

Dynamically Anchored Prompting for Task-Imbalanced Continual Learning

Chenxing Hong¹, Yan Jin^{1,2}, Zhiqi Kang⁴, Yizhou Chen¹, Mengke Li³, Yang Lu^{1,2*} and Hanzi Wang^{1,2}

¹Key Laboratory of Multimedia Trusted Perception and Efficient Computing, Ministry of Education of China, Xiamen University, Xiamen, China

²Fujian Key Laboratory of Sensing and Computing for Smart City, School of Informatics, Xiamen University, Xiamen, China

³Guangdong Laboratory of Artificial Intelligence and Digital Economy (SZ), Shenzhen, China

⁴Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France.

hongchenxing@stu.xmu.edu.cn, jinyan7973@gmail.com, zhiqi.kang@inria.fr, 36920231153184@stu.xmu.edu.cn, limengke@gml.ac.cn, {luyang, hanzi.wang}@xmu.edu.cn

Abstract

Existing continual learning literature relies heavily on a strong assumption that tasks arrive with a balanced data stream, which is often unrealistic in real-world applications. In this work, we explore task-imbalanced continual learning (TICL) scenarios where the distribution of task data is non-uniform across the whole learning process. We find that imbalanced tasks significantly challenge the capability of models to control the trade-off between stability and plasticity from the perspective of recent prompt-based continual learning methods. On top of the above finding, we propose Dynamically Anchored Prompting (DAP), a prompt-based method that only maintains a single general prompt to adapt to the shifts within a task stream dynamically. This general prompt is regularized in the prompt space with two specifically designed prompt anchors, called boosting anchor and stabilizing anchor, to balance stability and plasticity in TICL. Remarkably, DAP achieves this balance by only storing a prompt across the data stream, therefore offering a substantial advantage in rehearsal-free CL. Extensive experiments demonstrate that the proposed DAP results in 4.5% to 15% absolute improvements over state-of-the-art methods on benchmarks under task-imbalanced settings. Our code is available at <https://github.com/chenxing6666/DAP>.

1 Introduction

Human beings possess the remarkable ability to learn new tasks and solve evolving challenges by leveraging knowledge from their past experiences. Inspired by this, continual learning (CL) methods are designed to address a series of tasks using a singular model while preserving performance on tasks previously mastered [Wang *et al.*, 2022c; Zhang *et al.*, 2020; Aljundi *et al.*, 2018]. However, achieving

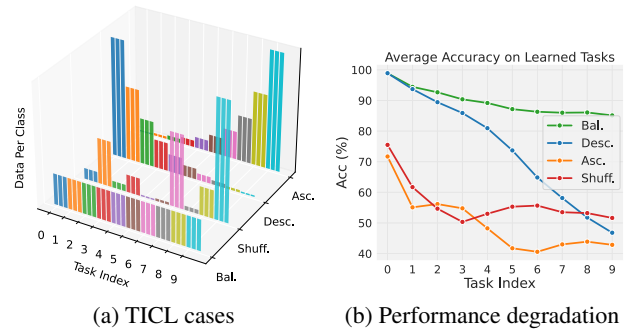


Figure 1: An illustration of the scenario of task-imbalanced continual learning (TICL). (a) Three cases in TICL compared to ordinary task-balanced CL. (b) Performance degradation of DualPrompt on TICL CIFAR-100. The number of balanced and imbalanced tasks is ensured to be the same.

this goal is challenging for deep models, as they tend to easily forget previously learned information, i.e., a phenomenon known as catastrophic forgetting [De Lange *et al.*, 2021; Kirkpatrick *et al.*, 2017]. This issue primarily arises from the network’s tendency to overwrite old knowledge with new data during the training process.

Although existing methods in CL have achieved notable progress, they generally assume a balanced distribution of training data across tasks, i.e., each task holds the same number of training samples. In practical applications, data streams often exhibit an imbalanced distribution [Huang *et al.*, 2019], where the data volume for each task can vary significantly, with some tasks presenting a large number of samples while others may have far fewer. This gives rise to task-imbalanced CL scenarios and introduces potential new issues for the CL process. The imbalance among tasks will amplify the difficulty of striking a balance between learning and retaining knowledge over time.

To address the aforementioned problem, this paper first investigates this more general and realistic scenario: task-imbalanced continual learning (TICL). TICL characterizes environments where the number of samples in each class is

*Corresponding Author: Yang Lu (luyang@xmu.edu.cn)

imbalanced across the data stream, reflecting the long-tailed nature observed in most real-world data distributions. In such distributions, a few classes dominate in terms of the number of samples, while many others have significantly fewer. In the context of CL, this imbalance manifests at the task level, leading to imbalanced tasks, as shown in Figure 1(a). We call the tasks with a relatively large number of samples the *data-rich tasks* and the tasks with a relatively small number of samples the *data-scarce tasks*. Specifically, we consider three cases of TICTL:

- **Descending TICTL:** The learning process starts with data-rich tasks followed by data-scarce ones.
- **Ascending TICTL:** The learning process starts with data-scarce tasks preceding data-rich ones.
- **Shuffled TICTL:** Tasks arrive in a random sequence without a prescribed order.

While descending and ascending TICTL are two extreme cases, shuffled TICTL can be regarded as the general form of TICTL.

To show the challenges brought by the proposed scenarios, we conduct a quick experiment using DualPrompt [Wang *et al.*, 2022d], a typical prompt-based CL method, in Figure 1(b). It can be observed that in the descending TICTL case, the model initially learns well on data-rich tasks but rapidly declines when training becomes scarce. This indicates the poor plasticity of the model in adapting to new tasks with fewer samples. On the contrary, in the ascending TICTL case, the model initially struggles with learning data-scarce tasks and remains poor in the following tasks. This is due to the severe forgetting caused by the following data-rich tasks (see Figure. 2 for more analysis). These results reveal a more realistic plasticity and stability dilemma: given an incoming task, the model learns poorly when its training data is scarce and forgets rapidly when there is abundant data. Undoubtedly, both learning and forgetting become more challenging in TICTL. Therefore, there is a clear need to design a specific method to address this issue.

