

Interactive Continual Learning: Fast and Slow Thinking

Biqing Qi^{1,2,4}, Xinquan Chen³, Junqi Gao³, Dong Li³, Jianxing Liu¹, Ligang Wu^{1,*}, Bowen Zhou^{1,2,4,*}

¹ Department of Control Science and Engineering, Harbin Institute of Technology,

² Department of Electronic Engineering, Tsinghua University,

³ School of Mathematics, Harbin Institute of Technology,

⁴ Frontis.AI, Beijing

{qibiqing7, xinquanchen0117, gjunqi97, arvinlee826}@gmail.com,

{jx.liu, ligangwu}@hit.edu.cn, {zhoubowen}@tsinghua.edu.cn

Abstract

Advanced life forms, sustained by the synergistic interaction of neural cognitive mechanisms, continually acquire and transfer knowledge throughout their lifespan. In contrast, contemporary machine learning paradigms exhibit limitations in emulating the facets of continual learning (CL). Nonetheless, the emergence of large language models (LLMs) presents promising avenues for realizing CL via interactions with these models. Drawing on Complementary Learning System theory, this paper presents a novel Interactive Continual Learning (ICL) framework, enabled by collaborative interactions among models of various sizes. Specifically, we assign the ViT model as System1 and multimodal LLM as System2. To enable the memory module to deduce tasks from class information and enhance Set2Set retrieval, we propose the Class-Knowledge-Task Multi-Head Attention (CKT-MHA). Additionally, to improve memory retrieval in System1 through enhanced geometric representation, we introduce the CL-vMF mechanism, based on the von Mises-Fisher (vMF) distribution. Meanwhile, we introduce the von Mises-Fisher Outlier Detection and Interaction (vMF-ODI) strategy to identify hard examples, thus enhancing collaboration between System1 and System2 for complex reasoning realization. Comprehensive evaluation of our proposed ICL demonstrates significant resistance to forgetting and superior performance relative to existing methods. Code is available at github.com/ICL.

1. Introduction

Advanced life forms exhibit continual learning (CL) and memory formation, facilitated by neural cognitive interactions that enable collaborative knowledge transfer [9, 26, 48]. These underlying mechanisms enhance memory

consolidation and utilization, as well as reasoning abilities in advanced life forms [16, 39]. However, current machine learning paradigms, particularly neural network-based models, face challenges in achieving CL. Specifically, neural networks learning from evolving data face a risk known as catastrophic forgetting [5], wherein the integration of new knowledge frequently disrupts existing knowledge, resulting in notable performance degradation [5, 17, 50].

To tackle this challenge, current CL methods strive to preserve and augment knowledge acquired throughout the learning process [14, 22, 28, 36]. In CL, rehearsal-based methods [32, 36, 40, 44–46], are the most direct strategy.

However, these methods often ignore the geometric structure [30] of memory representations and face challenges in open-class settings. Another perspective includes architecture-based methods [21, 22, 31], allocating distinct parameters for knowledge encoding from various tasks. Early studies centered on convolution-based architectures [28]. Recent advancements pivoted towards transformer-based methods like L2P [46], and Dualprompt [45].

From the perspective of Complementary Learning System (CLS) Theory in neurocognitive science [16], the current designs of CL frameworks may not be optimal. In a brain-like system, multiple memory modules dynamically maintain a balance between stability and plasticity, with each module possessing predictive capabilities [24, 33, 35]. However, most advanced CL frameworks lean towards more intuitive systems. Previous CLS-driven methods [25, 28] involve the separation and expansion of parameters to facilitate the learning of both fast and slow knowledge. For instance, DualNet [28] optimizes this process by emphasizing task-specific pattern separation. Similarly, BiMeCo [25] divides model parameters into two distinct components: a short-term memory module and a long-term memory module. However, these methods [25, 28] are limited to a single backbone model. In line with CLS prin-

*Corresponding authors

ciples, this underscores the importance of developing an interactive CL framework between models to consistently achieve higher performance levels.

Meanwhile, recent advancements in large language models (LLMs), as exemplified by ChatGPT [38] and GPT4 [5], have demonstrated remarkable reasoning capabilities. These models can employ chains of thought [47] to engage in complex reasoning, much like System2. Consequently, it raises an interesting question: Can we integrate intuitive models, such as ViT [8] as System1 alongside LLM-based as System2 to establish an interactive framework for CL?

To response this question, we reevaluated CL by exploring the interaction between ViT (System1) and Multimodal Large Language Model (System2). In alignment with the current CL setting, our focus is on adapting ViT parameters while keeping System2 parameters stable. System2 is responsible for handling hard examples and facilitates collaboration with System1. To enable the continual updating of ViT, we introduced the Class-Knowledge-Task Multi-Head Attention (CKT-MHA) module. CKT-MHA utilizes category features and the knowledge of ViT to aid System1 in acquiring task-related knowledge, facilitating knowledge retrieval through Class-Task collections. Furthermore, we introduce the CL-vMF mechanism, which employs von Mises-Fisher distribution modeling to improve memory geometry and enhance retrieval distinguishability through an Expectation-Maximization (EM) update strategy. This design enables System1 to retain old memory parameters, preventing unnecessary updates and addressing catastrophic forgetting issues. To realize the coordination between System1 and System2 during their reasoning transitions, we introduce the von Mises-Fisher Outlier Detection and Interaction (vMF-ODI) mechanism for assessing sample difficulty. This mechanism is designed to assist the System1 in adaptive identifying hard examples within each batch. Once identified, these hard examples undergo initial inference by System1, and the resulting prediction outcomes serve as background knowledge for System2 to facilitate more intricate reasoning.

We conduct experiments on various benchmarks, including the demanding Imagenet-R, to validate proposed interactive continual learning (ICL). The results illustrate that ICL significantly mitigates catastrophic forgetting, surpassing state-of-the-art methods. Moreover, it maintains consistently high accuracy across different task. In summary, our contributions are as follows:

- We propose an ICL framework from a novel perspective that emphasizes the interaction between fast intuitive model (ViT) and slow deliberate model (multimodal LLM), aligning with CLS principles.
- We propose the CKT-MHA module to acquire task-related information by leveraging category features and small model knowledge.

- We propose the CL-VMF mechanism, an optimization strategy guided by VMF distribution modeling with EM strategy updates to enhance the retrieval of geometric memory representations.
- We propose vMF-ODI, a batch-wise retrieval interaction strategy that enables the adaptive identification of hard examples within each batch, fostering collaborative reasoning between the two systems.

2. Related Works

We discuss three primary categories of CL methods.

Regularization-based methods incorporate regularization terms into the loss function to mitigate catastrophic forgetting for previously learned tasks. These methods [1, 14, 17, 20, 27, 50], primarily revolve around the development of metrics for assessing task importance, with additional research efforts dedicated to characterizing the significance of individual features [42]. Nevertheless, these methods tend to exhibit reduced performance when applied to more complex datasets.

Rehearsal based methods utilize previous task data to mitigate catastrophic forgetting within limited memory buffers. Reservoir sampling techniques [7, 34], randomly retain a fixed number of old training samples from each training batch. Further, [12] employs coefficient-based cosine similarity to address sample number imbalances among categories. To better recover past knowledge, GEM [36] constructs individual constraints based on old training samples for each task to ensure no increase in their loss. LOGD [40] decomposes gradients for each task into shared and task-specific components, capitalizing on inter-task information. CVT [44] explores online CL using external attention strategy.

Architecture-based methods aim at assigning independent parameters for new task data. These methods involve strategies such as parameter allocation, model division, and modular network models [21, 22, 31]. Previous studies concentrated on specific convolution-based architectures, with DualNet [28] optimizes memory through separation representations for specific tasks. Recent work focus to transformer-based models. L2P [46] enhances integration of knowledge by treating prompts as optimization parameters. Furthermore, Dualprompt [45] enhances knowledge memory by constructing dual orthogonal prompt spaces.

