

# Pareto Continual Learning: Preference-Conditioned Learning and Adaption for Dynamic Stability-Plasticity Trade-off

Song Lai<sup>1,2,5</sup>, Zhe Zhao<sup>1,6</sup>, Fei Zhu<sup>2</sup>, Xi Lin<sup>1</sup>, Qingfu Zhang<sup>1,5\*</sup>, Gaofeng Meng<sup>2,3,4\*</sup>

<sup>1</sup> Department of Computer Science, City University of Hong Kong

<sup>2</sup> Centre for Artificial Intelligence and Robotics, HK Institute of Science & Innovation, Chinese Academy of Sciences

<sup>3</sup> State Key Laboratory of Multimodal Artificial Intelligence Systems, Institute of Automation, Chinese Academy of Sciences

<sup>4</sup> School of Artificial Intelligence, University of Chinese Academy of Sciences

<sup>5</sup> City University of Hong Kong Shenzhen Research Institute

<sup>6</sup> University of Science and Technology of China

{songlai2-c, xi.lin}@my.cityu.edu.hk, zz4543@mail.ustc.edu.cn

qingfu.zhang@cityu.edu.hk, {fei.zhu, gaofeng.meng}@cair-cas.org.hk

## Abstract

Continual learning aims to learn multiple tasks sequentially. A key challenge in continual learning is balancing between two objectives: retaining knowledge from old tasks (stability) and adapting to new tasks (plasticity). Experience replay methods, which store and replay past data alongside new data, have become a widely adopted approach to mitigate catastrophic forgetting. However, these methods neglect the dynamic nature of the stability-plasticity trade-off and aim to find a fixed and unchanging balance, resulting in suboptimal adaptation during training and inference. In this paper, we propose Pareto Continual Learning (ParetoCL), a novel framework that reformulates the stability-plasticity trade-off in continual learning as a multi-objective optimization (MOO) problem. ParetoCL introduces a preference-conditioned model to efficiently learn a set of Pareto optimal solutions representing different trade-offs and enables dynamic adaptation during inference. From a generalization perspective, ParetoCL can be seen as an objective augmentation approach that learns from different objective combinations of stability and plasticity. Extensive experiments across multiple datasets and settings demonstrate that ParetoCL outperforms state-of-the-art methods and adapts to diverse continual learning scenarios.

**Code** — <https://github.com/laisong-22004009/ParetoCL>

## Introduction

Continual learning (CL) has emerged as a crucial paradigm, enabling artificial intelligence systems to learn and adapt in dynamic environments. Unlike traditional machine learning (LeCun, Bengio, and Hinton 2015; He et al. 2016), where models are trained on a fixed dataset, continual learning enables models to learn from a stream of tasks or data distributions sequentially. This ability is crucial for real-world applications with incremental data and evolving environments. However, a fundamental challenge in continual learning is the stability-plasticity dilemma (Wang et al.

2024), which refers to the trade-off between a model’s ability to retain learned knowledge (stability) and its ability to acquire new knowledge (plasticity).

Numerous CL methods have been proposed to tackle the stability-plasticity dilemma. Among them, experience replay (ER) methods (Rolnick et al. 2018; Buzzega et al. 2020; Arani, Sarfraz, and Zonooz 2021) stand as a straightforward yet effective research line. These methods store a subset of past data in a memory buffer and replay it alongside new data during training. However, existing ER methods aim to find a fixed trade-off between stability and plasticity, ignoring the dynamic nature of the stability-plasticity trade-off. This results in a single model that cannot dynamically adapt to evolving environments and changing task requirements during training. Second, the optimal stability-plasticity trade-off may vary not only during training but also during inference. For example, in the study by (Cai, Sener, and Koltun 2021), it was observed that the stability-plasticity trade-off evolves over time as the data distribution shifts. This highlights the need for dynamic adaptation of the trade-off based on the current learning context.

*Can we make a proper trade-off between stability and plasticity and integrate these two objectives to enhance the model’s generalization ability?* To address this question, we draw inspiration from multi-objective optimization (MOO) (Miettinen 1999). MOO aims to find a set of Pareto optimal solutions that represent different optimal trade-offs among conflicting objectives. Building upon this concept, we envision finding a set of Pareto optimal models that embody different stability-plasticity trade-offs rather than a single model with a fixed trade-off. However, training separate models for each trade-off is computationally expensive, and a finite number of Pareto solutions cannot fully capture the relationship between stability and plasticity. To overcome these challenges, we propose Pareto Continual Learning (ParetoCL), a novel framework that introduces an efficient preference-conditioned model to *learn* the trade-off between the two objectives and perform dynamic adaptation. Specifically, ParetoCL trains the preference-conditioned model by sampling preference vectors from a prior distribution and

\*Corresponding authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

minimizes a preference-based aggregation of the losses on the replay buffer (for stability) and the new data (for plasticity). During inference, the model dynamically selects the optimal trade-off for each sample based on the preference vector that yields the most confident prediction. Experiments on multiple benchmarks demonstrate that ParetoCL significantly outperforms state-of-the-art continual learning methods. By inputting different preferences, ParetoCL can obtain a set of well-distributed Pareto optimal solutions with varying trade-offs between stability and plasticity.

From a generalization perspective, our work can be seen as an extension of data augmentation (Buzzega et al. 2020) and class augmentation (Zhu et al. 2021) in the continual learning domain—*objective augmentation*. The model is trained using different combinations of preferences between stability and plasticity, equivalent to learning from different combinations of old and new task distribution. Since we cannot know which task distribution combination leads to better generalization in advance, and training models for each possible distribution is prohibitively expensive, we transform the problem of combining task distribution into that of combining preference-conditioned model hypotheses. In this way, we bypass suboptimal task combinations and promote positive transfer between past and current knowledge.

The main contributions of this work are threefold:

- *New perspective*: To the best of our knowledge, this work is the first attempt to explore the dynamic nature of the stability-plasticity trade-off in continual learning and reformulate it as a MOO problem.
- *Effective algorithm*: We propose an efficient preference-conditioned CL model that learns a set of Pareto solutions of the stability-plasticity trade-off during training and enables dynamic adaptation during inference.
- *Compelling empirical results*: We conduct extensive experiments on multiple benchmarks and settings, demonstrating that ParetoCL significantly outperforms state-of-the-art continual learning methods.