In this paper, we propose a novel prompt-based approach named Dynamically Anchored Prompting (DAP). DAP only maintains a general prompt to learn from imbalanced task streams by strategically address the dilemma of stability and plasticity. It tackles the dilemma by decoupling stability and plasticity through two specialized prompts: the boosting prompt and the stabilizing prompt. The stabilizing prompt focuses on preserving the knowledge of past tasks, thereby ensuring stability and mitigating catastrophic forgetting. In contrast, the boosting prompt enhances the model’s generalization ability to learn and adapt to new tasks, promoting plasticity. The general prompt is updated by a novel dynamic stability-plasticity regularization (DSPR) strategy, which dynamically regularizes the general prompt in the prompt space based on task attributes, ensuring a flexible and adaptive learning process. Since DAP only stores a general prompt, it achieves superior performance with lower memory requirements, aligning well with the objectives of rehearsal-free CL. In summary, our contributions are three-fold:

- We analyze a more realistic CL scenario with imbalanced tasks along with prompt-based learning algorithms, uncovering their defects caused by the stability and plasticity dilemma. This dilemma highlights a critical challenge:

an effective balance between preserving existing knowledge and accommodating new learning demands.

- We propose DAP, a novel approach to dynamically balance the stability and plasticity with a single regularized general prompt, which effectively addresses the challenges in TICTL.
- We evaluate the performance of DAP on benchmark datasets. Our proposed DAP exhibits significant improvements over previous state-of-the-art methods, with large margins ranging from 4.5% to 15%.

2 Related Work

2.1 Continual Learning

The existing research in CL within machine learning primarily assumes a balanced task distribution and focuses on addressing catastrophic forgetting, which can be categorized into three strategies [De Lange *et al.*, 2021; Wang *et al.*, 2023]: architectural expansion, regularization, and rehearsal. Architectural expansion adapts the model’s structure for new tasks, suitable where adding to the model is practical [Kang *et al.*, 2022; Wang *et al.*, 2022b]. Regularization, either in the weight or prediction space, aims to retain past task knowledge during new task training, with knowledge distillation being particularly effective in prediction space [Castro *et al.*, 2018; Cha *et al.*, 2021]. Rehearsal methods, using original or synthetic data, are efficient but face data privacy and storage issues [Chaudhry *et al.*, 2021; Wang *et al.*, 2022a; Kang *et al.*, 2023]. These issues emphasize the need for rehearsal-free methods in CL, addressing both privacy concerns and computational efficiency [Vaishnavh *et al.*, 2018].

2.2 Prompting for Continual Learning

The recent trend in CL research focuses on combining prompting techniques with Vision Transformers (ViTs). This approach involves using a pre-trained, frozen backbone model from ImageNet, circumventing the need for a replay buffer. Prompting, initially applied in transfer learning with pre-trained language models like GPT-3, involves adding language-based instructions to the input text to guide the model in understanding downstream tasks. Traditional prompting methods were heuristic, but recent developments like Prompt Tuning [Kemker and Kanan, 2017] and Prefix Tuning [Kemker *et al.*, 2018] introduced the concept of learnable prompts in a continuous space, becoming mainstream in prompt-based learning.

In the realm of prompt-based CL, various methods have been proposed. Specifically, L2P [Wang *et al.*, 2022e] introduced a prompt pool concept to adjust the frozen ViT backbone for CL tasks. Building on this, DualPrompt [Wang *et al.*, 2022d] employs two different prompt types: G-Prompt for learning task-invariant knowledge and E-Prompt for task-specific knowledge, drawing inspiration from complementary learning systems. CODA-Prompt [Smith *et al.*, 2023] adopts input-conditioned prompts through an innovative attention-based end-to-end key-query mechanism that integrates the entire training sequence.

2.3 Imbalanced Continual Learning

While there is existing research on addressing imbalance in CL, it primarily concentrates on specific aspects of imbalance. For instance, BIC [Wu *et al.*, 2019] focuses on the imbalance between limited stored samples and current task samples, addressing challenges in storage under highly imbalanced conditions. PRS [Kim *et al.*, 2020] investigates the long-tail problem in multi-label scenarios, also relying on sample storage. Two-Stage-CL [Liu *et al.*, 2022] delves into long-tail class imbalances in CL, storing substantial amounts of original data.

It is noteworthy that these studies are primarily rehearsal-based methods. This underscores the significant challenge of storing samples where class imbalance is inherent, as data selection itself is imbalanced. Furthermore, our approach to TICL diverges fundamentally from these existing works. While they primarily concentrate on addressing imbalances within samples or classes, our focus lies in the imbalance across different tasks.

3 Problem Formulation

3.1 Preliminaries

Our CL protocol adopts the class-incremental CL setting. The training data is denoted as a sequence of T tasks $\mathcal{D} = \{\mathcal{D}_0, \dots, \mathcal{D}_T\}$, where $\mathcal{D}_t = \{(x_i^t, y_i^t)\}, i = 0, \dots, N_t$ is sampled from a joint data distribution in the input and label space $X_t \times Y_t$ at task t , whose size (i.e., the number of samples in this task) is denoted as N_t . The target model is formulated as $f : X \rightarrow Y$, integrating a patch embedding layer f_p , and a backbone f_b consisting of a stack of transformer encoder layers followed by a classifier, thus $f = f_p \circ f_b$. We employ a ViT-Base model, pre-trained on ImageNet, as the frozen feature extractor. In this class-incremental setting, similar to recent prompt-based CL methods [Smith *et al.*, 2023; Wang *et al.*, 2022d; Wang *et al.*, 2022e], task boundaries are clearly defined with no shared classes between them, and task identity is provided only during training.

