

HiCL: Hippocampal-Inspired Continual Learning

Kushal Kapoor,¹ Wyatt Mackey,² Yiannis Aloimonos,¹ Xiaomin Lin^{1,3,4}

¹Perception and Robotics Group, University of Maryland

²DEVCOM Army Research Lab (ARL)

³Department of Electrical and Computer Engineering, Johns Hopkins University

⁴Department of Electrical Engineering, University of South Florida

kushalk@terpmail.umd.edu, wyatt.t.mackey.civ@army.mil, jyaloimo@umd.edu, xlin2@usf.edu

Abstract

We propose HiCL, a novel hippocampal-inspired dual-memory continual learning architecture designed to mitigate catastrophic forgetting by using elements inspired by the hippocampal circuitry. Our system encodes inputs through a grid-cell-like layer, followed by sparse pattern separation using a dentate gyrus-inspired module with top- k sparsity. Episodic memory traces are maintained in a CA3-like autoassociative memory. Task-specific processing is dynamically managed via a DG-gated mixture-of-experts mechanism, wherein inputs are routed to experts based on cosine similarity between their normalized sparse DG representations and learned task-specific DG prototypes computed through online exponential moving averages. This biologically grounded yet mathematically principled gating strategy enables differentiable, scalable task-routing without relying on a separate gating network, and enhances the model’s adaptability and efficiency in learning multiple sequential tasks. Cortical outputs are consolidated using Elastic Weight Consolidation weighted by inter-task similarity. Crucially, we incorporate prioritized replay of stored patterns to reinforce essential past experiences. Evaluations on standard continual learning benchmarks demonstrate the effectiveness of our architecture in reducing task interference, achieving near state-of-the-art results in continual learning tasks at lower computational costs.

Code — <https://github.com/kushalk173-sc/HiCL>

Extended version — <https://arxiv.org/pdf/2508.16651>

Introduction

Artificial neural networks excel at many single-task benchmarks but struggle when trained sequentially on multiple tasks, suffering from *catastrophic forgetting*—the tendency to overwrite previous knowledge with new information. In contrast, the human brain learns continuously throughout life, preserving decades of memories while acquiring new skills.

In continual learning, every new task is like a distinct episode in our life’s unfolding story. Just as the hippocampus rapidly encodes unique experiences into *episodic memory*, tagging them with contextual and temporal cues, a continual learner must capture each task’s defining features before

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

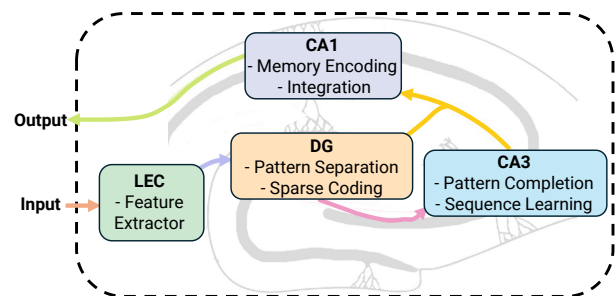


Figure 1: Schematic of the **HiCL architecture** overlaid with the hippocampal trisynaptic circuit. The entorhinal cortex (layers II/III) projects to the dentate gyrus (DG), which performs pattern separation. DG outputs feed into CA3, an autoassociative network responsible for pattern completion. CA3 then projects to CA1, which integrates inputs from both CA3 and the entorhinal cortex before sending output to other brain regions. This circuit underlies memory encoding and retrieval processes. The general background here of the hippocampal circuit is based on (Petranonakis and Poirazi 2014) and adapted from Moser, Kropff, and Moser (2008); Leutgeb et al. (2007); Rolls (2013b); Otmakhova, Otmakhov, and Griffith (2018).

they fade. The brain then *replays* these episodes—during sleep or quiet wakefulness—to gradually integrate them into neocortical networks, preventing older memories from being overwritten by new ones. We leverage this *complementary learning systems* principle by treating each task as an episodic trace: we store compact representations in a replay buffer and periodically revisit them during training, mirroring how hippocampal replay consolidates previous and precious episodes into lasting knowledge (McClelland, McNaughton, and O’Reilly 1995).

We draw on three core hippocampal subregions—the dentate gyrus (DG), CA3, and CA1—to inspire a *DG-Gated Mixture-of-Experts (MoE)* architecture (as shown in Figure 1, where each subcircuit is instantiated as an expert and inputs are dynamically routed based on their DG activations). This DG-based gating, eliminates the need for task labels at inference and directly tackles catastrophic forgetting.

Dentate Gyrus (DG): In biology, the DG performs *pattern separation* by mapping inputs onto a very sparse, typically only 3–5% of granule cells fire for any given pat-

tern (Hainmueller and Bartos 2020). In our model, the `SparseActivation` DG layer enforces top- k sparsity ($k = 5\%$), orthogonalizing features both for expert representations and for routing signals.

CA3: The DG’s orthogonal codes feed into CA3’s recurrent attractor network, which functions as *pattern completion*, reconstructing full memory traces from partial cues (Gilbert and Brushfield 2009). We mirror this with a lightweight two-layer MLP that refines and transforms DG outputs via a non-linear projection, functionally completing the pattern for downstream processing.

CA1: Finally, CA1 integrates completed patterns from CA3 with direct cortical inputs, acting as a comparator and coordinating memory consolidation through hippocampo-cortical dialogue (Bartsch et al. 2011). We emulate this via our `calintegration` block and a consolidation stage combining Elastic Weight Consolidation (EWC) with a replay buffer.

Bringing it together: By instantiating N such experts and routing via cosine similarities between live DG codes and stored DG prototypes—using either hard or soft-max gating—and then sharpening DG separation with a contrastive tuning phase.