## Related Work

### Continual Learning

Continual learning methods can be broadly categorized into regularization-based, expansion-based, and replay-based approaches. Regularization-based methods (Huszár 2017; Zenke, Poole, and Ganguli 2017; Yang et al. 2019, 2021) introduce additional regularization terms to the loss function to prevent catastrophic forgetting. These regularizers penalize changes to parameters that are important for previous tasks. Expansion-based methods (Li et al. 2019; Serra et al. 2018; Yoon et al. 2017; Mallya, Davis, and Lazebnik 2018; Hung et al. 2019) allocate new resources for new tasks while preserving the learned knowledge in the existing network. Replay-based methods (Rolnick et al. 2018; Buzzega et al. 2020; Aljundi et al. 2019) store and replay previous samples to maintain learned information, which has been widely adopted due to its effectiveness and simplicity. Methods like iCaRL (Rebuffi et al. 2017) combine replay with knowledge distillation to retain previous knowledge. MER (Riemer et al. 2018) and CLSER (Arani, Sarfraz,

and Zonoos 2021) introduce additional model parameters to consolidate task-specific knowledge. ER-ACE (Caccia et al. 2021) modifies the loss function to alleviate representation drift. However, these approaches seek a fixed trade-off between stability and plasticity, limiting their ability to adapt to dynamic learning scenarios.

### Multi-Objective Optimization

Multi-objective optimization (MOO) aims to find a set of Pareto optimal solutions representing different trade-offs among multiple conflicting objectives (Ehrgott 2005). MOO has found numerous applications in machine learning, including reinforcement learning (Liu, Xu, and Hu 2014), neural architecture search (Lu et al. 2019), and fairness (Wu et al. 2022). Gradient-based MOO algorithms, such as MDGA (Sener and Koltun 2018) and ParetoMTL (Lin et al. 2019), have been proposed to train large-scale neural networks for multi-task learning. These methods scalarize the MOO problem into a single-objective problem using a weighted sum of the objectives. However, they require a priori preference specification and can only find a single Pareto optimal solution. Recent works (Lin, Yang, and Zhang 2022; Navon et al. 2020) propose to approximate the entire Pareto front by training a hypernetwork conditioned on the preference vector. Our work shares the goal of approximating the Pareto front of multiple objectives with these MOO methods. However, we focus on the specific problem of continual learning and propose a novel preference-conditioned model to adapt the stability-plasticity trade-off dynamically during both training and inference.

## Preliminary

### Experience Replay

Continual learning requires the model to learn from a non-stationary data stream. At each time step  $t$ , the model receives a batch of labeled data  $\mathcal{D}_t = \{(x_i^t, y_i^t)\}_{i=1}^B$  drawn from an underlying distribution  $\mathcal{P}_t$  that gradually shifts over time. The goal is to learn a model  $f_\theta$  parameterized by  $\theta$  that performs well on all seen distributions  $\{\mathcal{P}_1, \dots, \mathcal{P}_t\}$  up to the current time step.

Let  $\mathcal{L}(f_\theta; \mathcal{D})$  denote the loss of the model  $f_\theta$  on a dataset  $\mathcal{D}$ . To tackle catastrophic forgetting, experience replay methods maintain a memory buffer  $\mathcal{M}$  that stores a subset of past data. At each time step  $t$ , the model is updated using the following loss:

$$\mathcal{L}_t = \lambda \mathcal{L}_{new}(f_\theta; \mathcal{D}_t) + (1 - \lambda) \mathcal{L}_{replay}(f_\theta; \mathcal{M}_t), \quad (1)$$

where  $\lambda \in [0, 1]$  is a hyperparameter that controls the trade-off between the losses on the new data  $\mathcal{D}_t$  and the memory buffer  $\mathcal{M}_t$ . The first term promotes plasticity to adapt to the new data, while the second term promotes stability from past data. The memory buffer  $\mathcal{M}_t$  is updated by randomly selecting samples to replace with the new data.

### Multi-Objective Optimization

Multi-objective optimization (MOO) problem is defined as:

$$\min_{\theta} \mathcal{F}(\theta) = (f_1(\theta), \dots, f_m(\theta))^T \quad \text{s.t.} \quad \theta \in \Theta, \quad (2)$$

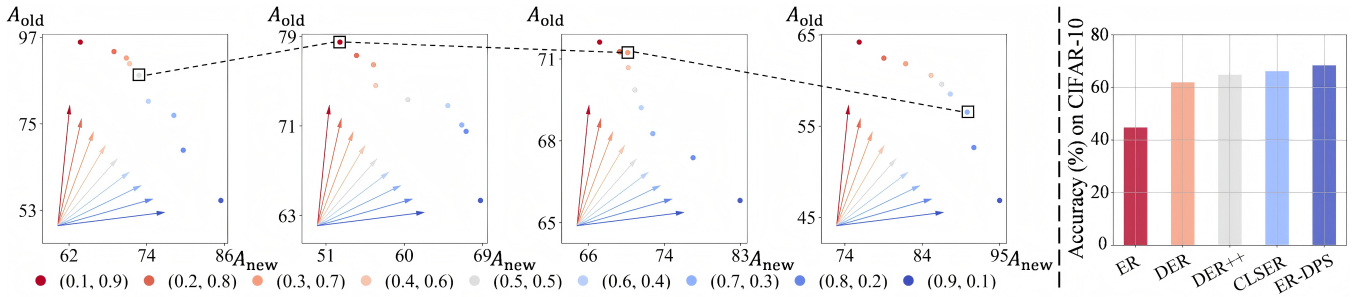


Figure 1: **Motivation.** (left) illustrates the multi-stage training process of a ResNet-18 model on 5 tasks of Seq-CIFAR10 while (right) compares ER-DPS with several SOTA replay-based methods. In (left), each stage builds upon the model trained on the previous task, incorporating different stability-plasticity trade-offs. Each point in (left) corresponds to the accuracy obtained when training the model under a specific preference vector, demonstrating how different trade-offs influence stability and plasticity. The metrics in the graph,  $A_{old}$  and  $A_{new}$ , represent the average accuracy on previous tasks and the accuracy on the current task, respectively, corresponding to stability and plasticity. The dashed line indicates the model selected at the current stage corresponding to the chosen preference. Detailed data is provided in supplementary material.