3.2 Task-Imbalanced Continual Learning

In this paper, we formally study task-imbalanced continual learning (TICL). TICL is characterized by the distribution of task sizes, i.e., N_t , throughout the learning process. Different from the conventional task-balanced CL where N_t is drawn from a uniform distribution, i.e., $N_i = N_j, \forall i \neq j$ and $i, j \in \{0, \dots, T\}$. Any distribution deviating from the uniform one introduces an inherent imbalance among tasks. By the long-tail nature of real-world data distributions, we generally assume that the task sizes also follow the long-tail distribution. Specifically, the task sizes follow $N_{I_0} > \dots > N_{I_T}, \forall I_i < I_j$, where I_0, \dots, I_T are the sorting index based on the task size by non-increasing order.

3.3 Case Study

We first study how deep the model performance might drop on two extreme cases of TICL: descending TICL and ascending TICL. To properly quantify plasticity and stability [Sun *et al.*, 2022], we adopt two metrics to evaluate a model on the task t :

$$P = Acc_{t,t}, \quad F = Acc_{t,t} - Acc_{T,t}. \quad (1)$$

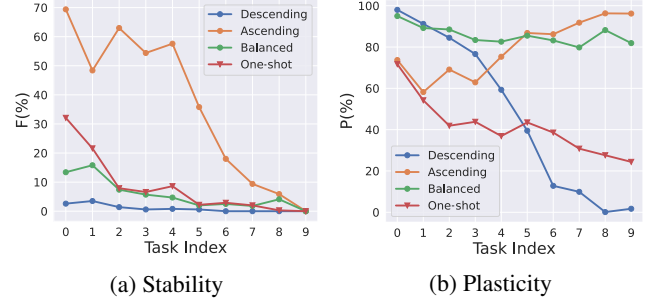


Figure 2: Performance analysis of DualPrompt on different TICL cases with evaluation metrics. (a) F shows the model stability to resist forgetting. (b) P shows the model plasticity to learn new tasks.

where $Acc_{i,j}$ represents the accuracy on the test set of the task j after finishing training on the task i . P measures the models' plasticity, i.e., the ability to learn new tasks, while F measures the model's stability, i.e., the resistance to catastrophic forgetting. Here, we examine the performance of DualPrompt [Wang *et al.*, 2022d] on descending TICL and ascending TICL with evaluation metrics P and F . As a recent prompt-based CL method, it achieves a good balance between stability and plasticity in task-balanced scenarios. In addition, to eliminate the factor of the absolutely small size of each task and focus on the task imbalance problem, we also compare with the case of one-shot, where each task contains only one sample for each class.

Stability. Stability is extremely hard to achieve by a model trained on ascending TICL as catastrophic forgetting occurs more easily on past data-scarce tasks. As shown in Figure 2(a), The model's performance decline on earlier tasks in the ascending case is significant, which indicates severe catastrophic forgetting. The result of the one-shot case reveals that this forgetting is not just due to limited data quantity. On the contrary, in the descending case, the model keeps good stability because the data-rich tasks come first, and the following data-scarce tasks can hardly overtake the knowledge of the data-rich tasks in the model.

Plasticity. On the opposite, plasticity is also extremely hard to achieve by a model trained on descending TICL as the model learns well on past data-rich tasks such that the following data-scarce tasks cannot be generalized well. As shown in Figure 2(b), the model performance in the descending case significantly drops. This raises the issue that the model initially trained on data-rich tasks struggles to adapt to the following data-scarce tasks, prioritizing stability over learning new information. We also compare with the one-shot case. The model shows consistent learning performance starting from the fifth task, surprisingly suggesting better adaptation if the last few tasks are not data-rich.

Accordingly, we propose a method to effectively balance the stability-plasticity dilemma in TICL scenarios.

4 Dynamically Anchored Prompting

To address the issues identified in the case study, we propose Dynamically Anchored Prompting (DAP) for TICL in this paper. The proposed method is designed to dynamically balance and optimize the trade-off between stability and plasticity,