This biology-inspired solution to the stability–plasticity dilemma offers actionable insights for continual learning in artificial systems. Our *DG-Gated Mixture-of-Experts* architecture translates hippocampal principles into three core mechanisms: (1) a top- k sparse DG layer ($k \approx 5\%$) that orthogonalizes inputs for both expert encoding and dynamic routing; (2) a CA3-inspired two-layer MLP autoencoder that completes partial patterns via attractor-like reconstruction; and (3) a consolidation module combining a prioritized replay buffer with Elastic Weight Consolidation to protect prior knowledge. By instantiating the hippocampal tripartite circuit across multiple experts and leveraging DG-based gating to minimize interference, CA3-style pattern completion for robust recall, and replay+EWC for selective consolidation, our DG-Gated MoE exemplifies how neuroscience-informed architectures can achieve efficient, adaptive continual learning.

Our primary contributions are:

- We **introduce HiCL, a DG-gated Mixture-of-Experts (MoE) model that routes inputs without a gating network**. By leveraging the sparse representations from a Dentate Gyrus-inspired layer, our model dynamically assigns tasks to experts based on cosine similarity to learned prototypes.
- We **developed a training strategy that explicitly improves pattern separation**. Our two-phase approach first trains individual experts using EWC, replay, and feature distillation, then applies a global contrastive loss to all DG layers, forcing them to learn more distinct representations for their respective tasks.
- We **achieved competitive accuracy at a fraction of the computational cost**. On benchmarks like Split CIFAR-10, our model demonstrates that a neuroscience-aligned architecture can be both computationally effective and efficient in a continual learning setting.

Related Work

Continual learning (CL) aims to develop systems that learn from a non-stationary data stream without catastrophic forgetting. In this section, we review key literature spanning computational neuroscience and machine learning, focusing on complementary learning systems, hippocampal memory motifs, and modern CL strategies.

Complementary Learning Systems (CLS). The CLS theory posits a dual-memory system: a fast hippocampal module for episodic learning and a slow neocortical system for generalization (McClelland, McNaughton, and O’Reilly 1995; Kemker et al. 2018). This foundational concept has influenced many biologically inspired CL frameworks. Examples include hybrid cortico-hippocampal networks that combine feedforward and spiking layers (Rothschild, Eban, and Frank 2017), DualNet’s distinct fast/slow learners (Pham, Liu, and Hoi 2021), FearNet’s use of generative replay (Kemker and Kanan 2017), and Triple Memory Networks that model hippocampus-PFC interactions with replay and consolidation (Wang et al. 2021).

Hippocampal Replay. Replay of past experiences is a powerful anti-forgetting mechanism, mirroring the brain’s hippocampal replay during sleep and wakefulness (Cairney et al. 2018). In machine learning, this is realized through **experience replay**, which explicitly rehearses stored examples (Rolnick et al. 2019), or **generative replay**, which trains a network to synthesize pseudo-data (Shin et al. 2017). We adopt prioritized experience replay, inspired by findings that emphasizing important memories enhances learning efficiency in both reinforcement learning (Schaul et al. 2016) and biological systems (Mattar and Daw 2018). The ability to generate contrastive data during replay is considered crucial for preventing representational collapse (Hinton 2022).

Modern CL Architectures and Strategies. Many CL frameworks are categorized as **parameter isolation**, **rehearsal based**, or **regularization based** (De Lange et al. 2021). **Regularization** methods add a loss penalty to prevent changes to weights important for past tasks, as seen in Elastic Weight Consolidation (EWC) (Kirkpatrick et al. 2017), Learning without Forgetting (LwF) (Li and Hoiem 2017), and Synaptic Intelligence (SI) (Zenke, Poole, and Ganguli 2017). **Parameter isolation** methods assign distinct weights to different tasks, using techniques such as network expansion in Progressive Nets (Rusu et al. 2016) or hard attention masks (Serrà et al. 2018; Madotto et al. 2021; Wang et al. 2024). **Rehearsal** methods use a buffer of past data (Rebuffi et al. 2017; Riemer et al. 2018; Zhang, Wang, and Yang 2022) or a generative model (Shin et al. 2017; Su et al. 2019; Skiers̄ and Deja 2025). Hybrid models such as Progress & Compress (Schwarz et al. 2018) and CLEAR (Riemer et al. 2018) combine these strategies effectively, and our approach synergizes these ideas by using architectural modularity, sparsity, and replay.

Mixture-of-Experts (MoE) and Gating. MoE models enable modular specialization by routing inputs to specialized subnetworks (Yuksel, Wilson, and Gader 2012; Shazeer et al. 2017), a principle recently applied to CL (Li et al.

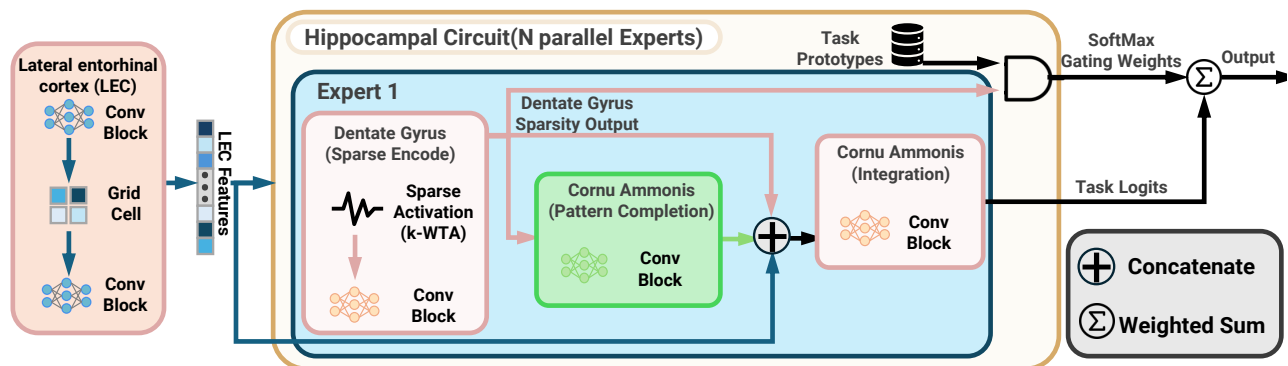


Figure 2: Overview of the HiCL architecture aligned with the hippocampal trisynaptic circuit: EC II/III → DG (pattern separation) → CA3 (pattern completion) → CA1 (integration).

