.5417	0.555	0.5148

Table 2: LH and RH generator temperature vs Split-Fashion-MNIST accuracy. Degraded performance observed for symmetrical temperatures.

stability-plasticity lateralisation described in the Novelty-Routine hypothesis [44]. Lateralising the generator temperature (τ) significantly improved performance (Table 2, Figure 6b). Symmetrical temperatures ($\tau = 1$ or $\tau = 2$ for both) caused degraded performance or limited sleep benefits, whereas asymmetrical configurations ($\tau = 1$ for LH, $\tau \in [2, 4]$ for RH) showed consistent performance gains after each sleep phase. This indicates that specialisation, elicited by asymmetrical parameterisation, provides an advantage and better use of total resources.

Applying a task-age class representation bias to the LH only (biasing training toward older tasks) improved the retention of earlier tasks by 15–26% (Figure 7a), allowing the ensemble to retain stability in LH while maintaining plasticity in RH.

		Confidence - RH						
		90%	**70%**	**50%**	**30%**	**10%**	**0%**	**K-Means**
Confidence - LH	**90%**	0.7562	0.7662	0.7668	0.7668	0.7536	0.7505	0.7656
	70%	0.7618	0.7724	0.7753	0.7725	0.7607	0.7589	0.7796
	50%	0.7695	0.7706	0.7733	0.7667	0.7648	0.7573	0.7779
	30%	0.7703	0.7718	0.7758	0.7627	0.7565	0.7247	0.7799
	10%	0.7673	0.7704	0.7675	0.7509	0.7351	0.727	0.7757
	0%	0.7547	0.7726	0.7613	0.7470	0.7285	0.7205	0.7662
	K-Means	0.7748	0.7776	0.7799	0.7764	0.7607	0.7588	0.776

Table 3: Comparison of mechanisms for selecting memories for STM buffer storage. Storing latent space centroids via K-Means clustering produced the most consistent results on Split-Fashion-MNIST, whereas highly asymmetric confidence selection ($L = 0.1/R = 0.5$) achieved peak performance on the more complex Split-CIFAR-100 benchmark.

4.4. STM buffer selection

Evaluating CCS and LSCC selection mechanisms (Table 3) showed that storing moderate-confidence samples ($\sim 50\%$) under CCS performed best, whereas selecting low-confidence outliers led to VAE posterior collapse. Storing latent space cluster centroids via K-Means (LSCC) yielded the best and most consistent results across MNIST and Fashion-MNIST tasks, indicating that the STM buffer is most effective when its primary role is grounding the latent space against representational drift rather than importing weakly-learned outliers. This aligns with neurobiological evidence showing that offline memory reactivation helps to preserve multiday representational stability [45].

The robust performance of the asymmetric threshold configuration ($L = 0.1, R = 0.5$) on CIFAR-100 is explained by a functional division of labour. The left hemisphere, operating at a low generator temperature ($\tau = 1$), behaves as a stable anchor that preserves core, high-confidence representations. By setting a very conservative threshold ($L = 0.1$), we prevent representational drift during consolidation. Conversely, the right hemisphere, operating at a high generator temperature ($\tau = 3$), acts as a flexible explorer. Storing intermediate-confidence boundary samples ($R = 0.5$) allows it to explore variations and refine task boundaries, leading to significantly enhanced ensemble consolidation.

4.5. Capacity and architectural scaling

To evaluate how model capacity influences continual learning performance, we systematically compared the parameter scaling behaviour of 4MAS

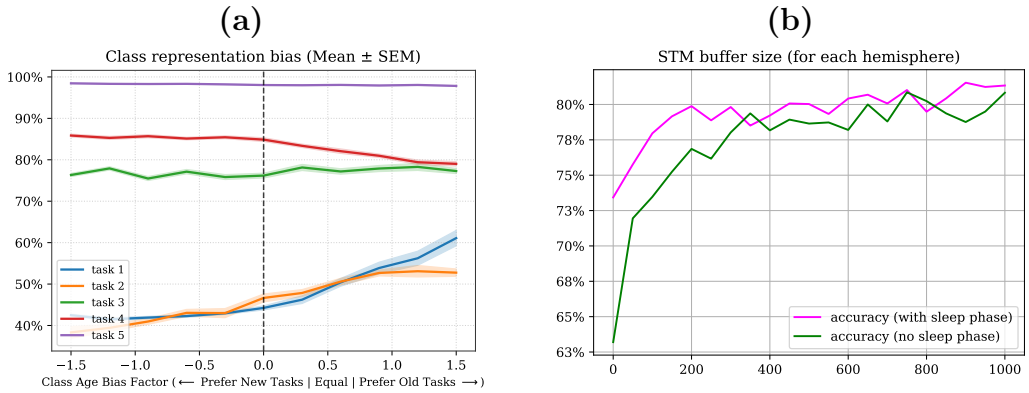

Figure 7: (a) Biasing class representation by recency; weighting training heavily toward older classes significantly boosts early-class accuracy with only minor performance degradation on recent classes. (b) Effect of STM memory buffer size on accuracy; gains from larger buffers plateau after 200 samples, while also diminishing the relative impact of the sleep phase.

against the unilateral B-IR baseline across Split-MNIST, Split-Fashion-MNIST, and Split-CIFAR-100 (Figure 8).