where  $\theta$  is the decision variable,  $\Theta$  is the feasible set, and  $f_1(\theta), \dots, f_m(\theta)$  are  $m$  objective functions. Due to the conflicting nature of objectives, no single solution  $\theta^*$  can optimize all objectives simultaneously. This leads to the definitions of Pareto dominance, Pareto optimality, and Pareto Set/Front for MOO problems:

**Definition 1 (Pareto Dominance)** A solution  $\theta_1$  dominates another solution  $\theta_2$ , denoted  $\theta_1 \prec \theta_2$ , if  $f_i(\theta_1) \leq f_i(\theta_2)$  for all  $i \in \{1, \dots, m\}$  and there exists  $j \in \{1, \dots, m\}$  such that  $f_j(\theta_1) < f_j(\theta_2)$ .

**Definition 2 (Pareto Optimality)** A solution  $\theta$  is Pareto optimal if there does not exist another solution  $\theta' \in \Theta$  that dominates  $\theta$ , i.e.,  $\nexists \theta' \in \Theta$  such that  $\theta' \prec \theta$ .

**Definition 3 (Pareto Set and Pareto Front)** The Pareto Set,  $\Theta^*$ , is the set of all Pareto optimal solutions, defined as  $\Theta^* = \{\theta \in \Theta \mid \nexists \theta' \in \Theta \text{ such that } \theta' \prec \theta\}$ . The Pareto Front,  $\mathcal{F}(\Theta^*)$ , is the image of the Pareto Set in the objective space, defined as  $\mathcal{F}(\Theta^*) = \{F(\theta) \mid \theta \in \Theta^*\}$ .

## Pareto Continual Learning

### Motivation

To illustrate the motivation behind our algorithm, we conduct an experiment on class incremental learning with Seq-CIFAR10 (Buzzega et al. 2020). The goal is to demonstrate the effectiveness of training a diverse set of solutions with different stability-plasticity trade-offs and dynamically selecting the appropriate solution during inference. Figure 1(left) illustrates the multi-stage training process of a ResNet-18 model on 5 tasks of Seq-CIFAR10. Each stage builds upon the model trained on the previous task and incorporates different stability-plasticity trade-offs by varying the preference vectors in the weighted aggregation of the replay loss and the cross-entropy loss on the current task. The models are evaluated on the test sets of both the current and previous tasks, and the model with the highest average accuracy is selected for the next stage. We call this method ER-DPS (Experience Replay with Dynamic Preference Selection), which can be considered as an **oracle** ER model

with optimal stability-plasticity trade-off in task level. Figure 1(right) shows the final mean accuracy of the models across all tasks, demonstrating the superiority of dynamic trade-off adaption over fixed weighted aggregation. The experimental results lead to two key observations:

- The optimal stability-plasticity trade-off varies and the model needs to prioritize plasticity in certain stages while emphasizing stability in others;
- Selecting the appropriate trade-off carefully improves the performance compared to using a fixed trade-off.

However, directly applying the above approach to continual learning faces two challenges:

- The optimal trade-off is an unknown priori.
- Training separate models for each trade-off is computationally expensive and memory-inefficient.

These challenges motivate the design of our Pareto Continual Learning framework, which learns a single model conditioned on preference vectors to approximate the Pareto front of stability-plasticity trade-offs.

### ER as Multi-Objective Optimization

Drawing insights from the above observations, we propose to formulate experience replay as a MOO problem. Specifically, let  $\mathcal{L}_{new}$  and  $\mathcal{L}_{replay}$  denote the losses on the new data and the replay buffer, respectively. We define two objectives corresponding to plasticity and stability:

$$\min_{\theta} F(\theta) = (\mathcal{L}_{replay}(f_{\theta}), \mathcal{L}_{new}(f_{\theta}))^T. \quad (3)$$

Our goal is to find a Pareto set that represent different trade-offs between these two objectives, rather than a single solution with fixed trade-off. This approach enables the model to adapt dynamically to varying degrees of knowledge retention and acquisition across different tasks and time periods.

### Preference-Conditioned Learning

To learn a set of models with different trade-offs, we introduce a preference-conditioned model  $f_{\theta}(x; \alpha)$  parameterized by  $\theta$  that takes as input a sample  $x$  and a preference




Figure 2: **An overview of our proposed ParetoCL framework.** Our model consists of two parts: a shared encoder and a preference-based classifier (with parameters generated by a hypernetwork). The model takes inputs from the data stream and memory buffer, which are then transformed into embeddings by the encoder. These embeddings are used by the preference-based classifier multiple times with different preferences, each of which is sampled from a prior distribution. The sampled preference is then fed into the hypernetwork to obtain the parameters of the preference-based classifier. For the embeddings from both the data stream and memory buffer, the classifier computes two losses:  $\mathcal{L}_{new}$  and  $\mathcal{L}_{replay}$ . The final optimization objective of the model is the expectation of the preference-based aggregation of these two losses. Through this approach, ParetoCL learns a mapping from the preference space to the objective space of different plasticity-stability trade-offs.