2025; Le et al. 2024) and large-scale models (Guo et al. 2025; Cai et al. 2025). Stability is often improved via adaptive routing, for instance by annealing a routing temperature (Nie et al. 2022; Fedus, Zoph, and Shazeer 2022). Our DG-Gated MoE is distinct in that expert selection emerges via cosine similarity to prototypes, bypassing an explicit learned gating network.

Prototype-Based Replay. Storing class prototypes instead of raw data is a memory-efficient rehearsal strategy pioneered by iCaRL (Rebuffi et al. 2017). Subsequent work has combined prototype anchoring with contrastive objectives to preserve the geometry of the feature space across tasks (Caccia et al. 2021; Aghasanli, Li, and Angelov 2025). We follow this line, maintaining online exponential moving average (EMA) prototypes and applying a push-pull contrastive loss during replay.

These prior threads motivate a more tightly coupled design: rather than treating biologically inspired modules as loose metaphors, we embed them directly into representation, routing, and training dynamics in a unified hippocampal-inspired dual-memory MoE.

Biologically Inspired Motifs and Dual-Memory Training

Our design is grounded in the hippocampal trisynaptic circuit (DG → CA3 → CA1) and the theory of complementary learning systems (CLS) (McClelland, McNaughton, and O’Reilly 1995; Kemker et al. 2018). We operationalize these motifs as inductive scaffolds: structured encoding, sparse separation, nonlinear refinement/completion, and integration, which together feed into prototype-based routing and consolidation. Replay-inspired mechanisms and prototype anchoring provide retention and interference mitigation, while modular gating arises naturally from the quasi-orthogonal codes produced by DG-like separation. Learning is organized as a two-phase dual-memory schedule—specialization followed by consolidation—whose detailed losses are defined in the Architecture section.

Dual-Memory Realization. We operationalize the complementary learning system as a two-phase training schedule: **Phase I** rapidly specializes, forming and updating task prototypes and shaping DG gating, while **Phase II** consolidates

prior knowledge via contrastive alignment, mitigating interference and preventing degradation.

Grid Cell Encoding

Grid-cell-like representations from the entorhinal cortex provide a structured relational prior that has been shown to enhance spatial and contextual reasoning in downstream tasks (Hafting et al. 2005; Banino et al. 2018). We implement this by applying $M = 4$ parallel 1×1 convolutions with learned phase offsets to backbone features \mathbf{f} :

$$\mathbf{g}_m = \sin(\mathbf{W}_m \mathbf{f} + \phi_m), \quad m = 1, \dots, 4, \mathbf{g} = [\mathbf{g}_1; \dots; \mathbf{g}_4].$$

This hexagonal-style embedding enhances discrimination and imposes structured relational priors before gating, similar in spirit to positional priors used in other continual and representation learning works.

Dentate Gyrus: Sparse Pattern Separation

To reduce interference between tasks, we draw on the role of the dentate gyrus in orthogonalizing inputs via sparse coding (Leutgeb et al. 2007; Rolls 2013a). The encoded representation \mathbf{g} is projected and sparsified:

$$\begin{aligned} \mathbf{z} &= \text{ReLU}(\mathbf{W}_{\text{DG}} \mathbf{g} + \mathbf{b}_{\text{DG}}), \\ \mathbf{p}_{\text{sep}} &= \text{TopK}(\text{LayerNorm}(\mathbf{z}), k), \quad k = \lfloor 0.05 \dim(\mathbf{z}) \rfloor. \end{aligned} \quad (1)$$

This top- k mechanism approximates inhibitory microcircuitry to produce quasi-orthogonal sparse codes, creating natural modularization and enabling routing without explicit task labels, akin to capacity partitioning in MoE and gating literature (Yuksel, Wilson, and Gader 2012; Shazeer et al. 2017).

CA3: Nonlinear Refinement / Completion

The sparse separation codes are refined through a feed-forward MLP to approximate the pattern completion function attributed to CA3 (Treves and Rolls 1994). Unlike classical recurrent attractor implementations (Hopfield 1982), we use:

$$\mathbf{p}_{\text{comp}} = \text{LN}(\text{ReLU}(\mathbf{W}_2 \text{ReLU}(\mathbf{W}_1 \mathbf{p}_{\text{sep}} + \mathbf{b}_1) + \mathbf{b}_2))$$

providing nonlinear task-specific refinement while maintaining stability in the face of non-i.i.d. data (Alonso and Krichmar 2024; Krotov and Hopfield 2016).

CA1: Integration

Separation and completion are fused to form a consolidated representation:

$$\mathbf{u} = [\mathbf{p}_{\text{sep}}; \mathbf{p}_{\text{comp}}],$$

which serves as the basis for both prediction and prototype anchoring. This mirrors CA1’s role in integrating pattern-separated and completed signals before driving downstream systems.

Biological Abstraction.

Our design philosophy emphasizes **principled abstraction**, extracting core computational principles from hippocampal circuits while trading biological fidelity for computational efficiency and trainability. The biological DG employs complex inhibitory microcircuitry to enforce sparse coding (Hainmueller and Bartos 2020). We replace this with a non-parametric TopK operation that, when combined with Phase II contrastive tuning, achieves effective pattern separation without simulating neuronal dynamics. Similarly, CA3’s recurrent attractor network (Treves and Rolls 1994) is approximated by a feedforward MLP that captures pattern completion functionality in a single stable forward pass, avoiding the computational cost and training instability of recurrent architectures (Hopfield 1982). Our learnable GridCellLayer adapts its spatial encoding to dataset geometry rather than using fixed Fourier bases, providing representational flexibility for the continual learning setting. These design choices demonstrate that capturing biological *function* enables effective continual learning without requiring literal anatomical replication.