At lower parameter ranges (e.g., $< 30\,\mathrm{M}$ parameters), monolithic generative models demonstrate superior sample and parameter efficiency. This is primarily because 4MAS splits its total parameter budget across two distinct hemispheric models (LH and RH) and requires sleep-phase cross-replay to consolidate knowledge, introducing an architectural capacity overhead. When the overall parameter budget is highly constrained, the split-hemisphere bottleneck limits the representation capacity of the individual generators.

However, 4MAS demonstrates superior scalability as model capacity increases. Unilateral networks typically suffer from severe representational drift and catastrophic interference when forced to represent a large number of conflicting class distributions in a single unified latent space [38]. Consequently, performance of B-IR (and other unilateral models) plateaus or degrades at larger scales. In contrast, 4MAS's bilateral hemispheric partitioning and stability-plasticity division of labour mitigate representational drift, enabling monotonic scaling. At larger parameter scales ($\geq 70\,\mathrm{M}$ parameters), 4MAS consistently outperforms B-IR on the more complex Fashion-MNIST benchmark and approaches parity on the CIFAR-100 benchmark.

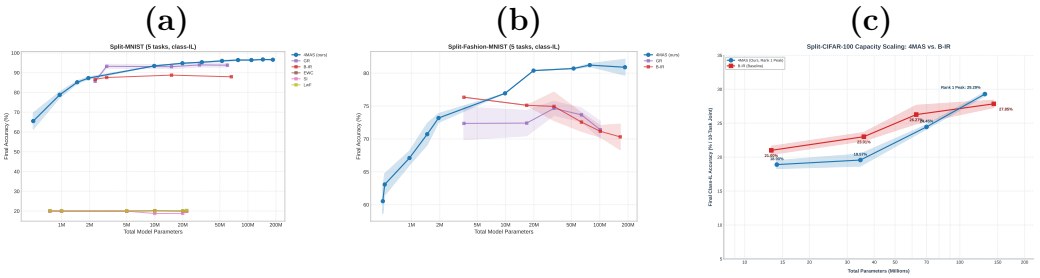

Figure 8: Capacity scaling curves comparing our bilateral 4MAS architecture against the unilateral B-IR baseline across (a) Split-MNIST, (b) Split-Fashion-MNIST, and (c) Split-CIFAR-100. The plots show final accuracy as a function of total model parameters (in millions).

4.6. Ablation studies

To evaluate the contribution of each component, we performed ablations on the STM buffer size, the sleep phase, and the dual-hemisphere structure. Varying the unilateral STM buffer size from 0 to 1,000 samples (Figure 7b) showed diminishing gains beyond 200 samples, improving Split-Fashion-MNIST accuracy from 73.4% (fully ablated STM) to 81.3%. In the absence of a sleep phase, small buffer sizes severely degraded performance (63.2% at size 0), while larger buffers offset this loss, narrowing the sleep phase benefit to $< 1\%$.

Ablation comparisons (Figure 9) show that configurations including the sleep phase consistently perform best on both datasets. Without sleep, the ensemble accuracy matches its single best-performing hemisphere, confirming that sleep-based cross-replay is crucial for bilateral knowledge integration. In particular, the comparative confidence levels of each hemisphere are aligned during sleep training, allowing for system accuracy to exceed either hemisphere's individual accuracy.

5. Limitations and future work

Generative models in this study were standard non-convolutional VAEs, chosen to evaluate task-agnostic macroarchitectures. Future research should scale generators to convolutional VAEs, GANs, or diffusion models to handle more complex image distributions. Additionally, regularisation mechanisms (e.g., EWC or SI) could be integrated in tandem with our replay structure to further boost stability, as demonstrated in hybrid models like B-IR [12].

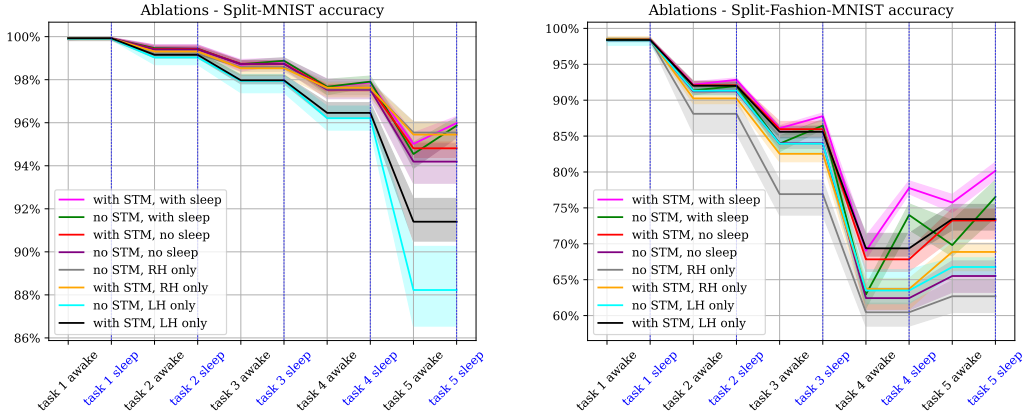

Figure 9: Ablation performance for Split-MNIST and Split-Fashion-MNIST. Configurations including the sleep phase consistently perform best on both datasets, where sleep informed hemispheric confidence levels lead to system accuracy exceeding individual hemisphere accuracies.