3. Methodology

3.1. Problem Setup

A standard paradigm for CL can be defined by a set of task descriptors $t \in \mathcal{T}$, and the corresponding distribution $p_t(\mathbf{x}, y)$ for each task. The task-specific dataset $\mathcal{D}_t := \{(\mathbf{x}_i^t, y_i^t) \mid (\mathbf{x}_i^t, y_i^t) \sim p_t, i \in [N_t]\}$, where N_t represents the number of samples in the training set of the t -th

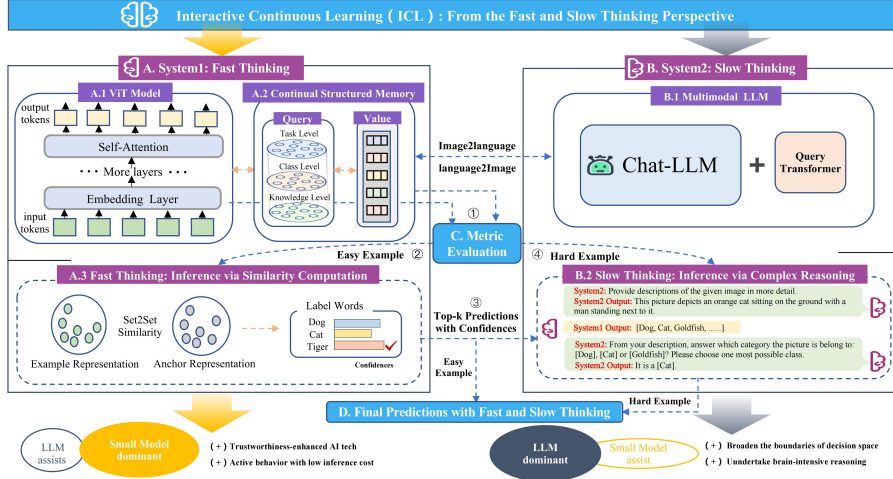


Figure 1. Comprehensive Training and Testing Illustration. In the training phase: we propose the CKT-MHA unified storage module for System1. And then use memory selection and updates through our CL-VMF mechanism with the EM strategy to optimize CL for the small model ViT. In the inference phase: 1) The process begins by assessing sample complexity using proposed vMF-ODI in System 1. 2) The System1 then swiftly generates inferential predictions. 3) If test samples surpass a complexity threshold, we activate collaborative inference. Specifically, the predictive results from System1 is used as background knowledge to narrow the scope of inference. 4) Subsequently, complex reasoning through the multimodal LLM is applied to achieve the final prediction.

task. The dataset is drawn i.i.d. from the sample space $\mathcal{X}^t \times \mathcal{Y}^t \in \mathcal{X} \times \mathcal{Y}$. During the training phase, the training samples are sequentially fed as inputs, following the task descriptors from 0 to $|\mathcal{T}|$. In formulating the ICL framework, we provide definitions for System1 and System2, respectively. System1 is instantiated by the model $f_{\theta}(\cdot) : \mathcal{X} \mapsto \mathbb{R}^d$ parameterized by θ which updates its parameters to θ_t in the t -th task. Our objective is to determine $\theta_{|\mathcal{T}|} = \arg \min_{\theta_{|\mathcal{T}|}} \mathbb{E}_{t \in \mathcal{T}} \mathbb{E}_{(\mathbf{x}^t, \mathbf{y}^t)} [\ell(f_{\theta_{|\mathcal{T}|}}(\mathbf{x}^t), \mathbf{y}_i^t)]$, ensuring the memory capacity of the System1. Here $\ell(\cdot, \mathbf{y}^t)$ represents loss function, with $\mathbf{y}^t \in \mathcal{Y}^t$. The system1 utilizes a memory buffer \mathcal{M} and updates θ_{t-1} to θ_t by using $\mathcal{D}_t \cap \mathcal{M}$. Furthermore, System2 is instantiated with a Multimodal LLM represented as g_{ψ} to model complex reasoning abilities. To enable collaborative inference, System2 must handle hard samples $\tilde{\mathcal{X}}$ that System1 struggles with, maximizing the probability $\mathbb{P}_{\tilde{\mathbf{x}} \in \tilde{\mathcal{X}}} (\hat{y} \neq \arg \max_i f_{\theta}(\tilde{\mathbf{x}})_i)$. This requires the ability of System1 to filter out these samples, thereby improving the inference of System2 by leveraging predicted results from System1 $I_{f_{\theta}}$ for second-stage inference:

$$\hat{y} = \mathcal{S}^{-1}(g_{\psi}(\tilde{\mathbf{x}}; \mathcal{S}(I_{f_{\theta}}(\tilde{\mathbf{x}})))), \quad (1)$$

here \mathcal{S} represents label to prompt operation, while \mathcal{S}^{-1} represents its inverse operation. System2 is expected to produce an output $\mathcal{S}(\hat{y})$ that minimizes $-\log p_{\psi}(\mathcal{S}(\hat{y}) | \tilde{\mathbf{x}}; \mathcal{S}(I_{f_{\theta}}(\tilde{\mathbf{x}})))$. Next, we will present detailed designs for each component of the ICL framework and discuss optimization strategies.

3.2. Query and Value Memory for System1

In general, deep neural networks implicitly encode data memories within their parameters, and unnecessary param-

eter changes can result in memory degradation. The prevailing methods to deploying models in downstream tasks involves utilizing pre-trained feature extractors and introducing new parameters for adaptation, which has been demonstrated to be effective in CL setting [45, 46]. Nevertheless, these methods face challenge when the precise number of classes in the downstream task is uncertain. Moreover, updating all parameters of the classification head for each new task worsens the issue of forgetting. To address these challenges, we propose separating the model's parameters into two distinct groups: value memory parameters, denoted as $\mathcal{Z} = \mathcal{Z}_1, \mathcal{Z}_2, \dots$ with class-specific representations $\mathbf{z}^{y^t} \in \mathcal{Z}_t$, and query memory parameters represented by θ . This decoupling strategy enhances operational flexibility. In simpler terms, we envision \mathcal{Z} as a collection of class-specific value memory variables, ensuring that

- Value memory parameters \mathcal{Z} can be augmented as the number of tasks increases, i.e., there is a memory increment. When training on task t , \mathcal{Z} will be updated from $\{\mathcal{Z}_1, \dots, \mathcal{Z}_{t-1}\}$ to $\{\mathcal{Z}_1, \dots, \mathcal{Z}_t\}$.
- The relevant value memory for each class is only updated when necessary, i.e., the value memory \mathbf{z}^{y^t} will not be updated if the input does not contain class y^t , thus $\mathbf{z}^{y^{t-1}}$ after trained on task t equals to that trained before training on task t .

Then we can ensure that the System1 can be equipped with persistent value memory of old data and update query parameters using a portion of the memory buffer as rehearsal samples, thus can effectively handle old data. At the same time, this enables System1 to be unconstrained by a predefined number of classes and allows for more flexible memory allocation for new tasks.