---

#### Algorithm 1: ParetoCL Training

---

**Require:** Prior distribution of preferences  $p(\alpha)$ , model parameter  $\theta$

- 1: Initialize  $\theta$
  - 2: **for** each iteration **do**
  - 3:   Sample data  $\mathcal{D}_t$  from the data stream
  - 4:   Sample data  $\mathcal{M}_t$  from the memory buffer
  - 5:   Feature embedding:  $h_{new} = f(\theta, \mathcal{D}_t)$
  - 6:   Feature embedding:  $h_{replay} = f(\theta, \mathcal{M}_t)$
  - 7:   **for**  $k = 1$  to  $K$  **do**
  - 8:     Sample preference vector:  $\alpha_k \sim p(\alpha)$
  - 9:     Compute output based on Eq.(5)
  - 10:   **end for**
  - 11:   Compute loss based on Eq.(4)
  - 12:   Update  $\theta$
  - 13: **end for**
- 

---

#### Algorithm 2: ParetoCL Inference

---

**Require:** Model parameter  $\theta$

- 1: Input sample:  $x$
  - 2: Feature embedding:  $h = f(\theta, x)$
  - 3: **for**  $k = 1$  to  $K$  **do**
  - 4:   Sample preference vector  $\alpha^k \sim p(\alpha)$
  - 5:   Compute each output based on Eq.(5)
  - 6: **end for**
  - 7: Compute final output based on Eq.(6)
- 

vector  $\alpha \in \mathbb{R}_+^2$  and outputs a prediction, as illustrated in Figure 2. The preference vector  $\alpha = (\alpha_1, \alpha_2)$  determines the trade-off between stability and plasticity, where a higher value of  $\alpha_1$  prioritizes stability and a higher value of  $\alpha_2$  prioritizes plasticity. The model is trained to minimize the fol-

lowing loss:

$$\mathcal{L}(\theta) = \mathbb{E}_{\alpha \sim p(\alpha)} [\alpha_1 \mathcal{L}_{replay}(\theta, \alpha) + \alpha_2 \mathcal{L}_{new}(\theta, \alpha)], \quad (4)$$

where  $p(\alpha)$  is a prior distribution over the preferences (e.g., a Dirichlet distribution).

To implement preference conditioning, we employ a hypernetwork  $\Psi$  that generates the parameters of the model's final linear layer based on the preference vector. Specifically, let  $h_\theta(x)$  denote the penultimate features of the model. The final output is computed as:

$$f_\theta(x; \alpha) = W(\alpha)h_\theta(x) + b(\alpha), \quad (5)$$

where  $W(\alpha) = \Psi_W(\alpha)$  and  $b(\alpha) = \Psi_b(\alpha)$  are the weight matrix and bias vector generated by the hypernetwork  $\Psi$  conditioned on the preference vector  $\alpha$ .

During training, we sample a new preference vector  $\alpha \sim p(\alpha)$  for each mini-batch and compute the losses  $\mathcal{L}_{replay}(\theta, \alpha)$  and  $\mathcal{L}_{new}(\theta, \alpha)$  using the preference-conditioned model  $f_\theta(x; \alpha)$ , which is outlined in Algorithm 1. The model parameters  $\theta$  and the hypernetwork parameters are jointly optimized to minimize the expected loss over the preference distribution, as shown in Eq.(4). By varying the preference vectors, the model learns to map different preferences to different points on the Pareto front of stability-plasticity trade-offs.

#### Dynamic Preference Adaptation

With the preference-conditioned model, we can adapt the stability-plasticity trade-off on the fly during inference. As shown in Algorithm 2, for each input sample  $x$ , we sample multiple preference vectors  $\{\alpha^1, \dots, \alpha^K\}$  and compute the model outputs  $\{f_\theta(x; \alpha^1), \dots, f_\theta(x; \alpha^K)\}$ . We then select the output with the highest confidence (i.e., lowest entropy) as the final prediction:

$$y^* = \arg \min_k H(f_\theta(x; \alpha^k)), \quad (6)$$

Method	Seq-CIFAR10				Seq-CIFAR100				Seq-TinyImageNet			
	Online		Offline		Online		Offline		Online		Offline	
	AAA	Acc	AAA	Acc	AAA	Acc	AAA	Acc	AAA	Acc	AAA	Acc
SGD	38.12	18.33	43.11	19.27	11.63	5.69	19.22	8.46	11.35	4.02	17.60	6.91
OnEWC (Huszár 2017)	38.75	17.89	43.25	21.33	12.14	6.23	17.84	6.64	11.92	4.73	14.64	4.92
LwF (Li and Hoiem 2016)	35.61	18.91	-	-	12.16	5.99	-	-	11.42	4.78	-	-
IS (Zenke, Poole, and Ganguli 2017)	38.89	20.44	40.52	18.78	13.21	7.33	16.56	5.25	10.82	4.23	11.42	2.83
La-MAML (Gupta, Yadav, and Paull 2020)	36.76	35.44	61.56	44.38	13.45	12.21	32.35	20.34	16.32	13.83	37.11	24.91
VR-MCL (Wu et al. 2024a)	69.57	58.43	77.95	68.14	30.46	22.31	42.36	30.09	24.41	20.55	41.13	28.02
ER (Rolnick et al. 2018)	52.74	33.14	75.24	64.13	18.70	15.12	36.73	22.19	19.14	14.21	37.87	25.95
A-GEM (Chaudhry et al. 2018)	39.25	18.54	43.63	20.62	14.66	6.24	19.89	7.94	11.66	5.26	19.23	8.40
GEM (Lopez-Paz and Ranzato 2017)	36.83	19.75	49.68	24.45	16.23	9.91	22.35	11.45	13.57	6.89	20.32	9.28
DER (Buzzega et al. 2020)	38.59	18.10	66.48	59.55	14.65	6.27	19.22	7.85	11.58	5.31	17.71	7.18
DER++ (Buzzega et al. 2020)	60.63	50.33	75.20	67.36	24.47	16.32	40.01	23.86	19.21	13.65	38.94	27.07
CLSER (Arani, Sarfraz, and Zonooz 2021)	61.88	49.03	76.25	67.86	27.83	19.45	41.28	27.06	23.47	19.72	38.64	26.48
ER-OBC (Chrysakis and Moens 2023)	65.77	54.64	76.42	65.93	26.11	17.31	42.26	<b>30.22</b>	24.11	19.86	39.41	27.67
OCM (Guo, Liu, and Zhao 2022)	66.22	53.89	-	-	26.03	15.88	-	-	17.55	8.33	-	-
<b>ParetoCL (Ours)</b>	<b>70.89</b>	<b>59.95</b>	<b>78.98</b>	<b>69.55</b>	<b>33.04</b>	<b>24.45</b>	<b>44.32</b>	28.79	<b>31.72</b>	<b>23.09</b>	<b>43.02</b>	<b>28.28</b>

Table 1: Performance of continual learning in Seq-CIFAR10 and longer task sequences Seq-CIFAR100, Seq-TinyImageNet with ResNet-18 under online and offline settings. ‘-’ indicates the implementation is unstable. All reported results are the average of 5 runs. Full table can be found in supplementary material.

where  $H(\cdot)$  denotes the entropy of the model’s output distribution. Intuitively, the output with the lowest entropy corresponds to the preference vector that yields the most confident prediction for the given sample. This dynamic preference adaptation allows the model to select the most suitable trade-off for each sample, leading to better performance than using a fixed trade-off.