While we evaluated 4MAS up to a 50-task Split-CIFAR-100 benchmark, autonomous agents in the real world require handling longer task horizons. Subsequent work should test 4MAS under unbounded dynamic streams, explore diverse forms of hemispheric asymmetry, and investigate additional biological mechanisms, such as slow-wave sleep modelling, to enhance persistent lifelong learning.

6. Conclusion

Lifelong learning in artificial networks remains constrained by catastrophic forgetting. Addressing this challenge, this paper introduced 4MAS, a novel continual learning macroarchitecture derived from biological memory consolidation. By modelling experience replay, bilateral sleep consolidation, and inter-hemispheric lateralisation, 4MAS employs two asymmetric LTM/STM hemispheres that coordinate via a consolidation sleep phase. Our empirical results on Split-MNIST, Split-Fashion-MNIST, and CIFAR-100 demonstrate competitive performance and robust knowledge retention against representational drift. our analysis reveals that functional specialisation is essential to these gains: while symmetrical configurations degraded performance or restricted sleep benefits, asymmetrical parameterisation consistently produced post-sleep performance enhancements, demonstrating that lateralised specialisation improves overall resource utilisation.

This work illustrates how system-level neurobiological structures can be abstracted to manage the stability-plasticity trade-off. Moving beyond local weight-level constraints, system-level bilateral consolidation offers a promising path toward persistent and adaptable artificial intelligence.

References

[1] B. Wickramasinghe, G. Saha, K. Roy, Continual Learning: A Review of Techniques, Challenges, and Future Directions, IEEE Transactions on Artificial Intelligence 5 (6) (2024) 2526–2546. doi:10.1109/TAI.2023.3339091.
URL https://ieeexplore.ieee.org/document/10341211

[2] M. McCloskey, N. J. Cohen, Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem, in: G. H. Bower (Ed.), Psychology of Learning and Motivation, Vol. 24, Academic Press, 1989, pp. 109–165. doi:10.1016/S0079-7421(08)60536-8.
URL https://www.sciencedirect.com/science/article/pii/S0079742108605368

[3] R. Kemker, M. McClure, A. Abitino, T. Hayes, C. Kanan, Measuring Catastrophic Forgetting in Neural Networks (Nov. 2017). doi:10.48550/arXiv.1708.02072.
URL http://arxiv.org/abs/1708.02072

[4] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, D. Hasabis, C. Clopath, D. Kumaran, R. Hadsell, Overcoming catastrophic forgetting in neural networks (2017). doi:10.1073/pnas.1611835114.
URL https://www.pnas.org/doi/10.1073/pnas.1611835114

[5] S. Corkin, What's new with the amnesic patient H.M.?, Nature Reviews Neuroscience 3 (2) (2002) 153–160. doi:10.1038/nrn726.
URL https://www.nature.com/articles/nrn726

[6] B. Giri, H. Miyawaki, K. Mizuseki, S. Cheng, K. Diba, Hippocampal Reactivation Extends for Several Hours Following Novel Experience, The Journal of Neuroscience 39 (5) (2019) 866–875. doi:10.1523/JNEUROSCI.1950-18.2018.

URL https://www.jneurosci.org/lookup/doi/10.1523/JNEUROSCI.1950-18.2018

[7] T. L. Hayes, G. P. Krishnan, M. Bazhenov, H. T. Siegelmann, T. J. Sejnowski, C. Kanan, Replay in Deep Learning: Current Approaches and Missing Biological Elements, Neural Computation (2021) 1–44doi:10.1162/neco_a_01433.
URL https://direct.mit.edu/neco/article/doi/10.1162/neco_a_01433/107071/Replay-in-Deep-Learning-Current-Approaches-and

[8] M. Tamaki, J. Bang, T. Watanabe, Y. Sasaki, Night Watch in One Brain Hemisphere during Sleep Associated with the First-Night Effect in Humans, Current Biology 26 (9) (2016) 1190–1194. doi:10.1016/j.cub.2016.02.063.
URL https://linkinghub.elsevier.com/retrieve/pii/S0960982216301749

[9] M. Ghosh, F.-C. Yang, S. P. Rice, V. Hetrick, A. L. Gonzalez, D. Siu, E. K. Brennan, T. T. John, A. M. Ahrens, O. J. Ahmed, Running speed and REM sleep control two distinct modes of rapid interhemispheric communication, Cell Reports 40 (1) (2022) 111028. doi:10.1016/j.celrep.2022.111028.
URL https://linkinghub.elsevier.com/retrieve/pii/S2211124722008221