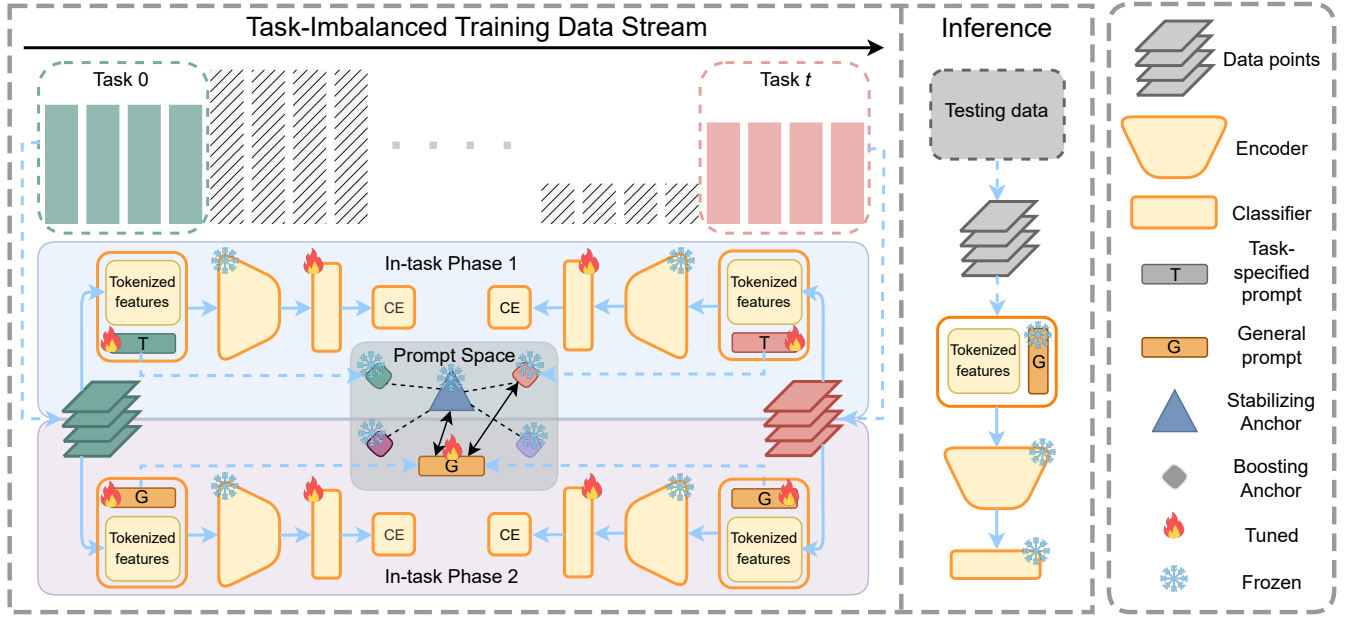


Figure 3: Overview of the proposed dynamic anchored prompting (DAP) for task-imbalanced continual learning. Task-imbalanced training data stream represents the sequential arrival of data from Task 0 to Task t . There are two phases for each task. In in-task phase 1, the initialized task-specific prompt learns knowledge related to the current task, serving as a boosting anchor. In in-task phase 2, the general prompt is trained, with the assistance of boosting anchor to ensure plasticity. Meanwhile, the centers of all previously learned task-specific prompts serve as stabilizing anchors, emphasizing stability. It is worth noting that the boosting anchors of past tasks are not stored as the stabilizing anchor is updated in an online manner.

empowering the model flexible to learn in TICL.

Different from the existing prompt-based CL methods [Wang *et al.*, 2022d; Smith *et al.*, 2023; Wang *et al.*, 2022e], the key idea of DAP is to maintain a prompt across all the tasks, which is called the *general prompt*. The general prompt aims to learn to generalize each task with the knowledge of the pre-trained model. It can be easily used during inference as there is no prompt pool, and therefore no prompt selection is needed. During the inference, the challenge of adopting only a general prompt is to balance the knowledge learned from each task, especially for imbalanced tasks. Therefore, to achieve the goal of using the general prompt to generalize well across all tasks, the proposed DAP adopts a two-phase learning scheme for each training task, termed in-task phases. Given a specific training task, during the first in-task phase, we optimize a *task-specific prompt*, which is then utilized to regularize the general prompt during the second in-task phase. To tackle the imbalance problem of TICL, we propose dynamic stability-plasticity regularization to make the general prompt learn well and stably no matter the amount of training data in the next incoming task. After two-phase tuning for each task, we only need to save the general prompt for inference. The overall framework of the proposed DAP is shown in Figure 3.

4.1 Anchored Prompting

The goal of DAP is to obtain a high-quality general prompt that properly addresses the intrinsic problem of TICL. For each task, we first optimize the task-specific prompt in the first in-task phase. The task-specific prompt is used to form two

key components: the *boosting anchor* \mathbf{p}_b and the *stabilizing anchor* \mathbf{p}_s . These anchors are designed to regularize the relationship between the *general prompt* \mathbf{p}_g and each anchor, effectively handling the plasticity and stability dilemma during training. The general prompt is then optimized in the prompt space with a new strategy called dynamic stability-plasticity regularization.

Boosting Anchor. Specifically, the boosting anchor \mathbf{p}_b aims to maintain model plasticity, ensuring adaptability to new tasks. It is especially useful when the size of the current task is small compared to the past tasks. The boosting anchor \mathbf{p}_b is simply set at the task-specific prompt optimized on the current task to capture the task-relevant information. Thus, \mathbf{p}_b functions as a critical focal point, guiding the model towards learning trajectories that maximize plasticity. To optimize \mathbf{p}_b , we formulate the following loss function:

$$\mathcal{L}_1 = \sum_{i=0}^{N_t} l_{CE}(f_b([\mathbf{p}_b, f_p(x_i^t)]), y_i^t) \quad (2)$$

Here, $h(x_i^t)$ and y_i^t represent the patched features of the i -th sample and its corresponding label from task t , respectively. \mathbf{p}_b is initialized in each new task. The concatenation $[\mathbf{p}_b, h(x_i^t)]$ is fed into the pre-trained transformer as the input. \mathcal{L}_1 represents the total loss for the first in-task phase, and l_{CE} is the cross-entropy loss. The goal of minimizing this loss is to make \mathbf{p}_b fully learn the knowledge for adapting the pre-trained model for the current task.

Stabilizing Anchor. To ensure model stability, the stabilizing anchor \mathbf{p}_s is designed to prevent knowledge forgetting from

past tasks by monitoring the learned task-specific prompts. To maintain the knowledge of all of the learned tasks, \mathbf{p}_s is calculated by the weighted center of the boosting anchors of all learned tasks in the prompt space from task 0 to the current task t . The weight is associated with the inverse of the size of each task. To make the algorithm rehearsal-free, i.e., without storing any data or prompt, \mathbf{p}_s can be updated in an online manner:

$$\mathbf{p}_s \leftarrow \frac{1}{\sum_{i=0}^t \frac{1}{N_i}} \left(\sum_{i=0}^{t-1} \frac{\mathbf{p}_s}{N_i} + \frac{\mathbf{p}_b}{N_t} \right). \quad (3)$$

The rationale behind this update is that task-specific prompts are optimized to generalize well in the corresponding task, thus, their weighted average represents stability. The weights in Eq. (3) associated with the corresponding task size N_i emphasize the past task with a smaller size to make the stabilizing prompt uniformly represent the knowledge of past tasks.