Prototype Maintenance and Replay Motivation

Each expert maintains a prototype \mathbf{u}_i updated during specialization via exponential moving average of its DG separation code:

$$\mathbf{u}_i \leftarrow (1 - \mu) \mathbf{u}_i + \mu \mathbf{p}_{\text{sep}}^{(i)}, \quad \mu = 0.01.$$

This echoes prototype-based rehearsal approaches (Rebuffi et al. 2017; Caccia et al. 2021) and supplies stable references for both gating and later consolidation. The use of replay-like alignment and contrastive objectives (detailed in Section) is inspired by hippocampal replay phenomena that reinforce important memories during offline phases (Cairney et al. 2018; Schaul et al. 2016; Mattar and Daw 2018).

The Architecture

Our system is a DG-Gated Mixture-of-Experts (MoE) designed to mitigate catastrophic forgetting by partitioning capacity into N specialized experts. Each expert implements a differentiable analogue of the hippocampal trisynaptic loop (DG \rightarrow CA3 \rightarrow CA1), combining sparse separation, non-linear refinement, and integration. Expert selection is driven by similarity between current DG codes and per-expert prototypes, enabling dynamic specialization without an explicit learned gating network.

A lightweight CNN backbone (e.g., a LeNet) maps each input to feature maps \mathbf{f} . The grid-cell encoding, sparse DG separation, refinement, and integration are as described in Section .

Gating & Routing

Each expert maintains its own DG module, enabling task-specific sparse representations. At inference, inputs are processed through all N DG modules in parallel to compute routing scores via cosine similarity to prototypes. Once the expert is selected, only that expert’s CA3 and CA1 modules are activated for prediction. The gating computation requires N DG forward passes, included in our reported FLOPs.

Each expert’s prototype \mathbf{u}_i is updated during Phase I via EMA of its DG separation code:

$$\mathbf{u}_i \leftarrow (1 - \mu) \mathbf{u}_i + \mu \mathbf{p}_{\text{sep}}^{(i)}, \quad \mu = 0.01.$$

At inference, similarity scores

$$s_i = \cos(\mathbf{p}_{\text{sep}}^{(i)}, \mathbf{u}_i)$$

are converted to gating weights via:

- *Soft gating*: $\alpha_i \propto \exp(s_i/\tau)$.
- *Hard gating*: $\alpha_i = \mathbf{1}\{i = \arg \max_j s_j\}$.
- *Top-2*: select top two and renormalize.
- *Hybrid*: softmax over top- k similarities.

Outputs from experts are aggregated according to $\{\alpha_i\}$ to form the final prediction.

Training Objectives

Training proceeds in two phases.

Phase I: Specialized Encoding The classification loss is augmented with auxiliary terms. The prototype proximity loss is denoted $\mathcal{L}_{\text{intra}}$ to reflect push-pull dynamics:

$$\begin{aligned} \mathcal{L}_{\text{intra}} = & \sum_{i,j} \mathbf{1}_{[y_i=y_j]} \|\mathbf{p}_i - \mathbf{p}_j\|^2 \\ & - \lambda \sum_{i,j} \mathbf{1}_{[y_i \neq y_j]} [\max\{0, m - \|\mathbf{p}_i - \mathbf{p}_j\|\}]^2. \end{aligned} \quad (2)$$

Additional losses are:

- $\mathcal{L}_{\text{replay}}$: cross-entropy on replayed examples.
- $\mathcal{L}_{\text{distill}}$: feature distillation aligning to frozen snapshots.
- \mathcal{L}_{EWC} : Elastic Weight Consolidation regularizer.
- $\mathcal{L}_{\text{sparsity}}$: encourages target DG activation sparsity.

The Phase I objective combines these terms:

$$\mathcal{L}_{\text{phase1}} = \mathcal{L}_{\text{cls}} + \alpha_{\text{intra}} \mathcal{L}_{\text{intra}} + \alpha_{\text{dist}} \mathcal{L}_{\text{distill}} + \alpha_s \mathcal{L}_{\text{sparsity}}. \quad (3)$$

Phase II: Consolidation & Contrastive Prototypes After freezing non-DG parameters, a prototype-based contrastive loss reinforces current-task alignment while suppressing interference:

$$\begin{aligned} \mathcal{L}_{\text{contrastive}} = & (1 - \cos(\mathbf{p}_{\text{sep}}^{(t)}, \mathbf{u}_t)) \\ & + \sum_{j \neq t} [\max\{0, \cos(\mathbf{p}_{\text{sep}}^{(j)}, \mathbf{u}_j) - m\}]. \end{aligned} \quad (4)$$

$$\mathcal{L}_{\text{phase2}} = \alpha_{\text{contrastive}} \mathcal{L}_{\text{contrastive}}.$$