Interactive Query and Value Memory with CKT-MHA

Firstly, we design such a memory module for System1. Specifically, we propose a Set2Set memory retrieval mechanism to further enhance memory retrieval stability. This mechanism involves first obtaining class (cls) information via a projector f_{θ_c} parameterized by θ_c : $\xi_c = f_{\theta_c, \varphi}(\mathbf{x}) \in \mathbb{R}^{L \times d_c}$, L is the obtained token length. Here we use the pre-trained ViT as image feature extractor f_φ , followed by an query interactor f_θ , parameterized by θ , introduced for memory matching. This results in $f_{\theta, \varphi} = f_\theta \circ f_\varphi$. Subsequently, we utilize ξ_c and pretrained knowledge $\xi_k = f_\varphi(\mathbf{x}) \in \mathbb{R}^{L \times d_c}$ to construct the Class-Knowledge-Task Multi-Head Attention (CKT-MHA) to capture the task information corresponding to the class:

$$\mathbf{h}_{\tau, i} = \text{Attention}_{\theta_{SA}} \left(\xi_c^{[:, R_i]}, \xi_k^{[:, R_i]}, f_{\theta_\tau}(\xi_k)^{[:, R_i]} \right), \quad (2)$$

$$\xi_\tau = f_{\theta_o}(\text{Concat}[\mathbf{h}_{\tau, 1}, \dots, \mathbf{h}_{\tau, N_h}]) \in \mathbb{R}^{L \times \tau}, \quad (3)$$

where N_h is the head number of MHA, $R_i = [(i-1) * \frac{d_c}{N_h}, i * \frac{d_c}{N_h}]$ is the corresponding attention head interval. Then we combine the class and task features set to obtain the classification feature: $\xi = \text{Concat}[\frac{1}{L} \sum_{l=1}^L \xi_\tau^l, \frac{1}{L} \sum_{l=1}^L \xi_c^l]$, where ξ_τ^l denotes the l -th token of ξ_τ . Consequently, the interactor parameters are denoted as $\theta = \{\theta_c, \theta_\tau, \theta_{SA}, \theta_o\}$. Finally, we set a group of task-specified value memory vector $\mathbf{z}_\tau \in \mathcal{Z}^\tau$ and class-specified $\mathbf{z}_c \in \mathcal{Z}^c$ corresponding to the class to form the final value memory variable \mathbf{z} for retrieval: $\mathbf{z} = \text{Concat}[\mathbf{z}_\tau, \mathbf{z}_c]$. During the inference stage, the proposed task-class Set2Set retrieval is performed from value memory as: $y_{\text{pred}} = \arg \max_{z \in \mathcal{Z}} p_{\theta, \varphi}(z|\mathbf{x})$.

With decoupled parameters, we also need to decouple the updates of query memory parameters and value memory parameters to ensure the flexibility of adjusting these two parts of parameters, avoiding unnecessary updates to value memory parameters, which prevents biases in memory variables for future tasks and mitigates catastrophic forgetting. Furthermore, prioritizing value memory parameter optimization guides the optimization of query parameters. Is it possible to achieve such decoupled optimization while ensuring a consistent optimization objective? The inherent nature of EM algorithm provides a framework for meeting these optimization requirements. Specifically, considering the Maximum Likelihood Estimation (MLE) under a classification setting, with the probability of task t denoted as $\mathbb{P}(t)$, the objective can be expressed as follows:

$$\text{Minimize}_{\theta|\tau} \sum_t \mathbb{P}(t) \mathbb{E}_{\mathbf{x}^t} [-\log p_{\theta|\tau_1}(y^t|\mathbf{x}^t)]. \quad (4)$$

For the sake of simplicity in the framework, here we only consider the parameters that are being updated. During training on task t , the System1 aims to determine $\theta_t = \arg \min_{\theta_t} \mathbb{E}_{\mathbf{x}^t} [-\log p_{\theta_t}(y^t|\mathbf{x}^t)]$. To achieve this, we model value memory $\mathbf{z} \in \mathcal{Z}$ as hidden variables,

alongside the hidden distribution $q(\mathbf{z}|\mathbf{x})$ which satisfies $\sum_{z \in \mathcal{Z}} q(\mathbf{z}|\mathbf{x}) = 1$. This allows us to obtain

$$\log p_{\theta_t}(y^t|\mathbf{x}^t) = \log \left[\frac{p_{\theta_t}(y^t, \mathbf{z}|\mathbf{x}^t)}{q(\mathbf{z}|\mathbf{x}^t)} \right] - \log \left[\frac{p_{\theta_t}(\mathbf{z}|y^t, \mathbf{x}^t)}{q(\mathbf{z}|\mathbf{x}^t)} \right], \quad (5)$$

taking the expectation with respect to $q(\mathbf{z}|\mathbf{x})$ on both sides,

$$\log p_{\theta_t}(y^t|\mathbf{x}^t) = \sum_{z \in \mathcal{Z}} q(\mathbf{z}|\mathbf{x}^t) \log \left[\frac{p_{\theta_t}(y^t, \mathbf{z}|\mathbf{x}^t)}{q(\mathbf{z}|\mathbf{x}^t)} \right], \quad (6)$$

$$+ \sum_{z \in \mathcal{Z}} q(\mathbf{z}|\mathbf{x}^t) \log \left[\frac{q(\mathbf{z}|\mathbf{x}^t)}{p_{\theta_t}(\mathbf{z}|y^t, \mathbf{x}^t)} \right], \quad (7)$$

the first term of R.H.S is referred to as the Evidence Lower Bound $\mathcal{L}_{\text{EL}}(p_{\theta_t}(y^t, \mathbf{z}|\mathbf{x}^t), q(\mathbf{z}|\mathbf{x}^t))$, the second term is $\text{KL}(q(\mathbf{z}|\mathbf{x}^t) \| p_{\theta_t}(\mathbf{z}|y^t, \mathbf{x}^t))$. Since the KL divergence is non-negative, we can achieve maximum likelihood by iteratively updating \mathcal{Z} and θ_t using the Generalized Expectation-Maximization algorithm (GEM) as follows:

$$\mathcal{Z}^{(i+1)} = \arg \min_{\mathcal{Z}} \mathbb{E}_{\mathbf{x}^t} \left[\text{KL} \left(q(\mathbf{z}^{(i)}|\mathbf{x}^t) \| p_{\theta_t^{(i)}}(\mathbf{z}^{(i)}|y^t, \mathbf{x}^t) \right) \right], \quad (8)$$

$$\theta_t^{(i+1)} = \arg \max_{\theta} \mathbb{E}_{\mathbf{x}^t} \left[\mathcal{L}_{\text{EL}} \left(p_{\theta_t^{(i)}}(y^t, \mathbf{z}^{(i+1)}|\mathbf{x}^t), q(\mathbf{z}^{(i+1)}|\mathbf{x}^t) \right) \right]. \quad (9)$$

In the context of supervised learning, we can define \mathbf{z} as class-specific $\mathbf{z}^{y^t} \in \mathcal{Z}_t$. The prior form of q is $q(\mathbf{z}|\mathbf{x}^t) = \sum_{z \in \mathcal{Z}} \mathbb{1}(z = \mathbf{z}^{y^t})$. Note that in this case,

$$p_{\theta_t}(\mathbf{z}|y^t, \mathbf{x}^t) = p_{\theta_t}(\mathbf{z}^{y^t}|\mathbf{x}^t) = p_{\theta_t}(\mathbf{z}, y^t|\mathbf{x}^t). \quad (10)$$

Hence, we can express the two distinct objectives in eq.(8) and (9) as a unified objective:

$$\theta_t^*, \mathcal{Z}_t^* = \arg \min_{\theta, \mathcal{Z}} \mathbb{E}_{\mathbf{x}^t} \left[-\log p_{\theta_t}(\mathbf{z}^{y^t}|\mathbf{x}^t) \right]. \quad (11)$$

When training on a sequence of tasks, if we directly optimize as in eq.(4) for the current task's optimal $\theta_t^* = \arg \min_{\theta} \mathbb{E}_{\mathbf{x}^t} [-\log p_{\theta_t}(y^t|\mathbf{x}^t)]$, as is typical, it will inevitably lead to $\theta_t^* \neq \theta_{t-1}^*$, resulting in catastrophic forgetting. However, our decoupled optimization strategy ensures that the value memory parameters remain optimal for their corresponding tasks, and with the integration of rehearsal, it guarantees that the optimization of query memory parameters can be consistently guided by such value memory parameters of old tasks, thereby mitigating catastrophic forgetting while continually adapting to new data. Next, we describe how to model $p_{\theta_t}(\mathbf{z}^{y^t}|\mathbf{x}^t)$ in eq.(11).