[10] L. F. Koziol, The Novelty-Routinization Principle of Brain Organization, in: L. F. Koziol (Ed.), The Myth of Executive Functioning: Missing Elements in Conceptualization, Evaluation, and Assessment, Springer International Publishing, Cham, 2014, pp. 27–31. doi:10.1007/978-3-319-04477-4_8.
URL https://doi.org/10.1007/978-3-319-04477-4_8

[11] C. S. Prat, J. Gallée, B. L. Yamasaki, Getting language right: Relating individual differences in right hemisphere contributions to language learning and relearning, Brain and Language 239 (2023) 105242. doi:10.1016/j.bandl.2023.105242.
URL https://www.sciencedirect.com/science/article/pii/S0093934X23000214

[12] G. M. van de Ven, H. T. Siegelmann, A. S. Tolias, Brain-inspired replay for continual learning with artificial neural networks, Nature Communications 11 (1) (2020) 4069. doi:10.1038/s41467-020-17866-2.
URL https://www.nature.com/articles/s41467-020-17866-2

[13] G. M. v. d. Ven, A. S. Tolias, Three scenarios for continual learning (Apr. 2019). doi:10.48550/arXiv.1904.07734.
URL http://arxiv.org/abs/1904.07734

[14] S. Gutstein, O. Fuentes, E. Freudenthal, Knowledge transfer in deep convolutional neural nets, International Journal on Artificial Intelligence Tools 17 (03) (2008) 555–567. doi:10.1142/S0218213008004059.
URL https://www.worldscientific.com/doi/abs/10.1142/S0218213008004059

[15] F. Zenke, B. Poole, S. Ganguli, Continual Learning Through Synaptic Intelligence, Proceedings of machine learning research (2017).

[16] R. Aljundi, F. Babiloni, M. Elhoseiny, M. Rohrbach, T. Tuytelaars, Memory Aware Synapses: Learning what (not) to forget, in: Proceedings of the European conference on computer vision (ECCV), 2018, pp. 139–154.
URL https://openaccess.thecvf.com/content_ECCV_2018/html/Rahaf_Aljundi_Memory_Aware_Synapses_ECCV_2018_paper.html

[17] S. I. Mirzadeh, M. Farajtabar, R. Pascanu, H. Ghasemzadeh, Understanding the Role of Training Regimes in Continual Learning, Advances in Neural Information Processing Systems 33 (2020) 7308–7320.
URL https://proceedings.neurips.cc/paper/2020/hash/518a38cc9a0173d0b2dc088166981cf8-Abstract.html?ref=https://githubhelp.com

[18] G. Shi, J. Chen, W. Zhang, L.-M. Zhan, X.-M. Wu, Overcoming Catastrophic Forgetting in Incremental Few-Shot Learning by Finding Flat Minima, Advances in neural information processing systems 34 (2021).

[19] Y. Kong, L. Liu, H. Chen, J. Kacprzyk, D. Tao, Overcoming Catastrophic Forgetting in Continual Learning by Exploring Eigenvalues of Hessian Matrix, IEEE Transactions on Neural Networks and Learning Systems 35 (11) (2024) 16196–16210. doi:10.1109/TNNLS.2023.

3292359.
URL https://ieeexplore.ieee.org/document/10190202

[20] M. D. Lange, R. Aljundi, M. Masana, S. Parisot, X. Jia, A. Leonardis, G. Slabaugh, T. Tuytelaars, A continual learning survey: Defying forgetting in classification tasks, IEEE Transactions on Pattern Analysis and Machine Intelligence (2021) 1–1ArXiv:1909.08383 [cs]. doi:10.1109/TPAMI.2021.3057446.
URL http://arxiv.org/abs/1909.08383

[21] A. S. Razavian, H. Azizpour, J. Sullivan, S. Carlsson, CNN Features off-the-shelf: an Astounding Baseline for Recognition (May 2014). doi:10.48550/arXiv.1403.6382.
URL http://arxiv.org/abs/1403.6382

[22] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, T. Darrell, DeCAF: A Deep Convolutional Activation Feature for Generic Visual Recognition (Oct. 2013). doi:10.48550/arXiv.1310.1531.
URL http://arxiv.org/abs/1310.1531

[23] A. V. Terekhov, G. Montone, J. K. O'Regan, Knowledge Transfer in Deep Block-Modular Neural Networks, in: S. P. Wilson, P. F. Verschure, A. Mura, T. J. Prescott (Eds.), Biomimetic and Biohybrid Systems, Springer International Publishing, Cham, 2015, pp. 268–279. doi:10.1007/978-3-319-22979-9_27.