## Comparison with Existing Work

While POCL (Wu et al. 2024b) also incorporates Pareto optimization into CL, ParetoCL differs in several aspects: First, ParetoCL explicitly models the stability-plasticity trade-off as a MOO problem, while POCL focuses on modeling relationships among past tasks to enhance stability; Second, instead of using gradient alignment like POCL, ParetoCL learns a set of Pareto solutions, which allows ParetoCL to handle the dynamic nature of the trade-off.

## Experiments

### Experimental Setup

**Datasets and Settings** To verify the effectiveness of the proposed ParetoCL, we conduct comprehensive experiments on commonly used datasets Seq-CIFAR10, as well as the longer task sequences Seq-CIFAR100 and Seq-TinyImageNet (Buzzega et al. 2020). Specifically, the Seq-CIFAR10 dataset comprises 5 tasks, with each task containing 2 classes. Seq-CIFAR100 consists of 10 tasks, each with 10 classes, while Seq-TinyImageNet includes 10 tasks, each encompassing 20 classes. In this paper, we focus on class incremental learning, where the model cannot get the task index during testing. Our evaluation includes metrics following previous works (Wu et al. 2024a). We choose the final average accuracy (Acc) across all tasks after sequential training on each task as the main metric for comparing approaches. Moreover, we use the averaged anytime accuracy

(AAA) to evaluate the model throughout the stream of tasks. Let  $AA_j$  denote the test average accuracy after training on  $T_j$ , then the evaluation metrics of AAA and Acc are:

$$AAA = \frac{1}{N} \sum_{j=1}^N AA_j, \quad Acc = AA_N. \quad (7)$$

**Baseline and Training Details** To effectively validate the efficacy of ParetoCL, we select three regularization-based methods (OnEWC (Huszár 2017), SI (Zenke, Poole, and Ganguli 2017), LwF (Li and Hoiem 2016)), two meta continual learning methods (La-MAML (Gupta, Yadav, and Paull 2020), VR-MCL (Wu et al. 2024a)), and six rehearsal-based methods (ER (Rolnick et al. 2018), GEM (Lopez-Paz and Ranzato 2017), A-GEM (Chaudhry et al. 2018), DER (Buzzega et al. 2020), DER++ (Buzzega et al. 2020), CLSER (Arani, Sarfraz, and Zonooz 2021), ER-OBC (Chrysakis and Moens 2023), OCM (Guo, Liu, and Zhao 2022)). For fair comparison among different continual learning methods, we follow the hyperparameter settings used in their respective original papers and code implementations. We further provide a lower bound using SGD without any countermeasure to forgetting. For fair comparison, we train all models using the Stochastic Gradient Descent (SGD) optimizer. Following DER (Buzzega et al. 2020), we utilize ResNet-18 as the backbone. The hypernetwork  $\Psi$  is an MLP with two hidden layers. The distribution for preferences during both training and inference is Dirichlet Distribution. The preferences sample number  $K$  is 5 for training and 20 for inference. The learning rate of ParetoCL is set to 0.05 and the size of the replay buffer is set to 32 for all experiments. Under the offline CL setting, we set the training epochs to 5 for all methods. Further details are provided in supplementary material.




Figure 3: **Effectiveness of Preference-Conditioned Learning and Dynamic Preference Adaptation.** (left) shows the Pareto front approximated by ParetoCL in each stage on Seq-CIFAR10 in the offline setting. The results demonstrate that ParetoCL can effectively explore various stability-plasticity trade-offs; (right) compares the average incremental accuracy of different methods on Seq-CIFAR100 in the online setting. Dynamic Preference Adaptation can effectively improve performance compared to ParetoCL-- and other baselines.