3.3. Optimizing Memory via CL-vMF

Modelling Posterior with vMF Distribution. To assure value memory vectors more discriminative that facilitate more explicit memory retrieval, necessitating the construction of more separable geometric relationships for them. Introducing improved geometric relationships for value memory vectors further ensures that the query features of each

class are more separable and closer to the class center when combined with the EM updating strategy. This also aids in more effectively filtering out outliers, laying the groundwork for screening hard samples for System1, as discussed in Section 3.4. Therefore, we opt for the von Mises-Fisher (vMF) distribution, which naturally excels in modeling geometric relationships in high-dimensional spaces [2, 18, 29] and has demonstrated effectiveness in downstream tasks [23, 51]. Its probability density function is as follows:

$$p_{\text{vMF}}(\mathbf{x}; \boldsymbol{\mu}, \kappa) = C_d(\kappa) \exp(\kappa \langle \boldsymbol{\mu}, \mathbf{x} \rangle), \quad (12)$$

where the correlated sample space is defined as $\{\mathbf{x} | \mathbf{x} \in \mathbb{R}^d, \|\mathbf{x}\| = 1\}$. $\boldsymbol{\mu} \in \mathbb{R}^d$ represent the mean direction, while κ is the concentration parameter. The constant $C_d(\kappa)$ is only related to κ and d . By setting normalized value memory parameter $\mathbf{z}^{y^t} / \|\mathbf{z}^{y^t}\|$ to the mean direction, we model the distribution of the normalized feature $f_{\boldsymbol{\theta}, \varphi}(\mathbf{x}^t) / \|f_{\boldsymbol{\theta}, \varphi}(\mathbf{x}^t)\|$ for class y^t as a vMF distribution with probability density $C_d(\kappa) \exp(\kappa \langle \frac{\mathbf{z}^{y^t}}{\|\mathbf{z}^{y^t}\|}, \frac{f_{\boldsymbol{\theta}, \varphi}(\mathbf{x}^t)}{\|f_{\boldsymbol{\theta}, \varphi}(\mathbf{x}^t)\|} \rangle)$. By utilizing the constructed probability density for Bayesian discrimination, we can ascertain the form of the posterior probability $p_{\boldsymbol{\theta}, \varphi}(\mathbf{z}^{y^t} | \mathbf{x}^t)$ as:

$$p_{\boldsymbol{\theta}, \varphi}(\mathbf{z}^{y^t} | \mathbf{x}^t) = \frac{\exp(\kappa \langle \frac{\mathbf{z}^{y^t}}{\|\mathbf{z}^{y^t}\|}, \frac{f_{\boldsymbol{\theta}, \varphi}(\mathbf{x}^t)}{\|f_{\boldsymbol{\theta}, \varphi}(\mathbf{x}^t)\|} \rangle)}{\sum_{\mathbf{z}^t \in \mathcal{Z}} \exp(\kappa \langle \frac{\mathbf{z}^t}{\|\mathbf{z}^t\|}, \frac{f_{\boldsymbol{\theta}, \varphi}(\mathbf{x}^t)}{\|f_{\boldsymbol{\theta}, \varphi}(\mathbf{x}^t)\|} \rangle)}. \quad (13)$$

When a sample pair of a new class $(x^{t+1}, y^{t+1}) \in \mathcal{D}^{t+1}$ is inputted, a new $\mathbf{z}^{y^{t+1}} \notin \mathcal{Z}$ can be assigned to it for memory expansion, thereby avoiding the limitation imposed by the predefined number of classes, i.e., class-free. Combined with the proposed memory module, we introduce CL-vMF, a class-free memory retrieval mechanism that accommodates new classes when the total number is unknown, it retrieves from memories based on the vMF posterior and ensures that class-specific value memory parameters are only updated when needed.

Implementation of CL-vMF. We set the value memory vectors \mathcal{Z} as learnable memory embeddings, enabling the model to use them for retrieval using the vMF posterior to ascertain a class of sample. The process resembles constructing a new "hippocampus" within the pre-trained model by utilizing \mathcal{Z} and learning query memory interactions via $\boldsymbol{\theta}$ to adapt and employ the value memory. In the training phase, upon encountering each new class y , we allocate a \mathbf{z}^y , incorporating it into \mathcal{Z} . To ensure updating the value memory parameter \mathbf{z}^y exclusively when the input contains a sample of class y , we introduce the subsequent batchwise vMF loss to maximize likelihood:

$$\mathcal{L}_{\text{vMF}}(\mathcal{B}_i^t) = \frac{1}{|\mathcal{B}_i^t|} \sum_{(\mathbf{x}^t, y^t) \in \mathcal{B}_i^t} -\log \left[p_{\boldsymbol{\theta}, \varphi}^{\mathcal{B}_i^t}(\mathbf{z}^{y^t} | \mathbf{x}^t) \right], \quad (14)$$

where $\mathcal{B}_i^t \in \mathcal{D}^t$ is the i -th input batch of task t and

$$p_{\boldsymbol{\theta}, \varphi}^{\mathcal{B}_i^t}(\mathbf{z}^{y^t} | \mathbf{x}^t) = \frac{\exp(\kappa \langle \Pi_{\mathcal{N}}(\mathbf{z}^{y^t}), \Pi_{\mathcal{N}}(\boldsymbol{\xi}^t) \rangle)}{\sum_{y \in \mathcal{Y}_i^t} \exp(\kappa \langle \Pi_{\mathcal{N}}(\mathbf{z}^y), \Pi_{\mathcal{N}}(\boldsymbol{\xi}^t) \rangle)}, \quad (15)$$

where $\boldsymbol{\xi}^t = f_{\boldsymbol{\theta}, \varphi}(\mathbf{x}^t)$, $\mathcal{Y}_i^t = \text{Unique}(\{y\}_{(\mathbf{x}, y) \in \mathcal{B}_i^t})$ is the set of classes that appears in \mathcal{B}_i^t , $\Pi_{\mathcal{N}}(\cdot) := \frac{(\cdot)}{\|\cdot\|}$ denotes the projection operator. κ is set as a hyperparameter. Hence, our training approach involves alternating between the EM steps based on eq.(8) and eq.(9) throughout the training process. Nonetheless, when delving deeper into the optimization process, it becomes imperative to consider the gradient of $\Pi_{\mathcal{N}}(\mathbf{z}^{y^t})$:

$$\begin{aligned} & \nabla_{\Pi_{\mathcal{N}}(\mathbf{z}^y)} \left[-\log \left(p_{\boldsymbol{\theta}, \varphi}^{\mathcal{B}_i^t}(\mathbf{z}^{y^t} | \mathbf{x}^t) \right) \right] \\ &= \begin{cases} \left[p_{\boldsymbol{\theta}, \varphi}^{\mathcal{B}_i^t}(\mathbf{z}^{y^t} | \mathbf{x}^t) - 1 \right] \kappa \Pi_{\mathcal{N}}(\boldsymbol{\xi}^t), & y = y^t \\ p_{\boldsymbol{\theta}, \varphi}^{\mathcal{B}_i^t}(\mathbf{z}^y | \mathbf{x}^t) \kappa \Pi_{\mathcal{N}}(\boldsymbol{\xi}^t), & y \neq y^t \end{cases}, \end{aligned} \quad (16)$$