[24] A. A. Rusu, N. C. Rabinowitz, G. Desjardins, H. Soyer, J. Kirkpatrick, K. Kavukcuoglu, R. Pascanu, R. Hadsell, Progressive Neural Networks (Oct. 2022). doi:10.48550/arXiv.1606.04671.
URL http://arxiv.org/abs/1606.04671

[25] S. Grossberg, Adaptive Resonance Theory: How a brain learns to consciously attend, learn, and recognize a changing world, Neural Networks 37 (2013) 1–47. doi:10.1016/j.neunet.2012.09.017.
URL https://www.sciencedirect.com/science/article/pii/S0893608012002584

[26] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. Le, G. Hinton, J. Dean, Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer (Jan. 2017). doi:10.48550/arXiv.1701.

06538.
URL http://arxiv.org/abs/1701.06538

[27] Z. Chen, A. Wuerkaixi, S. Cui, H. Li, D. Li, J. Zhang, B. Han, G. Niu, H. Liu, Y. Yang, S. Yang, C. Zhang, T. Ren, Learning without Isolation: Pathway Protection for Continual Learning, arXiv:2505.18568 [cs] (May 2025). doi:10.48550/arXiv.2505.18568.
URL http://arxiv.org/abs/2505.18568

[28] N. Omi, S. Sen, A. Farhadi, Load Balancing Mixture of Experts with Similarity Preserving Routers, arXiv:2506.14038 [cs] (Oct. 2025). doi:10.48550/arXiv.2506.14038.
URL http://arxiv.org/abs/2506.14038

[29] D. Lopez-Paz, M. A. Ranzato, Gradient Episodic Memory for Continual Learning, Advances in Neural Information Processing Systems 30 (2017).
URL https://proceedings.neurips.cc/paper/2017/hash/f87522788a2be2d171666752f97ddebb-Abstract.html

[30] R. Aljundi, E. Belilovsky, T. Tuytelaars, L. Charlin, M. Caccia, M. Lin, L. Page-Caccia, Online Continual Learning with Maximal Interfered Retrieval, Advances in Neural Information Processing Systems 32 (2019).
URL https://proceedings.neurips.cc/paper/2019/hash/15825aee15eb335cc13f9b559f166ee8-Abstract.html

[31] X. Li, B. Tang, H. Li, AdaER: An adaptive experience replay approach for continual lifelong learning, Neurocomputing 572 (2024) 127204. doi:10.1016/j.neucom.2023.127204.
URL https://www.sciencedirect.com/science/article/pii/S0925231223013279

[32] Y. Balaji, M. Farajtabar, D. Yin, A. Mott, A. Li, The Effectiveness of Memory Replay in Large Scale Continual Learning (Oct. 2020). doi:10.48550/arXiv.2010.02418.
URL http://arxiv.org/abs/2010.02418

[33] H. Shin, J. K. Lee, J. Kim, J. Kim, Continual Learning with Deep Generative Replay (Dec. 2017). doi:10.48550/arXiv.1705.08690.
URL http://arxiv.org/abs/1705.08690

[34] X. Liu, C. Wu, M. Menta, L. Herranz, B. Raducanu, A. D. Bagdanov, S. Jui, J. v. d. Weijer, Generative Feature Replay For Class-Incremental Learning (Apr. 2020). doi:10.48550/arXiv.2004.09199.
URL http://arxiv.org/abs/2004.09199

[35] E. T. Rolls, X. Yan, G. Deco, Y. Zhang, V. Jousmaki, J. Feng, A ventromedial visual cortical 'Where' stream to the human hippocampus for spatial scenes revealed with magnetoencephalography, Communications Biology 7 (1) (2024) 1–16. doi:10.1038/s42003-024-06719-z.
URL https://www.nature.com/articles/s42003-024-06719-z

[36] V. Khan, S. Cygert, K. Deja, T. Trzcinski, B. Twardowski, Looking Through the Past: Better Knowledge Retention for Generative Replay in Continual Learning, IEEE Access 12 (2024) 45309–45317. doi:10.1109/ACCESS.2024.3379148.
URL https://ieeexplore.ieee.org/document/10474374

[37] Y. Lecun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, Proceedings of the IEEE 86 (11) (1998) 2278–2324. doi:10.1109/5.726791.
URL https://ieeexplore.ieee.org/document/726791/

[38] G. M. Van De Ven, T. Tuytelaars, A. S. Tolias, Three types of incremental learning, Nature Machine Intelligence 4 (12) (2022) 1185–1197. doi:10.1038/s42256-022-00568-3.
URL https://www.nature.com/articles/s42256-022-00568-3

[39] H. Xiao, K. Rasul, R. Vollgraf, Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms (Sep. 2017). doi:10.48550/arXiv.1708.07747.
URL http://arxiv.org/abs/1708.07747

[40] G. Sokar, D. C. Mocanu, M. Pechenizkiy, SpaceNet: Make Free Space For Continual Learning, Neurocomputing 439 (2021) 1–11, arXiv:2007.07617 [cs]. doi:10.1016/j.neucom.2021.01.078.
URL http://arxiv.org/abs/2007.07617