Anchor Alignment. The purpose of maintaining the boosting anchor and the stabilizing anchor in each task is to dynamically regularize the learning of the general prompt, such that it can be flexible to balance the stability and plasticity in the TICL scenario. We employ cosine similarity to measure the proximity between the general prompt and anchors in each task for anchor alignment, as it focuses on the orientation rather than the magnitude of vector representations, effectively capturing the inherent relationships between prompts:

$$\mathcal{L}_a(\mathbf{p}_g, \mathbf{p}) = 1 - \frac{\mathbf{p}_g \cdot \mathbf{p}}{\|\mathbf{p}_g\| \|\mathbf{p}\|}, \quad (4)$$

where \mathbf{p} can be either \mathbf{p}_b or \mathbf{p}_s to address the plasticity or the stability, respectively. Different from the methods that simultaneously adopt two sets of prompts [Wang *et al.*, 2022d], the proposed DAP only utilizes the boosting anchors (i.e. the task-specific prompts) as a constant to regularize the general prompt by anchor alignment. Therefore, task-specific prompts are not involved in the final inference process. The advantages of only maintaining the general prompt through the CL process are two-fold. On one hand, it avoids the error produced by matching the improper prompts in the prompt pool [Wang *et al.*, 2022e]. On the other hand, it makes the proposed DAP fully rehearsal-free.

4.2 Dynamic Stability-Plasticity Regularization

Anchored prompting offers an opportunity to address either plasticity or stability individually, yet it falls short in dynamically adapting within complex, imbalanced task streams. Recall the three cases in TICL described in Figure 1(a). Updating the general prompt by anchor alignment with a single anchor may only work for the descending TICL or the ascending TICL. For shuffled TICL, the incoming task cannot be guaranteed to be a data-scarce or data-rich task, such that the emphasis on updating the general prompt cannot be determined. The challenge of achieving balance in such fluctuating scenarios remains unresolved.

To address the remaining issue in the DAP framework, we introduce a strategy that enables the model to adjust its focus between plasticity and stability in the nature of TICL tasks. In the second in-task phase, we optimize the general prompt

\mathbf{p}_g by considering the balance between stability and plasticity. Different from \mathbf{p}_b that is initialized in each new task, \mathbf{p}_g is initialized in task 0 and updated through the whole CL process in order to achieve the ability to generalize well on all learned classes. Thus, we propose dynamic stability-plasticity factor λ as the coefficient between two anchor alignments. The factor λ modulates the balance between stability and plasticity by considering the size of the current task t and the sizes of past tasks:

$$\lambda = \frac{N_t - N_{\min}}{N_{\max} - N_{\min} + \epsilon}, \quad (5)$$

where $N_{\min} = \min\{N_i\}$ and $N_{\max} = \max\{N_i\}$ are the minimum and maximum sizes of learned tasks for $i = 0, \dots, t$, and ϵ is a small positive constant to prevent division by zero. It is basically the min-max normalization that measures the difference between the size of the current task and the minimum size of the learned tasks, which is then normalized into the range of $[0, 1]$. As with many re-weighting techniques in long-tail learning [Peng *et al.*, 2023; Chen *et al.*, 2023; Zhang *et al.*, 2023b], the task size N_t is an important indicator to represent the learning difficulty. It is relatively easier to learn from a task with a larger size. Accordingly, we use the factor λ and $1 - \lambda$ as the regularization coefficient to reflect this relationship. In the second in-task phase, the general prompt updated in the current task t is given by the following loss function with the dynamic stability-plasticity factor λ :

$$\mathcal{L}_2 = \sum_{i=0}^{N_t} l_{CE}(f_b([\mathbf{p}_g, f_p(x_i^t)]), y_i^t) + \lambda \cdot \mathcal{L}_a(\mathbf{p}_g, \mathbf{p}_s) + (1 - \lambda) \cdot \mathcal{L}_a(\mathbf{p}_g, \mathbf{p}_b). \quad (6)$$

Therefore, a larger N_t indicates a smaller λ , enhancing stability to prevent forgetting. On the other hand, a smaller N_t , indicative of a more challenging task, prompts an increase in λ to ensure sufficient plasticity for learning new, complex tasks. If the size of the current task t is the largest or smallest ever, λ then becomes 1 or 0, respectively.

The effective utilization of the prompt anchors and the flexible adjustment of λ according to the task sizes are the core of DAP, which adapts the pre-trained model well in the scenario of TICL. It facilitates a balance between the acquisition of new knowledge and the retention of prior learning.

5 Experimental Results

In this section, we first introduce the experimental setup and then compare the proposed DAP with different existing CL methods applied on TICL benchmarks. Finally, we specifically evaluate the effectiveness of DAP and conduct ablation studies to key elements.

5.1 Implementation Details

Datasets. Given that long-tail distributions are the most prevalent form of imbalance in the real-world, we adopt the long-tailed setting to construct imbalances. The long-tailed distribution typically follows an exponential decay in sample size across classes [Cao *et al.*, 2019]. This decay is parameterized by ρ which is the ratio between the most and least frequent