thus when $p_{\boldsymbol{\theta}, \varphi}^{\mathcal{B}_i^t}(\mathbf{z}^{y^t} | \mathbf{x}^t) \rightarrow 1$, the gradient for $\Pi_{\mathcal{N}}(\mathbf{z}^y)$ ($y = y^t$) will tend to zero, meanwhile, as p approaches 0, $\nabla_{\Pi_{\mathcal{N}}(\mathbf{z}^y)} \left[-\log \left(p_{\boldsymbol{\theta}, \varphi}^{\mathcal{B}_i^t}(\mathbf{z}^{y^t} | \mathbf{x}^t) \right) \right]$ ($y \neq y^t$) will also tend to zero. Likewise, for $\Pi_{\mathcal{N}}(\boldsymbol{\xi}^t)$ we have:

$$\begin{aligned} & \nabla_{\Pi_{\mathcal{N}}(\boldsymbol{\xi}^t)} \left[-\log \left(p_{\boldsymbol{\theta}, \varphi}^{\mathcal{B}_i^t}(\mathbf{z}^{y^t} | \mathbf{x}^t) \right) \right] \\ &= \kappa \left[\sum_{y \in \mathcal{Y}_i^t} p_{\boldsymbol{\theta}, \varphi}^{\mathcal{B}_i^t}(\mathbf{z}^y | \mathbf{x}^t) \Pi_{\mathcal{N}}(\mathbf{z}^y) - \Pi_{\mathcal{N}}(\mathbf{z}^{y^t}) \right]. \end{aligned} \quad (17)$$

Note that when $p_{\boldsymbol{\theta}, \varphi}^{\mathcal{B}_i^t}(\mathbf{z}^{y^t} | \mathbf{x}^t) \rightarrow 1$,

$$\sum_{y \in \mathcal{Y}_i^t} p_{\boldsymbol{\theta}, \varphi}^{\mathcal{B}_i^t}(\mathbf{z}^y | \mathbf{x}^t) \Pi_{\mathcal{N}}(\mathbf{z}^y) \rightarrow \Pi_{\mathcal{N}}(\mathbf{z}^{y^t}). \quad (18)$$

This means that $\nabla_{\Pi_{\mathcal{N}}(\boldsymbol{\xi}^t)} \left[-\log \left(p_{\boldsymbol{\theta}, \varphi}^{\mathcal{B}_i^t}(\mathbf{z}^{y^t} | \mathbf{x}^t) \right) \right]$ will tend to zero. To ensure stable gradients during training and achieve consistent loss reduction, we introduce a gradient stabilization loss

$$\ell_{\text{GS}}^{\delta}(\mathbf{x}^t, y^t) = \ell_{\text{margin}}^{\delta}(\mathbf{z}^{y^t}, \boldsymbol{\xi}^t) + \sum_{y \in \mathcal{Y}_i^t, y \neq y^t} \ell_{\text{margin}}^{\delta}(-\mathbf{z}^y, \boldsymbol{\xi}^t), \quad (19)$$

where

$$\ell_{\text{margin}}^{\delta}(\mathbf{z}^y, \boldsymbol{\xi}^t) = \max(|1 - \langle \Pi_{\mathcal{N}}(\mathbf{z}^y), \Pi_{\mathcal{N}}(\boldsymbol{\xi}^t) \rangle| - \delta), \quad (20)$$

which provides a constant gradient as compensation, with δ regulating the threshold of gradient compensation. When the loss becomes small, no further compensation is provided due to the instability of the zero point for the absolute value function. Similarly, we incorporate the batchwise gradient stabilization loss $\mathcal{L}_{\text{GS}}^{\delta}(\mathcal{B}_i^t) = \frac{1}{|\mathcal{B}_i^t|} \sum_{(\mathbf{x}^t, y^t) \in \mathcal{B}_i^t} \ell_{\text{GS}}^{\delta}(\mathbf{x}^t, y^t)$ into the objective, leading to the overall loss:

$$\mathcal{L}(\mathcal{B}_i^t) = \mathcal{L}_{\text{vMF}}(\mathcal{B}_i^t) + \lambda \mathcal{L}_{\text{GS}}^{\delta}(\mathcal{B}_i^t), \quad (21)$$

here, δ and λ are hyperparameters.

The CL-vMF model possesses several advantageous features, including incremental value memory and the ability to handle an arbitrary number of classes. As a result, there is no need to retrain the classification head even when the number of classes exceeds the pre-defined limit. Since the value memory parameters \mathcal{Z}^t are frozen after the completion of task t , i.e. all memories $\mathcal{Z}^1, \dots, \mathcal{Z}^{t-1}$ remain unchanged both before and after training task t . This guarantees stable and persistent value memory. Moreover, the storage cost of class value memory parameters is calculated as $d_{z_\tau} \times N_C + d_{z_c} \times |\mathcal{T}|$, where d_{z_c}, d_{z_τ} and N_C represent the dimensions of z_τ, z_c and the number of classes, respectively. Importantly, this implies that the storage cost of value memory parameters scales linearly with N_C .

3.4. Collaborative Inference: System1 and System2

To align with the CLS, we aim for the ICL framework to activate System2 when System1 fails to perform fast thinking, i.e. when encounters hard samples. This activation leverages the complex reasoning capabilities of System2 to achieve collaborative inference. Specifically, we use MLLM to instantiate g_ψ . And we propose a hard sample filtering strategy, vMF-ODI, to screen data that challenges the System1. Specifically, we use batchwise normalization to filter outliers, thus identifying hard sample set $\tilde{\mathcal{X}}$:

$$\tilde{\mathcal{X}}_i = \{(\tilde{x}, \tilde{y}) \in \mathcal{B}_i \mid (\nu - \bar{\nu}_{\mathcal{B}_i})/\sigma_{\mathcal{B}_i} < \alpha\}, \quad (22)$$

where $\nu = \frac{\langle f_{\theta, \varphi}(x), z^{\hat{y}} \rangle}{\|f_{\theta, \varphi}(x)\| \|z^{\hat{y}}\|}$, \hat{y} is the predicted label, $\nu_{\mathcal{B}_i} = \frac{1}{|\mathcal{B}_i|} \sum_{j \in \mathcal{B}_i} \nu_j$, $\sigma = \sqrt{\frac{1}{|\mathcal{B}_i|-1} \sum_{j \in \mathcal{B}_i} (\nu_j - \bar{\nu})^2}$ and α is a detection threshold. For the filtered hard samples, we utilize the TopK outputs $[f_{\theta, \varphi}(\tilde{x})]^K$ from the System1 to construct an inquiry-based language prompt $\mathcal{S}([f_{\theta, \varphi}(\tilde{x})]^K)$. The operation \mathcal{S} converts labels to language and prompts the System2 to perform reasoning and rank the results based on the given context, and finally gives the prediction $\hat{\tilde{y}}$ like we stated in eq. (1). After completing the inference in stage 2, if the System2 provides a precise answer, we will use that answer. Otherwise, we will rely on the judgment result from the System1. This interactive scheme suggests the possibility of using fine-tuning strategies, such as LoRA [13], to minimize $-\log p_\psi(\mathcal{S}(y) \mid x; \mathcal{S}([f_{\theta, \varphi}(x)]^K))$ with respect to ψ . This adjustment would help align the System2 with the System1. Consistent with existing CL setups, we focus on the parameter updates of System1. Detailed algorithmic description can be found in Appendix A. During the training process, we perform T iterations of EM alternating steps for each task, with updates only applied to θ in subsequent updates for that task.

4. Experiments

4.1. Experimental Setups

Following [41, 43], we investigate two common CL setups: Task-Incremental Learning (Task IL) and Class-

Incremental Learning (Class IL). In Task IL, task identifiers are provided during both the training and testing phases. In contrast, the Class IL protocol assigns task identifiers only during the training phase. During the testing phase, the model faces the challenge of predicting all classes encountered up to that point, making it a more demanding scenario.

Datasets. Three datasets are used in our experiments: CIFAR10 [15], CIFAR100 [15], ImageNet-R [11]. Details of these datasets are in Appendix.