Method	$M = 600$		$M = 1000$		$M = 1400$	
	AAA	Acc	AAA	Acc	AAA	Acc
SGD	11.63 $\pm$ 0.31	5.69 $\pm$ 0.19	11.63 $\pm$ 0.31	5.69 $\pm$ 0.19	11.63 $\pm$ 0.31	5.69 $\pm$ 0.19
OnEWC (Huszár 2017)	12.14 $\pm$ 0.43	6.23 $\pm$ 0.27	12.14 $\pm$ 0.43	6.23 $\pm$ 0.27	12.14 $\pm$ 0.43	6.23 $\pm$ 0.27
LwF (Li and Hoiem 2016)	12.09 $\pm$ 0.66	6.32 $\pm$ 0.38	12.16 $\pm$ 0.49	5.99 $\pm$ 0.36	12.89 $\pm$ 0.21	6.35 $\pm$ 0.18
IS (Zenke, Poole, and Ganguli 2017)	13.21 $\pm$ 0.59	7.33 $\pm$ 0.66	13.21 $\pm$ 0.59	7.33 $\pm$ 0.66	13.21 $\pm$ 0.59	7.33 $\pm$ 0.66
La-MAML (Gupta, Yadav, and Paull 2020)	10.42 $\pm$ 0.97	9.32 $\pm$ 0.48	13.45 $\pm$ 0.29	12.21 $\pm$ 0.47	16.32 $\pm$ 0.41	15.52 $\pm$ 0.32
VR-MCL (Wu et al. 2024a)	25.31 $\pm$ 0.52	17.32 $\pm$ 0.48	30.46 $\pm$ 0.49	22.31 $\pm$ 0.45	32.22 $\pm$ 0.33	25.67 $\pm$ 0.40
ER (Rolnick et al. 2018)	19.76 $\pm$ 0.31	12.69 $\pm$ 0.44	18.70 $\pm$ 0.37	15.12 $\pm$ 0.76	21.58 $\pm$ 0.42	18.25 $\pm$ 0.62
A-GEM (Chaudhry et al. 2018)	14.60 $\pm$ 0.24	6.59 $\pm$ 0.17	14.66 $\pm$ 0.32	6.24 $\pm$ 0.23	14.40 $\pm$ 0.28	6.75 $\pm$ 0.27
GEM (Lopez-Paz and Ranzato 2017)	16.97 $\pm$ 0.28	9.91 $\pm$ 0.39	16.23 $\pm$ 0.58	9.91 $\pm$ 0.32	16.03 $\pm$ 0.46	9.72 $\pm$ 0.41
DER (Buzzega et al. 2020)	14.62 $\pm$ 0.34	6.75 $\pm$ 0.19	14.65 $\pm$ 0.47	6.27 $\pm$ 0.13	6.65 $\pm$ 0.32	14.27 $\pm$ 0.44
DER++ (Buzzega et al. 2020)	24.15 $\pm$ 0.47	12.78 $\pm$ 0.52	24.47 $\pm$ 0.55	16.32 $\pm$ 0.39	22.55 $\pm$ 0.55	18.52 $\pm$ 0.51
CLSER (Arani, Sarfraz, and Zonooz 2021)	27.75 $\pm$ 0.42	16.97 $\pm$ 0.31	27.83 $\pm$ 0.71	19.45 $\pm$ 0.35	26.92 $\pm$ 0.57	24.96 $\pm$ 0.35
ER-OBC (Chrysakis and Moens 2023)	24.02 $\pm$ 0.34	15.92 $\pm$ 0.39	26.11 $\pm$ 0.54	17.31 $\pm$ 0.88	28.91 $\pm$ 0.31	19.74 $\pm$ 0.83
OCM (Guo, Liu, and Zhao 2022)	18.96 $\pm$ 0.22	8.72 $\pm$ 0.17	26.03 $\pm$ 0.90	15.88 $\pm$ 0.95	28.23 $\pm$ 0.94	16.35 $\pm$ 0.81
<b>ParetoCL (Ours)</b>	<b>31.03<math>\pm</math>0.61</b>	<b>20.49<math>\pm</math>0.29</b>	<b>33.04<math>\pm</math>0.61</b>	<b>24.45<math>\pm</math>0.39</b>	<b>33.76<math>\pm</math>0.39</b>	<b>26.97<math>\pm</math>0.42</b>

Table 2: Ablation study on buffer size  $M$  for Seq-CIFAR100 under online setting. The backbone is ResNet-18.

## Improvements over Baselines

We conduct extensive baseline comparison on three datasets: Seq-CIFAR10, Seq-CIFAR100, and Seq-TinyImageNet with both online and offline settings and the results are summarized in Table 1. It is evident that our proposed ParetoCL significantly outperforms other replay-based methods (shown in the bottom row of the table), including methods with sophisticated replay strategies such as CLSER (Arani, Sarfraz, and Zonooz 2021). These results align with our analysis in Figure 1 (right), demonstrating the effectiveness of learning a set of Pareto optimal models and dynamic adaption for different stability-plasticity trade-offs. Moreover, ParetoCL surpasses other representative state-of-the-art continual learning methods by a considerable margin. For instance, compared to VR-MCL (Wu et al. 2024a), ParetoCL achieves an improvement of approximately 6.6%, 24.2%, and 12.3% in terms of final average accuracy on Seq-CIFAR10, Seq-CIFAR100, and Seq-TinyImageNet, respectively, under the online setting.

## Ablation Study

**Effectiveness of preference-conditioned learning** The primary motivation behind introducing Preference-Conditioned Learning is to efficiently learn a set of Pareto optimal models that embody different stability-plasticity trade-offs. To investigate whether ParetoCL successfully learns the mapping between preference vectors and corresponding models with different trade-offs, we conduct experiments on Seq-CIFAR10 in the offline setting. For each stage of the model, we use ten uniformly distributed preference vectors as input and evaluate the model’s performance on the test sets of previous tasks and the current task. Figure 3 (left) illustrates the Pareto front approximated by ParetoCL. The well-distributed points along the Pareto front demonstrate that ParetoCL can effectively explore various stability-plasticity trade-offs. Interestingly, we observe that most solutions are clustered around the corner region of the Pareto front, suggesting that ParetoCL tends to favor balanced trade-offs that achieve good performance on both stability and plasticity objectives. This is a desirable

property for CL systems, as overly prioritizing stability or plasticity can hinder the model’s generalization ability.

**Effectiveness of dynamic preference adaptation** The purpose of Dynamic Preference Adaptation is to adapt the stability-plasticity trade-off on the fly during inference. To verify its effectiveness, we design a variant termed ParetoCL--, which employs Preference-Conditioned Learning during training but uses a static preference vector set at (0.5, 0.5) during inference. We conduct experiments on Seq-CIFAR100 to compare ParetoCL-- and ParetoCL, along with ER, CLSER, and DER++ under online setting. Figure 3(right) visualizes the average incremental accuracy of different methods. We observe that Dynamic Preference Adaptation can effectively improve performance. It is worth noting that ParetoCL-- demonstrates competitive performance compared to other baselines, indicating that learning a set of models with different preferences during training allows the feature extractor to capture data distributions corresponding to various stability-plasticity trade-offs. The results also supports our view of generalization enhancement brought by objectives augmentation. On the other hand, ParetoCL-- can also serve as a simplified version of ParetoCL, reducing the computational cost of adapting to each sample.

**Impact of different memory buffer size** In real-world applications, the storage cost of the memory buffer can vary significantly. For instance, storing samples on cloud servers is relatively easy, while it can be expensive on edge computing devices. Therefore, an ideal CL method should perform well under various buffer size settings. To investigate if ParetoCL maintains its performance across different memory buffer sizes, we conduct experiments on Seq-CIFAR100 in an online setting with varying  $M$  as shown in Table 2. The results demonstrate that the performance of most methods improves as the memory buffer size  $M$  increases. Notably, our ParetoCL consistently outperforms other baselines, suggesting that ParetoCL is effective and robust to different memory constraints. And we observe that GEM (Lopez-Paz and Ranzato 2017) and A-GEM (Chaudhry et al. 2018), two gradient alignment replay methods, are insensitive to the buffer size, which is consistent with the result of Fork-Merge (Jiang et al. 2024).