	TICL-Cifar100						TICL-ImageNet-R						
	Descending		Ascending		Shuffled		Descending		Ascending		Shuffled		
Method	A_N (\uparrow)	A_L (\uparrow)	A_N (\uparrow)	A_L (\uparrow)	A_N (\uparrow)	A_L (\uparrow)	A_N (\uparrow)	A_L (\uparrow)	A_N (\uparrow)	A_L (\uparrow)	A_N (\uparrow)	A_L (\uparrow)	Buffer
BiC	27.92	38.02	36.08	37.90	27.11	33.63	21.91	18.9	13.52	16.01	16.36	16.32	20/cls
PODNET	26.48	23.82	32.31	28.24	28.49	26.21	22.32	18.90	16.61	16.20	17.11	18.70	20/cls
EEIL++	31.49	38.24	35.93	37.85	31.63	39.31	18.50	17.81	15.76	15.60	15.79	16.30	20/cls
LUCIR++	27.74	21.26	42.39	28.92	35.62	25.94	18.36	21.29	8.05	8.27	15.62	15.62	20/cls
Pre+FT	65.83	22.02	19.52	25.58	43.30	33.56	40.60	7.68	18.22	21.15	21.37	22.62	0
Pre+iCaRL	53.00	28.73	41.70	26.88	48.62	31.02	48.41	29.55	24.40	29.17	40.21	23.02	20/cls
CODA-P	81.91	58.98	54.54	41.84	60.90	42.56	52.39	35.21	28.21	32.62	40.02	34.78	0
L2-P	66.51	50.26	53.50	48.73	51.43	49.43	50.05	31.72	27.24	29.42	30.19	26.21	0
Dual-P	70.51	51.79	54.50	45.72	49.49	48.82	51.47	31.12	25.03	25.42	34.68	27.38	0
Ours	79.09	61.49	56.30	55.47	61.43	56.12	58.47	40.25	31.42	36.47	43.22	36.38	0

Table 1: Comparison Results (%) on TICL-Cifar100 and TICL-ImageNet-R. ‘Pre’ refers to pretraining and ‘P’ stands for prompt. A_N gives the accuracy averaged over tasks, A_L gives the last accuracy.

classes. $\rho = 1$ is the conventional CIL case and ρ in (0,1) indicates different degrees of long-tailed distribution.

We follow the experimental datasets used in previous works [Wang *et al.*, 2022d; Khan *et al.*, 2023], first conducting a long-tail division of the datasets, and then dividing them into 10 disjoint tasks. Specifically, the CIFAR-100 dataset [Krizhevsky *et al.*, 2009] includes 100 classes of natural images, with 500 training samples for the head classes, and subsequently decreasing for the remaining classes according to the long-tail division method. To ensure balance within each task, we select an equal number of samples from each class within a task, based on the maximum class quantity present in that task. The ImageNet-R dataset [Wang *et al.*, 2022d] contains 200 classes of images, divided in a similar manner for long-tail calculation. Aligned with this long-tailed distribution approach, we examine three cases in our study: Descending TICL, where learners first encounter data-rich tasks followed by data-scarce ones; Ascending TICL, featuring data-scarce tasks preceding data-rich ones; Shuffled TICL, where tasks arrive in a random sequence without a prescribed order of data volume.

Comparison Methods. In our experiments, we compare the proposed DAP with two groups of rehearsal-based methods. The first group includes classical approaches like PODNET [Douillard *et al.*, 2020] and BiC [Wu *et al.*, 2019]. The second group comprises methods designed for long-tailed distributions such as EEIL++ [Liu *et al.*, 2022] and LUCIR++ [Liu *et al.*, 2022]. Additionally, we include pretrained methods in our comparison, FineTune [Khan *et al.*, 2023], iCaRL [Rebuffi *et al.*, 2017], both these methods start from the same ImageNet pre-trained ViT-Base [Dosovitskiy *et al.*, 2020] model to ensure a fair comparison. Lastly, we compare DAP against the current state-of-the-art (SOTA) prompt-based methods, including L2P [Wang *et al.*, 2022e], Dual-Prompt [Wang *et al.*, 2022d], and CODA-Prompt [Smith *et al.*, 2023].

Evaluation Protocol. For the test set, we followed the setting of long-tailed learning research [Cao *et al.*, 2019] where the training set is imbalanced while the test set is balanced. Therefore, our test set is consistent with the balanced CL [Wang *et al.*, 2022c; Wang *et al.*, 2022d]. In this manner, the testing accuracy can be easily averaged over all classes to reflect the performance of each class with equal weight. For evaluation,

we report their average values with standard errors using two widely used CL metrics: average accuracy ($A_N \uparrow$) [Lopez-Paz *et al.*, 2017] of the final average accuracy by the model, last accuracy ($A_L \uparrow$) [Zhang *et al.*, 2023a] of the last accuracy at the end of the learning process.

Implementation. Following the settings of L2P [Wang *et al.*, 2022e], We train DAP using Adam with β_1, β_2 of 0.9, a learning rate of 0.01, and a batch size of 64. We resize the input images to a 224×224 resolution and normalize them between 0 and 1. To ensure models converge, we train TICL-CIFAR-100 for 5 epochs per task, TICL-ImageNet-R for 50 epochs each task.

5.2 Comparison to the State-of-The-Art

We compare various rehearsal-based and prompt-based methods for TICL-CIFAR-100 and TICL-ImageNet-R in Table 1. We observe that DAP consistently outperforms all rehearsal-based methods by a considerable margin, with a substantial improvement ranging from 10% to 30%, establishing a new state-of-the-art in all cases. It also demonstrates a significant advantage over other prompt-based methods, showing an increase of 4.5% to 15%.

5.3 Effectiveness of the General Prompt

To verify that adopting a single general prompt \mathbf{p}_g is able to accumulate knowledge across the data stream, we adopt a linear probing experiment [He *et al.*, 2020] to evaluate the performance of the representation layer. Specifically, following each incremental learning task, we freeze the representation layer and introduce an additional classification layer known as a linear probe, which is trained on all classes of the benchmark dataset.

We conducted a detailed analysis of the descending, ascending, and shuffled cases. As illustrated in Figure 4, in the descending case, we observe a rapid initial improvement, indicating that the model indeed learns significant global knowledge when presented with abundant data initially. However, as the tasks progress and data availability decreases, we notice a stagnation in learning, suggesting that under typical conditions, the model struggles to acquire new knowledge with limited data. In contrast, DAP continues to improve even with less data, overcoming this learning stagnation. In the ascending case, despite continuous learning, the final performance




Figure 4: Performance comparison between DAP and DualPrompt with linear probe.