[41] N. Niethard, M. Hasegawa, T. Itokazu, C. N. Oyanedel, J. Born, T. R. Sato, Sleep-Stage-Specific Regulation of Cortical Excitation and Inhibition, Current Biology 26 (20) (2016) 2739–2749.

doi:10.1016/j.cub.2016.08.035.
URL https://www.cell.com/current-biology/abstract/S0960-9822(16)30979-4

[42] P. Skierś, K. Deja, Joint Diffusion models in Continual Learning (Oct. 2025). doi:10.48550/arXiv.2411.08224.
URL http://arxiv.org/abs/2411.08224

[43] F. Zhou, T. Zhao, L. V. Nguyen, Z. Yao, A Parallel Gumbel-Softmax VAE Framework with Performance-Based Tuning, in: ECAI 2024, IOS Press, 2024, pp. 1784–1791. doi:10.3233/FAIA240689.
URL https://ebooks.iospress.nl/doi/10.3233/FAIA240689

[44] E. Goldberg, A New Look at the Old Riddle : Novelty, Routines and the Evolution of the Bicameral Brain 20 (3) (2018).

[45] A. D. Grosmark, F. T. Sparks, M. J. Davis, A. Losonczy, Reactivation predicts the consolidation of unbiased long-term cognitive maps, Nature Neuroscience 24 (11) (2021) 1574–1585. doi:10.1038/s41593-021-00920-7.
URL https://www.nature.com/articles/s41593-021-00920-7

Appendix A. Hyperparameter Searches and Optimisation

This appendix provides detail on the hyperparameter tuning sweeps conducted to optimise the 4MAS architecture across the three benchmarks: Split-MNIST, Split-Fashion-MNIST, and Split-CIFAR-100. These searches cover the grid-search sweeps logged in the main project database (comprising 4,328 unique trials) as well as the sequential Bayesian optimisation sweeps conducted using the Optuna framework. Hyperparameter sweeps reported in this Appendix are single runs.

Appendix A.1. Split-MNIST and Split-Fashion-MNIST Grid Searches

For both Split-MNIST and Split-Fashion-MNIST, we ran systematic parameter grids to evaluate the interaction of model size, latent representation dimensions, and generator capacity under our bilateral sleep-consolidation model.

The range of hyperparameters explored in these grid searches is summarised in Table A.4.

Table A.4: Hyperparameter search space for Split-MNIST and Split-Fashion-MNIST.

Hyperparameter	**Explored Values**
Model Capacity (Parameters)	248K, 477K, 736K, 970K, 4.9M, 9.9M, 11.8M, 16.0M, 26.3M, 39.2M, 54.5M, 72.4M, 92.7M
Short-Term Memory Size (N_{stm})	50 per hemisphere (10 samples per class)
VAE Generator Dimension	140, 250, 360, 450, 1450, 2250, 2500, 3000, 4000, 5000, 6000, 7000, 8000
VAE Latent Space Dimension	16, 32, 64, 128, 256
Awake-Phase Learning Rate (LR_{awake})	2×10^{-4} (Adam)
Sleep-Phase Learning Rate (LR_{sleep})	2×10^{-5} (Adam)
Awake-Phase Gating End-Weight	10.0
Left Generator Temperature (τ_L)	1.0
Right Generator Temperature (τ_R)	3.0
Oversampling Factor (O)	10, 50

Large model capacity paired with a larger VAE dimension and low latent dimensionality (16) consistently yielded optimal Class-IL accuracy by preserving high-fidelity reconstructions without latent-space drift.

Appendix A.2. Split-CIFAR-100 Optuna Optimisation Sweeps

Due to the increased complexity of CIFAR-100, we executed 14 distinct optimisation sweeps using the Optuna framework, focusing on architectural changes, pre-training steps, and latent alignment strategies. The best-performing trial configuration for each sweep is documented in Table A.5.

Appendix A.3. Split-CIFAR-100 Architectural Grid Search

To clarify the contribution of supervised contrastive loss, left-hemisphere parameter freezing, and bilateral sleep-consolidation (BSS) active gating, we conducted a 16-configuration grid search on Split-CIFAR-100. The complete experimental results of this grid search are summarised in Table A.6.

Appendix A.4. Split-CIFAR-100 Ablation and Asymmetry Trials

To systematically investigate the roles of comparative confidence selection (CCS), latent space cluster centroids (LSCC), pre-training tasks (1 vs. 5), and synaptic intelligence (SI) regularisation on the Split-CIFAR-100 benchmark, we evaluated several design configurations. The results are summarised in Table A.7.

Appendix A.5. Split-CIFAR-100 Architecture Capacity and Parameter Scaling

To investigate the capacity limits and scaling robustness of 4MAS relative to B-IR, we evaluated both architectures across a range of total parameter capacities (from 14M to 140M parameters). The comparative results are summarized in Table A.8.