Method	Buffer Size	Split CIFAR 10			Split TinyImageNet		
		Task IL \uparrow	Class IL \uparrow	FLOPs \downarrow	Task IL \uparrow	Class IL \uparrow	FLOPs \downarrow
<i>Upper Bound</i>							
Joint	–	97.18 \pm 0.12	87.00 \pm 1.03	219.60	82.04 \pm 0.3	60.21 \pm 0.3	850.72
<i>Regularization-Based Methods</i>							
SGD (Fine-tuning)	–	66.08 \pm 7.30	19.53 \pm 0.07	235.66	4.79 \pm 0.43	0.32 \pm 0.02	866.77
LwF	–	63.89 \pm 1.54	19.40 \pm 0.15	235.66	21.43 \pm 0.28	8.54 \pm 0.20	866.77
SI	–	62.16 \pm 4.91	19.29 \pm 0.18	219.60	21.40 \pm 0.19	8.56 \pm 0.08	850.72
<i>Replay-Based Methods</i>							
A-GEM	200	85.35 \pm 2.54	19.66 \pm 0.19	235.66	22.31 \pm 0.94	8.89 \pm 0.10	866.77
	500	85.89 \pm 2.12	19.83 \pm 0.46	235.66	23.17 \pm 0.49	8.70 \pm 0.08	866.77
ER	200	90.07 \pm 0.83	47.77 \pm 1.69	235.66	31.20 \pm 0.10	8.44 \pm 0.25	866.77
	500	92.62 \pm 0.28	61.48 \pm 3.54	235.66	31.74 \pm 0.29	8.62 \pm 0.15	866.77
DER++	200	89.25 \pm 0.52	56.73 \pm 1.46	235.66	33.10 \pm 0.94	8.82 \pm 0.20	866.77
	500	91.21 \pm 0.25	66.49 \pm 1.53	235.66	32.74 \pm 0.92	8.78 \pm 0.07	866.77
<i>Architectural & Hybrid Methods</i>							
iCaRL	200	94.03 \pm 0.46	47.54 \pm 3.28	219.60	24.00 \pm 0.37	8.87 \pm 0.12	850.72
	500	94.36 \pm 0.26	48.34 \pm 1.49	219.60	23.25 \pm 0.91	8.71 \pm 0.05	850.72
FDR	200	91.23 \pm 0.50	29.51 \pm 1.95	236.36	35.90 \pm 0.33	8.38 \pm 0.10	867.48
	500	92.34 \pm 1.26	28.73 \pm 2.82	236.36	41.37 \pm 0.50	8.56 \pm 0.12	867.48
HAL	200	77.97 \pm 2.42	34.11 \pm 5.67	219.60	21.92 \pm 0.57	8.52 \pm 0.18	850.72
	500	84.55 \pm 1.32	41.98 \pm 3.01	219.60	21.49 \pm 0.62	8.45 \pm 0.20	850.72
E2Net	200	93.11 \pm 0.46	67.23 \pm 1.57	219.60	40.94 \pm 0.43	15.35 \pm 0.08	850.72
	500	93.48 \pm 0.37	72.19 \pm 0.82	219.60	45.51 \pm 0.10	22.47 \pm 0.15	850.72
SparCL (Sparsity: 0.25)	200	85.63 \pm 0.62	45.21 \pm 1.10	37.22	39.14 \pm 0.85	8.98 \pm 0.38	149.40
	500	87.38 \pm 0.46	50.45 \pm 1.22	37.22	50.83 \pm 0.69	10.48 \pm 0.29	149.40
<i>Our Method</i>							
HiCL Small (Sparsity: 0.15)	200	90.24 \pm 0.84	57.65 \pm 0.49	38.21	59.16 \pm 0.37	23.28 \pm 0.46	152.9
	500	90.36 \pm 0.54	60.68 \pm 1.05	38.21	59.63 \pm 0.85	23.51 \pm 0.91	152.9
HiCL Large (Sparsity: 0.15)	200	91.74 \pm 0.81	60.48 \pm 0.15	84.18	62.76 \pm 0.26	25.45 \pm 0.14	336.71
	500	92.38 \pm 0.49	65.58 \pm 0.45	84.18	62.60 \pm 0.79	25.71 \pm 0.64	336.71

Table 1: Comparison of accuracy (%), and FLOPs (M) on Split CIFAR-10 and Split TinyImageNet under different memory and compute settings. The baselines have been updated based on new results for Split CIFAR-10. FLOPs are reported in millions.

Full Objective The overall loss combines both phases and retention mechanisms:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{\text{phase1}} + \lambda_2 \mathcal{L}_{\text{phase2}} + \lambda_3 \mathcal{L}_{\text{EWC}} + \lambda_4 \mathcal{L}_{\text{replay}}.$$

Results

Benchmarks and Tasks

We evaluate our model, HiCL, on two standard image-classification benchmarks, each split into a sequence of distinct tasks to simulate a continual learning environment.

- **Split CIFAR-10** (Krizhevsky, Hinton et al. 2009): The 10 classes are divided into 5 sequential tasks, with 2 new classes introduced in each task.
- **Split Tiny-ImageNet** (Deng et al. 2009): A more complex benchmark where 200 classes are split into 10 tasks of 20 classes each.

For all experiments, we follow the **task-incremental learning** protocol, where the model learns tasks sequentially. We

use a ResNet backbone and a fixed-size memory buffer for all methods, where applicable to ensure fair and controlled comparisons.

Evaluation Metrics

Our goal is to develop a model that learns new information effectively while forgetting as little as possible, all with high computational efficiency. To measure this, we report the following metrics:

- **Accuracy (Goal: Higher is Better):** We measure final average accuracy in two scenarios.
 - **Task-IL Accuracy (A_T):** Measures performance when the task identity is provided at test time. This evaluates the model’s ability to preserve task-specific knowledge.
 - **Class-IL Accuracy (A_C):** Measures performance when the task identity is unknown, forcing the model to classify among all previously seen classes. This is

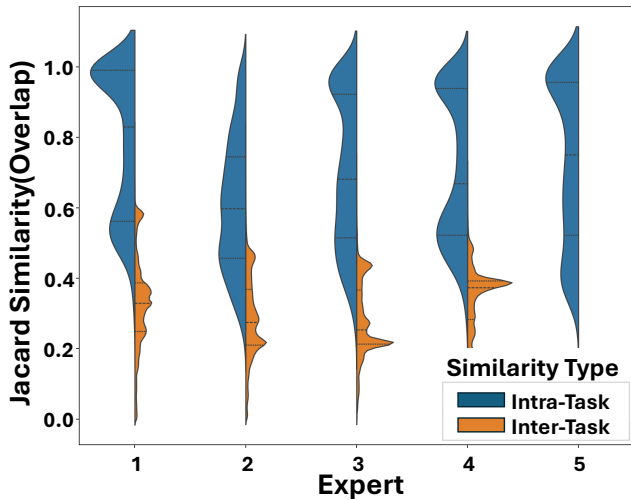


Figure 3: **DG-Gated Routing is Enabled by Effective Feature Separation.** The Class-IL routing matrix shows that inputs are correctly routed to their corresponding expert. For example, inputs from classes 0 and 1 (Task 0) are primarily routed to Expert 0. This routing is enabled by strong pattern separation in the DG layer, where the Jaccard similarity between patterns from different tasks (‘Inter-Task’) is significantly lower than within the same task (‘Intra-Task’).

a more challenging and realistic measure of a model’s ability to avoid catastrophic forgetting.