## Discussion

In this section, we delve deeper into the characteristics and performance of ParetoCL by addressing two key questions. Further analysis can be found in supplementary material.

**How does ParetoCL’s algorithm efficiency compare to other algorithms?** Balancing computational cost and performance is a practical challenge in continual learning, particularly in online settings or resource-constrained environments (Verwimp et al. 2023). We analyze the time complexity of ParetoCL by comparing its training and testing time with other representative CL methods on the Seq-CIFAR10 dataset, as shown in Table 3. Despite achieving superior performance, ParetoCL requires less than one-fourth of the training time compared to the state-of-the-art method VR-MCL. As ParetoCL is an extension of the basic ER method, we also increased the training epochs of ER from 1 to 3 for a fair comparison. With nearly the same training time, ER

Method	Time (s)	AAA	Acc
GEM (Lopez-Paz and Ranzato 2017)	436.14	36.83	19.75
La-MAML (Gupta, Yadav, and Paull 2020)	693.42	36.76	35.44
VR-MCL (Wu et al. 2024a)	1073.61	69.57	58.43
DER++ (Buzzega et al. 2020)	108.81	60.63	50.33
ER (Rolnick et al. 2018)	79.21	52.74	33.14
ER (3 epochs)	288.71	59.21	46.30
<b>ParetoCL (Ours)</b>	<b>224.62</b>	<b>70.89</b>	<b>59.95</b>

Table 3: Training Time Comparison of Different Continual Learning Methods on Seq-CIFAR10.

Method	Seq-CIFAR10		Seq-CIFAR100	
	AAA	Acc	AAA	Acc
MGDA (Sener and Koltun 2018)	62.61	52.27	24.23	16.54
Tchebycheff (Choo and Atkins 1983)	52.03	46.71	23.15	15.58
<b>ParetoCL (Ours)</b>	<b>70.89</b>	<b>59.95</b>	<b>33.04</b>	<b>24.45</b>

Table 4: Comparison of methods on Seq-CIFAR10 and Seq-CIFAR100 datasets.

achieves an accuracy of 46.30%, while ParetoCL reaches 59.95%, demonstrating that the performance gain is not solely due to multiple forward passes.

**Does learning a set of pareto optimal solutions lead to better generalization?** As mentioned in the Introduction, our method can be viewed as a form of objective augmentation that enhances generalization. This hypothesis is preliminarily validated in Figure 3(right). To further verify this conjecture and exclude the influence of optimization, we directly applied multi-objective optimization methods MGDA (Sener and Koltun 2018) and Tchebycheff Scalarization (Choo and Atkins 1983) to the plasticity and stability objectives ( $\mathcal{L}_{replay}$  and  $\mathcal{L}_{new}$ ). While these methods guarantee Pareto optimality, they do not simultaneously learn multiple objectives. Experiments on Seq-CIFAR10 and Seq-CIFAR100 demonstrate that ParetoCL outperforms these MOO variants, emphasizing the necessity of learning a set of models with different preferences.

## Conclusion

In this work, we investigate the dynamic stability-plasticity trade-off in continual learning and reformulate it as a multi-objective optimization problem. We propose ParetoCL, a novel framework that efficiently learns a set of Pareto optimal solutions representing different trade-offs through preference-conditioned learning and enables dynamic adaptation during inference. Extensive experiments on multiple benchmarks demonstrate the superiority of ParetoCL over state-of-the-art continual learning methods.

**Limitations and Future Work.** ParetoCL relies on a hypernetwork to generate the parameters of the preference-conditioned model. Exploring more sophisticated conditioning mechanisms may improve the performance. Additionally, applying the ParetoCL framework to other learning paradigms, such as continual reinforcement learning, is a promising direction for future research.

## Acknowledgments

The work was supported by the Research Grants Council of the Hong Kong Special Administrative Region, China [GRF Project No. CityU 11215622], the National Natural Science Foundations of China (Grants No.62376267) and the innoHK project.