Baselines. We combine the proposed ICL framework with the following advanced baselines: Selected rehearsal-based methods: ER [34], A-GEM [6], iCaRL [32], CVT [44], SCoMMER [37], BiMeCo [25]. architecture-based methods: DualNet [28], L2P [46], DualPrompt [45]. Additionally, we perform comparisons with well-known regularization-based techniques, namely EWC [14], LwF [19], and SI [50]. Additionally, we assess JOINT, which entails supervised fine-tuning across all task training sets and represents an upper performance limit. We also examine FT-seq, a sequential fine-tuning technique that partially freezes pre-training weights and generally serves as a performance lower bound. Both JOINT and FT-seq have two variants, one using ViT [8] and the other using ResNet18 [10] as the backbone. Our main focus on comparing these methods with rehearsal- and prompt-based methods.

Metrics. We evaluate CL methods in terms of accuracy, following the accuracy definition presented in [3, 43, 45]: $\mathbf{A}_T = \frac{1}{T} \sum_{t=1}^T a_{T,t}$, where $a_{T,t}$ represents the testing accuracy for task \mathcal{T}_t when the model has completed learning task \mathcal{T}_T .

Implementation Details. Following [3, 4, 32, 44, 50], we use RestNet18 [10], ViT [8] as the backbone in System1. For System2, we choose MiniGPT4 [53], Inf-MLLM [52], Pure-MM as the backbone. More details are in Appendix.

4.2. Results

Extensive experiments are conducted on CIFAR10, CIFAR100, ImageNet-R. Specifically, we add our ICL to several state-of-the-art methods to evaluate its effectiveness.

Results of comparisons with State-of-the-Art Methods.

The quantitative comparisons are summarized in Tab. 1. In the rehearsal-based method, we experiment with different buffer sizes for comparison. The results consistently show that our method outperforms others. Notably, to better simulate the CL scenario, we restrict the number of epochs to one, allowing the model to encounter the data only once during incremental task learning. This restriction significantly affects the efficacy of regularization techniques like EWC [14], LwF [19], as well as rehearsal-based methods such as ER [34], A-GEM [6], iCaRL [32] and CVT [44], among others. However, transformer-based approaches like L2P [46] and DualPrompt [45] maintain certain performance. Additionally, even in the absence of integrating the System2, its incorporation leads to further performance en-

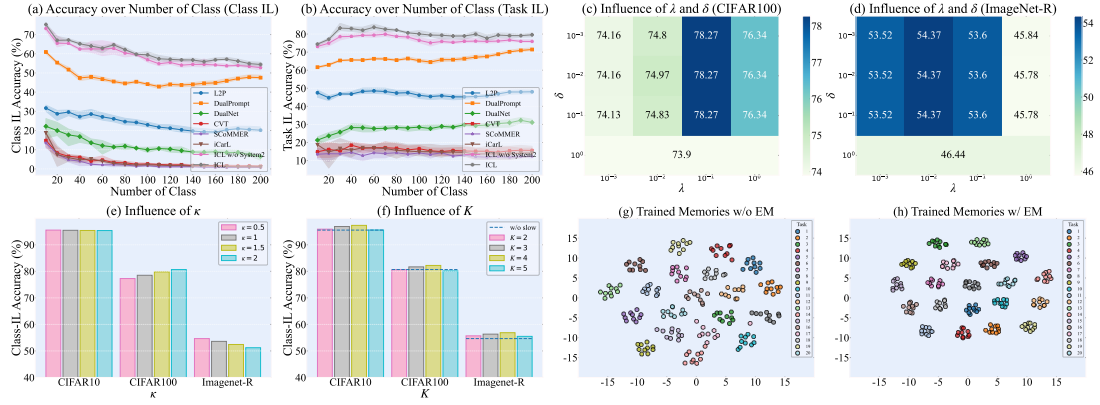


Figure 2. Further analysis of the proposed ICL. (a) and (b) are forgetting curves of different methods on the ImageNet-R in the Class IL and Task IL scenario respectively. (c) and (d) are the impact of regularization parameters λ and δ on CIFAR100 and ImageNet-R respectively. (e) The impact of concentration κ , is evaluated at values of 0.5, 1, 1.5, and 2, respectively. (f) The impact of the number of category choices K in the prompt, is evaluated at values of 2, 3, 4, and 5 respectively. (g) The impact of training memory without EM strategy. (h) The impact of training memory with EM strategy.

Backbone	Method Type	Memory Buffer	Method	CIFAR10		CIFAR100		ImageNet-R	
				Class-IL	Task-IL	Class-IL	Task-IL	Class-IL	Task-IL
ResNet18	Non-Rehearsal based	0	JOINT	92.20	98.31	70.62	86.19	7.72	25.48
			FT-seq	19.62	61.02	17.58	40.46	0.59	10.82
			EWCI[14]	17.82	83.52	7.62	55.14	1.08	21.34
			LwF[19]	18.52	84.72	8.88	61.32	1.24	45.68
			SIJ[50]	18.41	84.74	6.73	50.44	3.31	22.72
			ER[34]	44.79	91.19	21.40	61.36	1.01	15.36
	Rehearsal based	200	A-GEM [6]	18.58	80.19	7.97	55.20	1.23	16.24
			iCaRL [32]	23.80	67.82	7.31	33.10	0.81	9.20
			CVT [44]	30.74	75.92	12.09	43.14	1.60	9.01
			SCoMMER [37]	66.35	92.66	38.89	67.62	1.73	10.65
			DualNet [28]	24.50	90.70	25.30	54.60	7.01	20.70
			BiMeCo [25]	27.92	92.75	28.71	56.65	10.41	22.75
ViT	Non-Rehearsal based	0	JOINT	97.49	99.54	87.23	97.64	74.75	83.39
			FT-seq	22.32	86.33	20.48	83.97	32.56	49.62
			L2P [46]	92.22	98.99	79.68	96.24	48.68	65.38
			DualPrompt [45]	94.43	99.32	79.98	95.92	52.20	69.22
			L2P [46]	67.13	96.39	65.29	92.16	36.70	55.35
			DualPrompt [45]	70.60	97.78	65.97	92.85	38.79	59.32
	Rehearsal based	200	ICL w/o System2	94.60	99.43	77.34	94.81	49.87	68.62
			ICL w MiniGPT4	95.34	99.56	78.28	95.70	52.46	69.87
			ICL w Inf-MLLM	95.42	99.56	78.55	95.83	53.20	72.96
			ICL w Pure-MM	95.94	99.57	79.12	95.99	53.64	73.59
			L2P [46]	71.23	96.78	69.43	93.92	40.17	57.89
			DualPrompt [45]	73.56	98.12	69.98	93.76	43.77	61.24
Rehearsal based	500/600	ICL w/o System2	95.54	99.52	80.67	95.24	54.65	76.02	
		ICL w MiniGPT4	96.49	99.58	81.38	95.62	55.99	79.68	
		ICL w Inf-MLLM	96.69	99.64	82.29	96.14	57.47	81.82	
		ICL w Pure-MM	96.83	99.68	82.43	96.35	58.18	82.64	
		L2P [46]	71.23	96.78	69.43	93.92	40.17	57.89	
		DualPrompt [45]	73.56	98.12	69.98	93.76	43.77	61.24	

Table 1. Comparison of average incremental accuracy (%) with different continual learning method in two scenarios. Memory Buffer denotes the size of the memory buffer area (0 means no rehearsal is used). Note that when training with the ViT in System1, the buffer size of the Imagenet-R is 600, due to the need to allocate an equal amount of buffer area for each class.

hancement. This is demonstrated by a notable increase of over 3% in CL accuracy on the ImageNet-R dataset, observed across 10 consecutive splits, with a memory capacity of 600. The introduction of the System2 enhances System1 ability to effectively recognize and address previously forgotten images or information. To ensure fair comparisons, we incorporate a buffer into the L2P [46] and DualPrompt [45] methods, using a fixed strategy akin to CVT [44]. Subsequently, we observe that the performance of the L2P [46] and DualPrompt [45] methods declined after integrating the buffer. This decline is likely due to interference from the replayed samples in the buffer, which hampers the training of task-specific prompts and leads to greater forgetting.