- **Efficiency (Goal: Lower is Better):** We report the computational cost in Mega-FLOPs (MFLOPs) for a single forward pass. This metric serves as a proxy for the model’s inference speed and complexity.

Training Details

We use a 2-conv + 2-FC LeNet backbone for CIFAR-10 and a 4-conv + 2-FC CNN for Tiny-ImageNet (conv layers: $64 \rightarrow 128 \rightarrow 256 \rightarrow 512$ filters with 3×3 kernels and max pooling). The GridCellLayer has $M = 4$ sinusoidal units, DG size = 1024 with top- k sparsity $k = 5\%$, CA3/CA1 MLP widths (512, 256) and (512, 256, 128) respectively, and replay buffer of 200 samples per task. Models are trained with Adam for 10 epochs per CIFAR-10 task (20 for Tiny-ImageNet), $\lambda_{ewc} = 0.1$, and contrastive margin $m = 0.2$.

Table 1 reports Task-IL accuracy, Class-IL accuracy, and inference FLOPs per forward pass for both Small and Large HiCL variants alongside a suite of baselines. We fix the buffer at 200 or 500 samples per task and evaluate on Split CIFAR-10 (columns 4–6) and Split Tiny-ImageNet (columns 7–9). Baselines are grouped by methodology: regularization-based, replay-based, architectural/hybrid, and the joint offline upper bound.

Discussion

Summary of Key Results

HiCL demonstrates that hippocampal-inspired architecture can achieve state-of-the-art continual-learning performance with high computational efficiency. A key advantage of our

DG-gated design is its support for conditional computation, where the sparse DG codes select a single expert for activation per input, avoiding the need to execute all N expert pipelines.

This approach drastically lowers the total computational footprint. For instance, on Split CIFAR-10, our Small model requires 10.81 MFLOPs for the shared backbone and 5.48 MFLOPs for a single expert, resulting in an efficient inference cost of only 16.3 MFLOPs. The Large model’s cost under the same conditions is 49.7 MFLOPs. On the more complex Split Tiny-ImageNet, these costs are 49.1 MFLOPs for the Small model and 172.7 MFLOPs for the Large model. While activating a single expert may offer less GPU parallelism than a batch-parallel approach, the significant reduction in total operations makes HiCL well-suited for continual learning on resource-constrained hardware. Table 1 reports the full model capacity for fair comparison with baselines that execute all parameters, while conditional computation enables the 16.3 MFLOPs efficiency advantage described above.

The Mechanism of DG-Gated Routing

The success of our routing mechanism stems from effective pattern separation in the DG layer, as evidenced by our ablation study (Table 3). Removing DG gating or the contrastive loss collapses Class-IL accuracy from 57.99 to approximately 19, while omitting top- k sparsity reduces Task-IL accuracy from 92.67 to 55.85, demonstrating that sparse separation is critical to the model’s performance.

Each expert’s Dentate Gyrus (DG) module learns a unique feature vocabulary tailored to its specific task. For an expert trained on Task 1 (e.g., distinguishing cats from dogs), its DG module becomes highly attuned to features like fur texture, whiskers, and ear shape. In contrast, an expert trained on Task 2 (e.g., airplanes vs. ships) will learn to amplify features related to wings, fuselages, and water.

The expert’s prototype or abstraction, which is the average of all its DG codes, becomes a compact summary of this specialized vocabulary. When a new image of a cat is presented to the model, it is processed by all experts’ DG modules in parallel. However, only the DG module from the “cat/dog” expert will produce a strong and characteristic sparse activation, as the input’s features strongly resonate with its learned vocabulary. The other experts, lacking this specialization, will produce weaker and less distinct patterns.

This separation is illustrated by the two dimensional t-SNE embeddings in Figures 4 and 5. Each colored group corresponds to a learned prototype, and the points form compact clusters with clear gaps between groups. The minimal overlap across clusters indicates that prototypes associated with different tasks occupy distinct regions of the representation space, which is consistent with low cross task similarity and reduced interference during continual learning. Therefore, the routing decision, a simple cosine similarity check against the prototypes is not just a superficial pattern match. It is a measure of **representational resonance**. The input is naturally routed to the expert whose internal feature vocabulary is best suited to understand and process it,

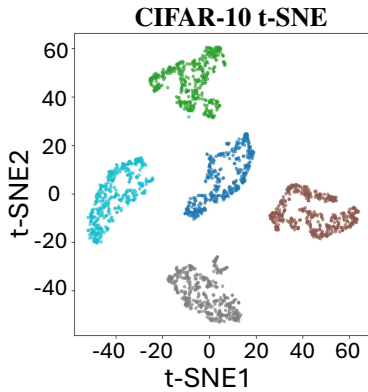


Figure 4: **DG Feature Separation and Prototype Orthogonality.** t-SNE plot for Split CIFAR-10.

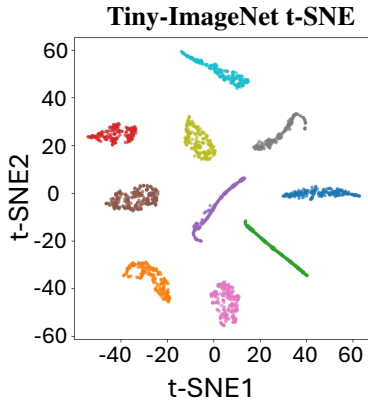


Figure 5: **DG Feature Separation and Prototype Orthogonality.** t-SNE plot for Split Tiny-ImageNet.

a choice made simple and reliable by the clear, quantifiable separation.

Memory Robustness

As shown in Table 2, HiCL degrades gracefully as the replay buffer size shrinks. For instance, reducing the buffer from 500 to 100 examples per task causes Task-IL accuracy to drop by only 2.7 percentage points (from 92.8% to 90.1%). This robustness stems from the DG’s pattern separation: each stored exemplar occupies a distinct sparse subspace, maximizing its utility during replay. Moreover, the model’s feature-distillation and EWC regularization further stabilize representations, mitigating the effect of limited memory. These results suggest that efficient episodic encoding—rather than large buffer capacity—is a key driver of continual-learning performance.