	Desc.	Asc.	Shuff.
(1) w/ only task-specific prompt	64.75	52.56	52.23
(2) w/ only general prompt	66.79	50.49	47.50
(3) DAP	79.09	56.30	61.43

Table 2: Ablation on general prompt and task-specific prompt.


is lower than in the descending case, implying some degree of knowledge forgetting. Yet, DAP still shows an upward trend at the end. A similar pattern is observed in the shuffled case. This demonstrates DAP successfully accumulated knowledge, effectively navigating the numerical disparities in TICL environments. Moreover, DAP consistently outperforms DualPrompt [Wang *et al.*, 2022d] through all the tasks.

5.4 Ablation Study

In this section, we delve into an in-depth ablation study to validate the effectiveness and contributions of different components in our model.

Dynamic Factor λ . λ is essential for calibrating the balance between stability and plasticity. Therefore, to assess the dynamic regularization’s effectiveness, we compare the results in the shuffled case to the use of the fixed values of λ from 0 to 1. As illustrated in Figure 5, the dynamically adjusted λ consistently outperforms any fixed value of λ . This supports our premise that a dynamic λ is more flexible to adapt to the evolving requirements of learning new information while retaining previously acquired knowledge. Generalization on all classes is ensured irrespective of the data abundance in each task during learning.

Task-specific Prompt and General Prompt. We further examine the roles of prompts within the DAP framework. Since DAP optimizes the task-specific prompt for each task in the first in-task phase and continually updates the general prompt in the second in-task phase, we can investigate the performance of using either prompt exclusively. This purpose of the study aims to discern whether each prompt can sustain the model’s efficacy across various learning scenarios on its own. We compare DAP with two exclusively designed methods: (1) To solely use the task-specific prompt. The task-specific prompts for all tasks are stored during inference, and it is assumed that the prompt is properly selected for each test sample. (2) To solely use the general prompt. The general prompt is continually optimized without dynamic stability-plasticity


 Figure 5: Ablation on dynamic factor λ .

	$A_N(\uparrow)$	$F(\downarrow)$	$P(\uparrow)$
(1) w/ only general prompt	47.50	18.28	63.29
(2) w/ only boosting anchor	58.15	15.45	70.38
(3) w/ only stabilizing anchor	59.47	6.19	64.86
(4) DAP	61.43	8.04	68.04

 Table 3: Ablation on boosting anchor and stabilizing anchor in shuffled case A_N gives the accuracy averaged over tasks and F gives the average forgetting. P gives the average plasticity.

regularization. As shown in Table 2, employing either the task-specific prompt or the general prompt solely yields results that are much inferior to DAP. The task-specific prompt cannot harness the information across tasks while updating the general prompt without regularization cannot well balance the stability and plasticity. The strength of DAP’s design of updating the general prompt with regularization is verified in addressing the demands of TICL.

Ablation on Anchors. To ablate the boosting anchor and stabilizing anchor, we focused exclusively on the shuffled case of TICL, because this case uniquely demands the model to effectively balance both stability and plasticity. Therefore, we conducted experiments in this setting using each anchor type in isolation without the dynamic stability-plasticity regularization, aiming to assess the capability of each anchor. As seen from Table 3, employing the boosting anchor or the stabilizing anchor in isolation only enhances the model’s plasticity or stability but at a significant cost to each other. The best performance is attained by combining both boosting and stabilizing anchors with the dynamic factor.

6 Conclusion

In this paper, we formally define task-imbalanced continual learning and systematically study its three cases. We discovered that imbalanced tasks significantly deteriorate the performance of prompt-based CL methods because they raise a new challenge to consider the dilemma of stability and plasticity. To counteract this, we introduced dynamically anchored prompting. DAP addresses the challenge by separating stability and plasticity with two prompts, one for stabilization and the other for plasticity, serving as anchors to guide the learning process of a general prompt. DAP improves the performance of baseline prompt-based TICL methods to set a new state-of-the-art.

Acknowledgements

This study was supported in part by the National Natural Science Foundation of China under Grants 62376233, U21A20514, and 62306181; in part by the FuXiaQuan National Independent Innovation Demonstration Zone Collaborative Innovation Platform under Grant 3502ZCQXT2022008; in part by the Natural Science Foundation of Guangdong Province under Grant 2024A1515010163; in part by the China Fundamental Research Funds for the Central Universities under Grant 20720230038; and in part by Xiaomi Young Talents Program.