Results of different task settings. To assess CL strategies across varying numbers of data streams, following the protocol outlined in [44, 49]. This method entails dividing the ImageNet-R dataset, which consists of 200 classes, into subsets containing 5, 10, and 20 classes each, thereby creating incremental tasks. Tab. 2 offers a comprehensive analysis of the accuracy achieved by various methods across different task configurations. The results unequivocally establish the significant superiority of our method over both regularization-based and rehearsal-based method in a wide range of incremental division scenarios. Notably, even when the memory capacity is held constant at 200, our method outperforms architecture-based methods such

Memory Buffer	Method	5 splits		10 splits		20 splits	
		Class-IL	Task-IL	Class-IL	Task-IL	Class-IL	Task-IL
0	ER[34]	1.56	11.35	1.08	21.34	8.10	12.68
	LwF[19]	1.38	14.66	1.24	45.68	0.77	14.34
	SI[50]	1.78	11.50	3.31	22.72	3.35	40.29
	L2P[46]	29.87	38.58	48.68	65.38	20.08	47.98
	DualPrompt[45]	54.43	66.86	52.20	69.22	47.13	71.43
	ICL (ours)	54.43	66.86	52.20	69.22	47.13	71.43
200	ER[34]	1.29	9.75	1.01	15.36	1.38	19.81
	A-GEM[6]	1.30	4.23	1.23	16.24	1.34	21.78
	iCaRL[32]	0.41	3.22	0.81	9.20	0.83	15.23
	CVT[44]	1.47	5.65	1.6	9.01	1.01	13.19
	SCoMMER[37]	0.80	3.78	1.73	10.65	0.32	10.36
	DualNet[28]	9.32	13.14	7.01	20.70	4.92	25.53
	BiMeCo[25]	11.18	14.27	10.41	22.75	5.86	26.33
	ICL w/o System2 (ours)	49.91	74.36	49.87	68.62	48.75	73.95
	ICL (ours)	54.85	75.57	52.46	69.87	49.98	74.63
	ICL (ours)	54.85	75.57	52.46	69.87	49.98	74.63
500/600	ER[34]	1.30	13.02	1.27	22.84	1.56	20.42
	A-GEM[6]	1.31	9.55	1.23	19.35	1.58	21.89
	iCaRL[32]	0.43	3.62	1.01	7.60	1.62	14.23
	CVT[44]	1.95	5.68	1.24	6.97	1.45	15.58
	SCoMMER[37]	0.61	3.35	1.40	10.05	0.56	12.60
	DualNet[28]	10.03	13.44	8.70	20.40	6.40	31.80
	BiMeCo[25]	11.89	14.57	12.13	22.45	7.34	32.73
	ICL w/o System2 (ours)	54.60	75.57	54.65	76.02	52.46	77.05
	ICL (ours)	56.34	78.36	55.99	79.68	53.60	80.47
	ICL (ours)	56.34	78.36	55.99	79.68	53.60	80.47

Table 2. Comparison under different task number settings on ImageNet-R dataset.

as L2P [46] and DualPrompt [45]. Furthermore, the results illustrate a consistent enhancement in our method performance as the buffer size increases.

Results of forgetting curve comparison. To illustrate the forgetting process within each compared methods in the CL data stream, we record the average test accuracy for both the current and preceding tasks upon completing the training of each task. Subsequently, we create a line chart to visualize the change in accuracy with the addition of each task, offering a visual representation of the forgetting process. Fig. 2(a) and Fig. 2(b) provide clear illustrations that, as new tasks are introduced, most methods exhibit a decline in performance. However, our method consistently outperforms these methods in terms of accuracy at every stage.

4.3. Ablation Study

Analysis of hyperparameters λ and margin δ .

Fig. 2(c) and Fig. 2(d) depict the impact of parameters λ and δ on our method performance, using the CIFAR100 and ImageNet-R datasets, with memory sizes set at 500 and 600, respectively. An analysis of the heat maps reveals that our method’s sensitivity to λ varies between datasets. Specifically, CIFAR100 demonstrates optimal performance when $\lambda = 0.1$, while ImageNet-R maintains consistent performance with λ values within the range [0.001, 0.01, 0.1]. In contrast, variations in δ do not result in substantial accuracy differences between the datasets. This suggests that the margin regularization term has only a marginal impact on overall model performance. Importantly, when $\lambda = 0$ and $\delta = 1$ (as shown in the bottom), the absence of the regularization term leads to diminished performance, underscoring the value of incorporating margin regularization to enhance the model’s effectiveness.

Analysis of concentration parameter Selection. We conduct an investigation into the influence of the concentration parameter κ on model performance. Fig. 2(e) indicates that there is no significant variation in performance across dif-

ferent values of κ for the CIFAR10 dataset. However, for datasets such as CIFAR100 and ImageNet-R, it becomes important to estimate the concentration in advance to ensure a more accurate modeling of the concentration. Consequently, these datasets demonstrate a noticeable sensitivity to the concentration parameter.

Impact of Top-k results in collaboration inference. To investigate the influence of the number of prompt categories K on the reasoning process of System2, a series of experiments are conducted using different values of K on multiple datasets, as depicted in Fig. 2(f). The results clearly indicate that memory vectors possessing well-defined geometric structures facilitate stable memory retrieval, thereby preventing System1 from deviating significantly from the central data point during the reasoning phase. Consequently, when a smaller number of top-k choices are provided, it can ensure almost guaranteed inclusion of the correct categories. Conversely, when K is larger, there is an increased tendency for erroneous category information to be included in the prompt, further intensifying the potential for confusion in System2. Fig. 2(f) also presents the accuracy achieved without System2 across various datasets (blue horizontal line in the figure), demonstrating the performance enhancement by the incorporation of System2.

Impact of separate query and value memory optimization. To assess the query-value separation strategy impact on stable value memory modeling in persistent scenarios, we compare its use (Fig. 2(g)) and absence (Fig. 2(h)). Visualizing the value memory via tsne reduction reveals that query-value parameter optimization can create more focused value memory modeling, enhancing task discrimination. The corresponding results consistently validate the effectiveness of our proposed CL-vMF mechanism.

5. Conclusion

In our paper, we introduced ICL, a groundbreaking continual learning (CL) paradigm inspired by Complementary Learning System theory in neurocognitive science. ICL combines ViT with an interactive query and value memory module powered by CKT-MHA, enhancing the efficiency of fast thinking (System1). Additionally, it leverages our CL-vMF mechanism to improve memory representation distinction. ICL also integrates multi-modal Large Language Models (System2) with System1 for advanced reasoning, dynamically modulated by hard examples detected through our VMF-ODI strategy. Our experiments confirmed the effectiveness of our framework in reducing forgetting, surpassing contemporary state-of-the-art methods.

Acknowledgement This work was supported in part by the National Key R&D Program of China (No. 2023YFC3305102). We extend our gratitude to the anonymous reviewers for their insightful feedback, which has greatly contributed to the improvement of this paper.