As shown in Table A.8, when properly initialized with pre-trained convolutional features (using the '–pre-convE' flag), the B-IR baseline scales robustly and monotonically with parameter capacity, rising from 21.0% (13.3M scale) to 27.85% (143.8M scale). Similarly, 4MAS scales monotonically and robustly across all parameter scales, rising from 13.16% (14M scale) to 18.41% (35M scale), $24.45 \pm 0.32\%$ (70M scale), and achieving a peak accuracy of **29.26%** at 130.8M capacity (`fb=1.0, mb=20.0, bg=0.90`) under strict single-task pre-training constraints, outperforming the unilateral B-IR baseline as well as the 147.8M v5 configuration (29.01%) which utilized multi-task pre-training. This breakthrough is achieved by resolving representation anchoring bottlenecks at scale via an expanded rehearsal buffer (`stmsize` = 4000). Importantly, 4MAS maintains significantly lower forgetting ($25.8 \pm 1.2\%$ at 70M scale and 23.4% at 140M scale) than B-IR ($52.82 \pm 1.97\%$ and 52.92%), verifying that bilateral sleep consolidation consistently provides superior retention of past task knowledge across all capacity scales.

Appendix A.6. VAE Hyperparameter Seed Replications and Task Trajectories

To verify the statistical consistency and trajectory stability of top-performing VAE hyperparameter configurations and Quality-Gated rehearsal regimes, multiple random seed replications were evaluated across all 10 tasks on Split-CIFAR-100. Figure A.10 illustrates the post-REM joint accuracy trajectory as a function of class scale (10 to 100 classes) with standard deviation error bars.

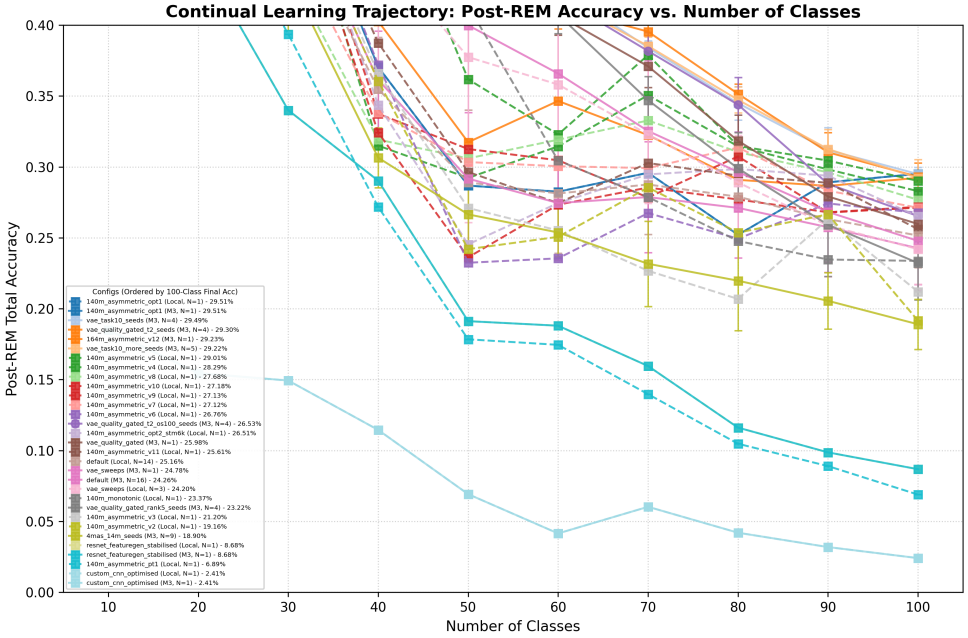

Figure A.10: Continual learning post-REM accuracy trajectories across class scale (10 to 100 classes) for top VAE configurations and Quality-Gated replay regimes, showing mean performance and standard deviation error bars across random seeds.

Table A.5: Best configurations and parameters across Split-CIFAR-100 Optuna sweeps.

Optuna Study Name	Trials	Best Class-IL (%)	Key Parameters and Configurations
fourmas_cifar100_v1	89	0.00	Default VAE, latent dim 64, LR 0.0014, SGD
fourmas_cifar100_v2	34	28.82	Task-IL scenario, latent dim 128, buffer 2000, Adam
fourmas_cifar100_v3	10	4.11	Class-IL, buffer 5000, VAE temp 1.15 (LH) / 1.54 (RH)
fourmas_cifar100_v4	26	7.63	Class-IL, buffer 5000, SWS/REM epochs 75/50, LR 0.0022
fourmas_cifar100_multi_arch	4	3.77	Pre-trained tap layer 2 feature extractor, latent dim 64
fourmas_cifar100_scratch_sota	15	6.85	Scratch training (no ImageNet), buffer 10000, latent dim 128
fourmas_cifar100_novel_designs	24	9.54	Evaluated classifier freezing policies, distill weight 0.16
fourmas_cifar100_refined_sota	17	9.49	VAE dim 400, latent dim 64, gating end-weight 2.0
fourmas_cifar100_final_sota	25	9.56	Optimised lr 0.0023, REM/SWS epochs 150/100
fourmas_cifar100_fixed_sota	13	9.99	Cross-entropy loss function, linear classifier, LR 0.0018
fourmas_cifar100_sota_push	34	11.06	Latent dim 192, gating end-weight 1.5, LR 0.0019
fourmas_cifar100_sota_breakthrough_v2	25	9.56	VAE dim 800, latent dim 192, gating end-weight 1.0
fourmas_cifar100_pretrain_multi_task	16	12.25	Bootloaded multi-task pre-training (1 task), latent dim 192
fourmas_cifar100_vqvae	62	10.35	Vector Quantized VAE, gating end-weight 2.0, LR 0.0017
fourmas_cifar100_vqvae_high_cap	32	9.47	High capacity VQ-VAE, codebook size 2048, latent dim 128