Limitations and Failure Modes

Despite its strengths, HiCL has several limitations. First, it operates under a Task IL assumption, requiring known task boundaries and labels during training. While Class IL inference is label free, Phase 1 still depends on task identities; fully unsupervised gating remains an open direction. Second, we use $\rho = 0.05$ and $N = 5/10$ experts across all

Memory Size	Task-IL	Class-IL	Routing
100	90.1	53.4	56.0
200	92.6	60.6	63.2
300	91.5	63.1	65.6
400	92.2	64.2	67.7
500	92.8	65.4	67.8
600	91.3	64.6	67.6
1000	92.5	67.2	69.1
5000	92.8	63.7	64.1

Table 2: Accuracy vs Memory Size (Task-IL, Class-IL, Routing)

Configuration	Task-IL	Class-IL
Full HiCL	92.67	57.99
w/o Contrastive	91.56	19.19
w/o DG Gating	91.82	18.39
w/o top k Sparsity	55.85	30.71
w/o CA1	92.64	52.39
w/o EMA	90.63	56.29
w/o CA3	92.29	58.73
w/o EWC	89.48	58.92

Table 3: Ablation study on Split CIFAR-10 (Replay size = 200)

tasks. These hyperparameters may not generalize well to domains with different levels of task heterogeneity or complexity. Third, routing can be sensitive in rare cases, with less than five percent of samples showing class confusion due to overlapping DG codes. This could potentially be addressed with adaptive gating thresholds or multi prototype routing. Lastly, our experiments are limited to vision datasets with LeNet style backbones. The model’s performance on larger scale datasets or nonvisual domains such as language or reinforcement learning remains unexplored.

Future Directions

HiCL opens up several exciting paths for future research. A key direction is enabling unsupervised task discovery by leveraging clustering or contrastive pretraining on DG codes to infer task identities, moving toward fully label free Class IL. Another promising avenue is learning adaptive sparsity, where the sparsity ratio ρ is dynamically adjusted per input or expert to balance representation separation and capacity as needed. For prototype maintenance, replacing exponential moving averages with more biologically grounded Hebbian updates or even transformer based attention could improve both interpretability and performance. Scaling HiCL to deeper architectures such as ResNet or Vision Transformers, and evaluating on larger benchmarks like ImageNet Split, will be essential to assess its effectiveness beyond small scale vision tasks. Finally, translating DG gating and expert routing into neuromorphic implementations using spiking neurons could enable real time, low power continual learning on edge hardware.

References

- Aghasanli, A.; Li, Y.; and Angelov, P. 2025. Prototype-Based Continual Learning with Label-free Replay Buffer and Cluster Preservation Loss. *arXiv preprint arXiv:2504.07240*.
- Alonso, N.; and Krichmar, J. L. 2024. A sparse quantized hopfield network for online-continual memory. *Nature Communications*, 15(1): 3722.
- Banino, A.; Barry, C.; Uribe, B.; et al. 2018. Vector-based navigation using grid-like representations in artificial agents. *Nature*, 557(7705): 429–433.
- Bartsch, T.; Döhring, J.; Rohr, A.; Jansen, O.; and Deuschl, G. 2011. CA1 neurons in the human hippocampus are critical for autobiographical memory, mental time travel, and auto-noetic consciousness. *Proceedings of the National Academy of Sciences*, 108(42): 17562–17567.
- Caccia, L.; Aljundi, R.; Tuytelaars, T.; Pineau, J.; and Belilovsky, E. 2021. Reducing Representation Drift in Online Continual Learning. *CoRR*, abs/2104.05025.
- Cai, W.; Jiang, J.; Wang, F.; Tang, J.; Kim, S.; and Huang, J. 2025. A survey on mixture of experts in large language models. *IEEE Transactions on Knowledge and Data Engineering*.
- Cairney, S. A.; El Marj, N.; Staresina, B. P.; et al. 2018. Memory consolidation is linked to spindle-mediated information processing during sleep. *Current Biology*, 28(6): 948–954.
- De Lange, M.; Aljundi, R.; Masana, M.; Parisot, S.; Jia, X.; Leonardis, A.; Slabaugh, G.; and Tuytelaars, T. 2021. A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern analysis and machine intelligence*, 44(7): 3366–3385.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Fedus, W.; Zoph, B.; and Shazeer, N. 2022. Switch Transformers: Scaling to Trillion Parameter Models with Simple and Efficient Sparsity. *Journal of Machine Learning Research*, 23(120): 1–39.
- Gilbert, P. E.; and Brushfield, A. M. 2009. The role of the CA3 hippocampal subregion in spatial memory: a process oriented behavioral assessment. *Progress in Neuro-Psychopharmacology and Biological Psychiatry*, 33(5): 774–781.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Hafting, T.; Fyhn, M.; Molden, S.; Moser, M.-B.; and Moser, E. I. 2005. Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052): 801–806.
- Hainmueller, T.; and Bartos, M. 2020. Dentate gyrus circuits for encoding, retrieval and discrimination of episodic memories. *Nature Reviews Neuroscience*, 21(3): 153–168.
- Hinton, G. 2022. The Forward-Forward Algorithm: Some Preliminary Investigations. *arXiv:2212.13345*.
- Hopfield, J. J. 1982. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8): 2554–2558.
- Kemker, R.; and Kanan, C. 2017. Fearnnet: Brain-inspired model for incremental learning. *arXiv preprint arXiv:1711.10563*.
- Kemker, R.; McClure, M.; Abitino, A.; Hayes, T. L.; and Kanan, C. 2018. Measuring catastrophic forgetting in neural networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13): 3521–3526.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Krotov, D.; and Hopfield, J. J. 2016. Dense associative memory for pattern recognition. *Advances in neural information processing systems*, 29.
- Le, M.; Nguyen, A.; Nguyen, H.; Nguyen, T.; Pham, T.; Van Ngo, L.; and Ho, N. 2024. Mixture of experts meets prompt-based continual learning. *Advances in Neural Information Processing Systems*, 37: 119025–119062.
- Leutgeb, S.; Leutgeb, J. K.; Treves, A.; Moser, M.-B.; and Moser, E. I. 2007. Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science*, 315(5814): 961–966.
- Li, H.; Lin, S.; Duan, L.; Liang, Y.; and Shroff, N. 2025. Theory on Mixture-of-Experts in Continual Learning. In *International Conference on Learning Representations (ICLR)*.
- Li, Z.; and Hoiem, D. 2017. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12): 2935–2947.
- Madotto, A.; Lin, Z.; Zhou, Z.; Moon, S.; Crook, P. A.; Liu, B.; Yu, Z.; Cho, E.; Fung, P.; and Wang, Z. 2021. Continual learning in task-oriented dialogue systems.
- Mattar, M. G.; and Daw, N. D. 2018. Prioritized replay of salient experiences enables efficient learning in humans and machines. *Nature Human Behaviour*, 2(10): 737–747.
- McClelland, J. L.; McNaughton, B. L.; and O’Reilly, R. C. 1995. 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.
- Moser, E. I.; Kropff, E.; and Moser, M.-B. 2008. Place cells, grid cells, and the brain’s spatial representation system. *Annual review of neuroscience*, 31: 69–89.
- Nie, X.; et al. 2022. EvoMoE: An Evolutional Mixture-of-Experts Training Framework via Dense-to-Sparse Gate. *arXiv preprint arXiv:2112.14397*.
- Otmakhova, N. A.; Otmakhov, N.; and Griffith, L. C. 2018. Memories of John Lisman.