References

- [1] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 139–154, 2018. **2**
- [2] Arindam Banerjee, Inderjit S. Dhillon, Joydeep Ghosh, and Suvrit Sra. Clustering on the unit hypersphere using von mises-fisher distributions. *J. Mach. Learn. Res.*, 6: 1345–1382, 2005. **5**
- [3] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems*, 33:15920–15930, 2020. **6, 1**
- [4] Pietro Buzzega, Matteo Boschini, Angelo Porrello, and Simone Calderara. Rethinking experience replay: a bag of tricks for continual learning. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 2180–2187. IEEE, 2021. **6, 1**
- [5] Yihan Cao, Siyu Li, Yixin Liu, Zhiling Yan, Yutong Dai, Philip S Yu, and Lichao Sun. A comprehensive survey of ai-generated content (aigc): A history of generative ai from gan to chatgpt. *arXiv preprint arXiv:2303.04226*, 2023. **1, 2**
- [6] Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. *arXiv preprint arXiv:1812.00420*, 2018. **6, 7, 8**
- [7] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc’Aurelio Ranzato. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486*, 2019. **2**
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. **2, 6**
- [9] Jonathan St BT Evans. In two minds: dual-process accounts of reasoning. *Trends in cognitive sciences*, 7(10):454–459, 2003. **1**
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. **6**
- [11] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8349, 2021. **6**
- [12] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 831–839, 2019. **2**
- [13] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2021. **6**
- [14] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, 2017. **1, 2, 6, 7, 8**
- [15] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. **6**
- [16] Dharshan Kumaran, Demis Hassabis, and James L McClelland. What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends in cognitive sciences*, 20(7):512–534, 2016. **1**
- [17] Janghyeon Lee, Hyeong Gwon Hong, Donggyu Joo, and Junmo Kim. Continual learning with extended kronecker-factored approximate curvature. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9001–9010, 2020. **1, 2**
- [18] Omer Levy, Yoav Goldberg, and Ido Dagan. Improving distributional similarity with lessons learned from word embeddings. *Transactions of the Association for Computational Linguistics*, 3:211–225, 2015. **5**
- [19] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. **6, 7, 8**
- [20] Xialei Liu, Marc Masana, Luis Herranz, Joost Van de Weijer, Antonio M Lopez, and Andrew D Bagdanov. Rotate your networks: Better weight consolidation and less catastrophic forgetting. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 2262–2268. IEEE, 2018. **2**
- [21] Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7765–7773, 2018. **1, 2**
- [22] Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *Proceedings of the European conference on computer vision (ECCV)*, pages 67–82, 2018. **1, 2**
- [23] Yu Meng, Jiaming Shen, Chao Zhang, and Jiawei Han. Weakly-supervised hierarchical text classification. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 6826–6833, 2019. **5**
- [24] Martial Mermillod, Aurélie Bugaiska, and Patrick Bonin. The stability-plasticity dilemma: Investigating the continuum from catastrophic forgetting to age-limited learning effects, 2013. **1**
- [25] Xing Nie, Shixiong Xu, Xiyan Liu, Gaofeng Meng, Chunlei Huo, and Shiming Xiang. Bilateral memory consolidation for continual learning. In *Proceedings of the IEEE/CVF Con-*

- ference on Computer Vision and Pattern Recognition, pages 16026–16035, 2023. 1, 6, 7, 8
- [26] Randall C O’Reilly and Kenneth A Norman. Hippocampal and neocortical contributions to memory: Advances in the complementary learning systems framework. *Trends in cognitive sciences*, 6(12):505–510, 2002. 1
- [27] Dongmin Park, Seokil Hong, Bohyung Han, and Kyoung Mu Lee. Continual learning by asymmetric loss approximation with single-side overestimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3335–3344, 2019. 2
- [28] Quang Pham, Chenghao Liu, and Steven Hoi. Dualnet: Continual learning, fast and slow. *Advances in Neural Information Processing Systems*, 34:16131–16144, 2021. 1, 2, 6, 7, 8
- [29] Biqing Qi, Bowen Zhou, Weinan Zhang, Jianxing Liu, and Ligang Wu. Improving robustness of intent detection under adversarial attacks: A geometric constraint perspective. *IEEE transactions on neural networks and learning systems*, PP, 2023. 5
- [30] Biqing Qi, Bowen Zhou, Weinan Zhang, Jianxing Liu, and Ligang Wu. Improving robustness of intent detection under adversarial attacks: A geometric constraint perspective. *IEEE Transactions on Neural Networks and Learning Systems*, 2023. 1
- [31] Qi Qin, Wenpeng Hu, Han Peng, Dongyan Zhao, and Bing Liu. Bns: Building network structures dynamically for continual learning. *Advances in Neural Information Processing Systems*, 34:20608–20620, 2021. 1, 2
- [32] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017. 1, 6, 7, 8
- [33] Blake A Richards and Paul W Frankland. The persistence and transience of memory. *Neuron*, 94(6):1071–1084, 2017. 1
- [34] Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. Learning to learn without forgetting by maximizing transfer and minimizing interference. *arXiv preprint arXiv:1810.11910*, 2018. 2, 6, 7, 8
- [35] Tomás J Ryan and Paul W Frankland. Forgetting as a form of adaptive engram cell plasticity. *Nature Reviews Neuroscience*, 23(3):173–186, 2022. 1
- [36] Gobinda Saha, Isha Garg, and Kaushik Roy. Gradient projection memory for continual learning. *arXiv preprint arXiv:2103.09762*, 2021. 1, 2
- [37] Fahad Sarfraz, Elahe Arani, and Bahram Zonooz. Sparse coding in a dual memory system for lifelong learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9714–9722, 2023. 6, 7, 8
- [38] John Schulman, Barret Zoph, Christina Kim, Jacob Hilton, Jacob Menick, Jiayi Weng, Juan Felipe Ceron Uribe, Liam Fedus, Luke Metz, Michael Pokorny, et al. Chatgpt: Optimizing language models for dialogue. *OpenAI blog*, 2022. 2
- [39] Advani Sun, Weinan et al. Organizing memories for generalization in complementary learning systems. *Nature neuroscience*, 26(8):1438–1448, 2023. 1
- [40] Shixiang Tang, Dapeng Chen, Jinguo Zhu, Shijie Yu, and Wanli Ouyang. Layerwise optimization by gradient decomposition for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9634–9643, 2021. 1, 2
- [41] Gido M Van de Ven and Andreas S Tolias. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*, 2019. 6
- [42] Liyuan Wang, Xingxing Zhang, Hang Su, and Jun Zhu. A comprehensive survey of continual learning: Theory, method and application. *arXiv preprint arXiv:2302.00487*, 2023. 2
- [43] Zhen Wang, Liu Liu, Yiqun Duan, Yajing Kong, and Dacheng Tao. Continual learning with lifelong vision transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 171–181, 2022. 6
- [44] Zhen Wang, Liu Liu, Yajing Kong, Jiaxian Guo, and Dacheng Tao. Online continual learning with contrastive vision transformer. In *ECCV*, 2022. 1, 2, 6, 7, 8
- [45] Zifeng Wang, Zhang, et al. Dualprompt: Complementary prompting for rehearsal-free continual learning. In *ECCV*, pages 631–648, 2022. 1, 2, 3, 6, 7, 8
- [46] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 139–149, 2022. 1, 2, 3, 6, 7, 8
- [47] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022. 2
- [48] Gordon Winocur, Morris Moscovitch, and Melanie Sekeres. Memory consolidation or transformation: context manipulation and hippocampal representations of memory. *Nature neuroscience*, 10(5):555–557, 2007. 1
- [49] Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3014–3023, 2021. 7
- [50] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pages 3987–3995. PMLR, 2017. 1, 2, 6, 7, 8
- [51] Chao Zhang, Liyuan Liu, Dongming Lei, Quan Yuan, Honglei Zhuang, Timothy Hanratty, and Jiawei Han. Triovevent: Embedding-based online local event detection in geo-tagged tweet streams. pages 595–604, 2017. 5
- [52] Qiang Zhou, Zhibin Wang, Wei Chu, Yinghui Xu, Hao Li, and Yuan Qi. Infmlm: A unified framework for visual-language tasks, 2023. 6
- [53] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language

understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023. [6](#), [2](#)