Table A.6: Grid search performance over generative model types, supervised contrastive (SupCon) weights, left-hemisphere parameter freezing, and BSS active gating.

ID	Model Type	SupCon Weight	LH Freeze	BSS Active	Class-IL (%)	Task-IL (%)
1	vae	0.0	False	False	1.52	12.04
2	vae	0.0	False	True	1.02	9.61
3	vae	0.0	True	False	1.19	11.17
4	vae	0.0	True	True	1.13	9.26
5	vae	0.5	False	False	**2.13**	**17.33**
6	vae	0.5	False	True	2.01	14.90
7	vae	0.5	True	False	1.03	10.32
8	vae	0.5	True	True	1.06	10.51
9	vqvae	0.0	False	False	1.41	10.31
10	vqvae	0.0	False	True	1.07	14.19
11	vqvae	0.0	True	False	1.11	10.45
12	vqvae	0.0	True	True	1.08	9.38
13	vqvae	0.5	False	False	1.52	12.41
14	vqvae	0.5	False	True	1.00	10.87
15	vqvae	0.5	True	False	1.04	10.47
16	vqvae	0.5	True	True	1.01	9.96

Table A.7: Split-CIFAR-100 performance across selection strategies (CCS, LSCC, Hybrid), pre-training task horizons, and local SI regularisation.

Selection Strategy (LH / RH)	Left Centroid (L)	Right Centroid (R)	Pre-training Tasks	Synaptic Intelligence (SI)	Class-IL Acc (%)
CCS (Symmetric Baseline)	0.5	0.5	1	No	21.38
CCS (Symmetric)	0.5	0.5	2	No	25.34
CCS (Symmetric)	0.5	0.5	3	No	29.18
CCS (Symmetric)	0.5	0.5	4	No	**30.46**
CCS (Symmetric)	0.5	0.5	5	No	29.01
CCS (Left-Skewed)	0.7	0.3	1	No	21.78
CCS (Right-Skewed)	0.3	0.7	1	No	23.80
CCS (Asymmetric Champion)	0.1	0.5	1	No	25.97
CCS (Asymmetric Low)	0.1	0.1	1	No	22.96
CCS (Asymmetric High)	0.2	0.8	1	No	22.34
CCS (Sweeter Spot)	0.05	0.55	1	No	25.63
CCS (Asymmetric + LH SI)	0.1	0.5	1	Yes	16.27
LSCC (Symmetric K-Means)	K-Means	K-Means	1	No	8.79
LSCC (Symmetric K-Means)	K-Means	K-Means	5	No	10.39
Hybrid (LSCC / CCS Asymmetric)	K-Means	0.5	1	No	13.20
Hybrid (LSCC / CCS Asymmetric)	K-Means	0.5	2	No	17.58
Hybrid (LSCC / CCS Asymmetric)	K-Means	0.5	3	No	20.09
Hybrid (LSCC / CCS Asymmetric)	K-Means	0.5	4	No	18.72
Hybrid (LSCC / CCS Asymmetric)	K-Means	0.5	5	No	**22.75**
Hybrid (LSCC / CCS, Seed 1001)	K-Means	0.5	5	No	20.14

Table A.8: Split-CIFAR-100 Class-IL performance across parameter capacity scales.

Model Size	**Architecture**	**Total Parameters**	**Class-IL Acc (%)**	**Forgetting (%)**
14M Scale	4MAS (Ours, Scaled-down)	14.1M	13.16	**25.8**
	B-IR (Standard Baseline)*	13.3M	**21.0**	61.1
35M Scale	4MAS (Ours, Scaled-down)	34.5M	18.41	**25.8**
	B-IR (Ours, Scaled-down)	35.8M	**23.01**	58.9
70M Scale	4MAS (Ours, Champion)	70.0M	24.45 ± 0.32	25.8 ± 1.2
	B-IR (Ours, Scaled)	62.7M	26.27 ± 1.37	52.82 ± 1.97
140M Scale	4MAS (Ours, Rank 1 Peak)	130.8M	29.29 ± 0.29	25.8
	4MAS (Ours, Peak v5)*	147.8M	29.01	**23.4**
	4MAS (Ours, v10)	130.8M	27.18	73.7
	4MAS (Ours, v9)	130.8M	27.13	51.0
	B-IR (Ours, Scaled)	143.8M	27.85	52.92