- Petrantonakis, P. C.; and Poirazi, P. 2014. A compressed sensing perspective of hippocampal function. *Frontiers in systems neuroscience*, 8: 141.
- Pham, Q.; Liu, C.; and Hoi, S. 2021. Dualnet: Continual learning, fast and slow. *Advances in Neural Information Processing Systems*, 34: 16131–16144.
- Rebuffi, S.-A.; Kolesnikov, A.; Sperl, G.; and Lampert, C. H. 2017. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2001–2010.
- Riemer, M.; Cases, I.; Ajemian, R.; Liu, M.; Rish, I.; Tu, Y.; and Tesauro, G. 2018. Learning to learn without forgetting by maximizing transfer and minimizing interference. *International Conference on Learning Representations*.
- Rolls, E. T. 2013a. The mechanisms for pattern completion and pattern separation in the hippocampus. *Frontiers in Systems Neuroscience*, 7: 74.
- Rolls, E. T. 2013b. Pattern separation, pattern completion, and new inputs to the hippocampus. *Neural Networks*, 48: 5–14.
- Rolnick, D.; Ahuja, A.; Schwarz, J.; Lillicrap, T.; and Wayne, G. 2019. Experience replay for continual learning. *Advances in neural information processing systems*, 32.
- Rothschild, G.; Eban, E.; and Frank, L. M. 2017. A cortical–hippocampal–cortical loop of information processing during memory consolidation. *Nature neuroscience*, 20(2): 251–259.
- Rusu, A. A.; Rabinowitz, N. C.; Desjardins, G.; Soyer, H.; Kirkpatrick, J.; Kavukcuoglu, K.; Pascanu, R.; and Hadsell, R. 2016. Progressive neural networks. *arXiv preprint arXiv:1606.04671*.
- Schaul, T.; Quan, J.; Antonoglou, I.; and Silver, D. 2016. Prioritized experience replay. *International Conference on Learning Representations*.
- Schwarz, J.; Czarnecki, W.; Luketina, J.; Grabska-Barwinska, A.; Teh, Y. W.; Pascanu, R.; and Hadsell, R. 2018. Progress & compress: A scalable framework for continual learning. *Proceedings of the 35th International Conference on Machine Learning*, 4528–4537.
- Serrà, J.; Surís, D.; Miron, M.; and Karatzoglou, A. 2018. Overcoming Catastrophic Forgetting with Hard Attention to the Task. In *ICML*, 4548–4557.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; et al. 2017. Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer. *arXiv preprint arXiv:1701.06538*.
- Shin, H.; Lee, J. K.; Kim, J.; and Kim, J. 2017. Continual learning with deep generative replay. *Advances in neural information processing systems*, 30.
- Skierś, P.; and Deja, K. 2025. Joint diffusion models in continual learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4380–4390.
- Su, X.; Guo, S.; Tan, T.; and Chen, F. 2019. Generative memory for lifelong learning. *IEEE transactions on neural networks and learning systems*, 31(6): 1884–1898.
- Treves, A.; and Rolls, E. T. 1994. Computational Analysis of the Role of the Hippocampus in Memory. *Hippocampus*, 4(3): 374–391.
- Wang, L.; Lei, B.; Li, Q.; Su, H.; Zhu, J.; and Zhong, Y. 2021. Triple-memory networks: A brain-inspired method for continual learning. *IEEE Transactions on Neural Networks and Learning Systems*, 33(5): 1925–1934.
- Wang, M.; Adel, H.; Lange, L.; Strötgen, J.; and Schütze, H. 2024. Rehearsal-free modular and compositional continual learning for language models. *arXiv preprint arXiv:2404.00790*.
- Yuksel, S. E.; Wilson, J. N.; and Gader, P. D. 2012. Twenty years of mixture of experts. *IEEE transactions on neural networks and learning systems*, 23(8): 1177–1193.
- Zenke, F.; Poole, B.; and Ganguli, S. 2017. Continual Learning Through Synaptic Intelligence. In *ICML*, 3987–3995.
- Zhang, Y.; Wang, X.; and Yang, D. 2022. Continual sequence generation with adaptive compositional modules. *arXiv preprint arXiv:2203.10652*.