References

- [Aljundi *et al.*, 2018] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *ECCV*, pages 139–154, 2018.
- [Cao *et al.*, 2019] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. *NeurIPS*, 32, 2019.
- [Castro *et al.*, 2018] Francisco M Castro, Manuel J Marin-Jimenez, Nicolas Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *ECCV*, 2018.
- [Cha *et al.*, 2021] Hyuntak Cha, Jaeho Lee, and Jinwoo Shin. Co2l: Contrastive continual learning. In *CVPR*, pages 9516–9525, 2021.
- [Chaudhry *et al.*, 2021] Arslan Chaudhry, Albert Gordo, Puneet Dokania, Philip Torr, and David Lopez-Paz. Using hindsight to anchor past knowledge in continual learning. In *AAAI*, volume 35, pages 6993–7001, 2021.
- [Chen *et al.*, 2023] Xiaohua Chen, Yucan Zhou, Dayan Wu, Chule Yang, Bo Li, Qinghua Hu, and Weiping Wang. Area: adaptive reweighting via effective area for long-tailed classification. In *ICCV*, pages 19277–19287, 2023.
- [De Lange *et al.*, 2021] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *TPAMI*, 44(7):3366–3385, 2021.
- [Dosovitskiy *et al.*, 2020] 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.
- [Douillard *et al.*, 2020] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *ECCV*, pages 86–102. Springer, 2020.
- [He *et al.*, 2020] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020.
- [Huang *et al.*, 2019] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Deep imbalanced learning for face recognition and attribute prediction. *TPAMI*, 42(11):2781–2794, 2019.
- [Kang *et al.*, 2022] Haeyong Kang, Rusty John Lloyd Mina, Sultan Rizky Hikmawan Madjid, Jaehong Yoon, Mark Hasegawa-Johnson, Sung Ju Hwang, and Chang D Yoo. Forget-free continual learning with winning subnetworks. In *ICML*, pages 10734–10750. PMLR, 2022.
- [Kang *et al.*, 2023] Zhiqi Kang, Enrico Fini, Moin Nabi, Elisa Ricci, and Karteek Alahari. A soft nearest-neighbor framework for continual semi-supervised learning. In *ICCV*, pages 11868–11877, 2023.
- [Kemker and Kanan, 2017] Ronald Kemker and Christopher Kanan. Fearnnet: Brain-inspired model for incremental learning. *arXiv preprint arXiv:1711.10563*, 2017.
- [Kemker *et al.*, 2018] Ronald Kemker, Marc McClure, Angelina Abitino, Tyler Hayes, and Christopher Kanan. Measuring catastrophic forgetting in neural networks. In *AAAI*, volume 32, 2018.
- [Khan *et al.*, 2023] Muhammad Gul Zain Ali Khan, Muhammad Ferjad Naeem, Luc Van Gool, Didier Stricker, Federico Tombari, and Muhammad Zeshan Afzal. Introducing language guidance in prompt-based continual learning. In *ICCV*, pages 11463–11473, 2023.
- [Kim *et al.*, 2020] Chris Dongjoo Kim, Jinseo Jeong, and Gunhee Kim. Imbalanced continual learning with partitioning reservoir sampling. In *ECCV*, pages 411–428. Springer, 2020.
- [Kirkpatrick *et al.*, 2017] 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.
- [Krizhevsky *et al.*, 2009] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [Liu *et al.*, 2022] Xialei Liu, Yu-Song Hu, Xu-Sheng Cao, Andrew D Bagdanov, Ke Li, and Ming-Ming Cheng. Long-tailed class incremental learning. In *European Conference on Computer Vision*, pages 495–512. Springer, 2022.
- [Lopez-Paz *et al.*, 2017] David Lopez-Paz, Robert Nishihara, Soumith Chintala, Bernhard Scholkopf, and Léon Bottou. Discovering causal signals in images. In *CVPR*, pages 6979–6987, 2017.
- [Peng *et al.*, 2023] Hanyu Peng, Weiguo Pian, Mingming Sun, and Ping Li. Dynamic re-weighting for long-tailed semi-supervised learning. In *WACV*, pages 6464–6474, 2023.
- [Rebuffi *et al.*, 2017] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *CVPR*, pages 2001–2010, 2017.

- [Smith *et al.*, 2023] James Seale Smith, Leonid Karlinsky, Vyshnavi Gutta, Paola Cascante-Bonilla, Donghyun Kim, Assaf Arbelle, Rameswar Panda, Rogerio Feris, and Zsolt Kira. Coda-prompt: Continual decomposed attention-based prompting for rehearsal-free continual learning. In *CVPR*, pages 11909–11919, 2023.
- [Sun *et al.*, 2022] Qing Sun, Fan Lyu, Fanhua Shang, Wei Feng, and Liang Wan. Exploring example influence in continual learning. *NIPS*, 35:27075–27086, 2022.
- [Vaishnavh *et al.*, 2018] N Vaishnavh, C Raffel, and IJ Goodfellow. Theoretical insights into memorization in gans. In *NeurIPS*, 2018.
- [Wang *et al.*, 2022a] Fu-Yun Wang, Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Foster: Feature boosting and compression for class-incremental learning. In *ECCV*, pages 398–414. Springer, 2022.
- [Wang *et al.*, 2022b] Liyuan Wang, Xingxing Zhang, Qian Li, Jun Zhu, and Yi Zhong. Coscl: Cooperation of small continual learners is stronger than a big one. In *ECCV*, pages 254–271. Springer, 2022.
- [Wang *et al.*, 2022c] Zhen Wang, Liu Liu, Yiqun Duan, Yajing Kong, and Dacheng Tao. Continual learning with life-long vision transformer. In *CVPR*, pages 171–181, 2022.
- [Wang *et al.*, 2022d] Zifeng Wang, Zizhao Zhang, Sayna Ebrahimi, Ruoxi Sun, Han Zhang, Chen-Yu Lee, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, et al. Dualprompt: Complementary prompting for rehearsal-free continual learning. In *ECCV*, pages 631–648, 2022.
- [Wang *et al.*, 2022e] 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 *CVPR*, pages 139–149, 2022.
- [Wang *et al.*, 2023] 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.
- [Wu *et al.*, 2019] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *CVPR*, pages 374–382, 2019.
- [Zhang *et al.*, 2020] Junting Zhang, Jie Zhang, Shalini Ghosh, Dawei Li, Serafettin Tasci, Larry Heck, Heming Zhang, and C-C Jay Kuo. Class-incremental learning via deep model consolidation. In *CVPR*, pages 1131–1140, 2020.
- [Zhang *et al.*, 2023a] Gengwei Zhang, Liyuan Wang, Guoliang Kang, Ling Chen, and Yunchao Wei. Slca: Slow learner with classifier alignment for continual learning on a pre-trained model. *arXiv preprint arXiv:2303.05118*, 2023.
- [Zhang *et al.*, 2023b] Yifan Zhang, Bingyi Kang, Bryan Hooi, Shuicheng Yan, and Jiashi Feng. Deep long-tailed learning: A survey. *TPAMI*, 2023.