## References

- Aljundi, R.; Lin, M.; Goujaud, B.; and Bengio, Y. 2019. Gradient based sample selection for online continual learning. *Advances in neural information processing systems*, 32.
- Arani, E.; Sarfraz, F.; and Zonooz, B. 2021. Learning Fast, Learning Slow: A General Continual Learning Method based on Complementary Learning System. In *International Conference on Learning Representations*.
- Buzzega, P.; Boschini, M.; Porrello, A.; Abati, D.; and Calderara, S. 2020. Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems*, 33: 15920–15930.
- Caccia, L.; Aljundi, R.; Asadi, N.; Tuytelaars, T.; Pineau, J.; and Belilovsky, E. 2021. New Insights on Reducing Abrupt Representation Change in Online Continual Learning. *ArXiv*, abs/2203.03798.
- Cai, Z.; Sener, O.; and Koltun, V. 2021. Online continual learning with natural distribution shifts: An empirical study with visual data. In *Proceedings of the IEEE/CVF international conference on computer vision*, 8281–8290.
- Chaudhry, A.; Ranzato, M.; Rohrbach, M.; and Elhoseiny, M. 2018. Efficient Lifelong Learning with A-GEM. *ArXiv*, abs/1812.00420.
- Choo, E. U.; and Atkins, D. R. 1983. Proper efficiency in nonconvex multicriteria programming. *Mathematics of Operations Research*, 8(3): 467–470.
- Chrysakakis, A.; and Moens, M.-F. 2023. Online Bias Correction for Task-Free Continual Learning. In *International Conference on Learning Representations*.
- Ehrgott, M. 2005. *Multicriteria optimization*, volume 491. Springer Science & Business Media.
- Guo, Y.; Liu, B.; and Zhao, D. 2022. Online Continual Learning through Mutual Information Maximization. In *International Conference on Machine Learning*.
- Gupta, G.; Yadav, K.; and Paull, L. 2020. La-MAML: Look-ahead Meta Learning for Continual Learning. *ArXiv*, abs/2007.13904.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hung, C.-Y.; Tu, C.-H.; Wu, C.-E.; Chen, C.-H.; Chan, Y.-M.; and Chen, C.-S. 2019. Compacting, picking and growing for unforgetting continual learning. *Advances in neural information processing systems*, 32.
- Huszár, F. 2017. Note on the quadratic penalties in elastic weight consolidation. *Proceedings of the National Academy of Sciences*, 115: E2496 – E2497.
- Jiang, J.; Chen, B.; Pan, J.; Wang, X.; Liu, D.; Jiang, J.; and Long, M. 2024. Forkmerge: Mitigating negative transfer in auxiliary-task learning. *Advances in Neural Information Processing Systems*, 36.
- LeCun, Y.; Bengio, Y.; and Hinton, G. 2015. Deep learning. *nature*, 521(7553): 436–444.
- Li, X.; Zhou, Y.; Wu, T.; Socher, R.; and Xiong, C. 2019. Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In *International conference on machine learning*, 3925–3934. PMLR.
- Li, Z.; and Hoiem, D. 2016. Learning without Forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40: 2935–2947.
- Lin, X.; Yang, Z.; and Zhang, Q. 2022. Pareto set learning for neural multi-objective combinatorial optimization. *arXiv preprint arXiv:2203.15386*.
- Lin, X.; Zhen, H.-L.; Li, Z.; Zhang, Q.-F.; and Kwong, S. 2019. Pareto multi-task learning. *Advances in neural information processing systems*, 32.
- Liu, C.; Xu, X.; and Hu, D. 2014. Multiobjective reinforcement learning: A comprehensive overview. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 45(3): 385–398.
- Lopez-Paz, D.; and Ranzato, M. 2017. Gradient Episodic Memory for Continual Learning. In *Neural Information Processing Systems*.
- Lu, Z.; Whalen, I.; Boddeti, V.; Dhebar, Y.; Deb, K.; Goodman, E.; and Banzhaf, W. 2019. Nsga-net: neural architecture search using multi-objective genetic algorithm. In *Proceedings of the genetic and evolutionary computation conference*, 419–427.
- Mallya, A.; Davis, D.; and Lazebnik, S. 2018. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *Proceedings of the European conference on computer vision (ECCV)*, 67–82.
- Miettinen, K. 1999. *Nonlinear multiobjective optimization*, volume 12. Springer Science & Business Media.
- Navon, A.; Shamsian, A.; Chechik, G.; and Fetaya, E. 2020. Learning the pareto front with hypernetworks. *arXiv preprint arXiv:2010.04104*.
- Rebuffi, S.-A.; Kolesnikov, A.; Sperl, G.; and Lampert, C. H. 2017. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2001–2010.
- Riemer, M.; Cases, I.; Ajemian, R.; Liu, M.; Rish, I.; Tu, Y.; and Tesauro, G. 2018. Learning to learn without forgetting by maximizing transfer and minimizing interference. *arXiv preprint arXiv:1810.11910*.
- Rolnick, D.; Ahuja, A.; Schwarz, J.; Lillicrap, T. P.; and Wayne, G. 2018. Experience Replay for Continual Learning. In *Neural Information Processing Systems*.
- Sener, O.; and Koltun, V. 2018. Multi-task learning as multi-objective optimization. *Advances in neural information processing systems*, 31.

- Serra, J.; Suris, D.; Miron, M.; and Karatzoglou, A. 2018. Overcoming catastrophic forgetting with hard attention to the task. In *International conference on machine learning*, 4548–4557. PMLR.
- Verwimp, E.; Ben-David, S.; Bethge, M.; Cossu, A.; Geperth, A.; Hayes, T. L.; Hüllermeier, E.; Kanan, C.; Kudithipudi, D.; Lampert, C. H.; et al. 2023. Continual learning: Applications and the road forward. *arXiv preprint arXiv:2311.11908*.
- Wang, L.; Zhang, X.; Su, H.; and Zhu, J. 2024. A comprehensive survey of continual learning: Theory, method and application. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Wu, H.; Ma, C.; Mitra, B.; Diaz, F.; and Liu, X. 2022. A multi-objective optimization framework for multi-stakeholder fairness-aware recommendation. *ACM Transactions on Information Systems*, 41(2): 1–29.
- Wu, Y.; Huang, L.-K.; Wang, R.; Meng, D.; and Wei, Y. 2024a. Meta Continual Learning Revisited: Implicitly Enhancing Online Hessian Approximation via Variance Reduction. In *The Twelfth International Conference on Learning Representations*.
- Wu, Y.; Wang, H.; Zhao, P.; Zheng, Y.; Wei, Y.; and Huang, L.-K. 2024b. Mitigating Catastrophic Forgetting in Online Continual Learning by Modeling Previous Task Interrelations via Pareto Optimization. In *International Conference on Machine Learning*.
- Yang, Y.; Zhou, D.-W.; Zhan, D.-C.; Xiong, H.; and Jiang, Y. 2019. Adaptive deep models for incremental learning: Considering capacity scalability and sustainability. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 74–82.
- Yang, Y.; Zhou, D.-W.; Zhan, D.-C.; Xiong, H.; Jiang, Y.; and Yang, J. 2021. Cost-effective incremental deep model: Matching model capacity with the least sampling. *IEEE Transactions on Knowledge and Data Engineering*, 35(4): 3575–3588.
- Yoon, J.; Yang, E.; Lee, J.; and Hwang, S. J. 2017. Lifelong learning with dynamically expandable networks. *arXiv preprint arXiv:1708.01547*.
- Zenke, F.; Poole, B.; and Ganguli, S. 2017. Continual Learning Through Synaptic Intelligence. *Proceedings of machine learning research*, 70: 3987–3995.
- Zhu, F.; Zhang, X.-Y.; Wang, C.; Yin, F.; and Liu, C.-L. 2021. Prototype augmentation and self-supervision for incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5871–5880.