

Rotation Augmented Distillation for Exemplar-Free Class Incremental Learning with Detailed Analysis

Xiuwei Chen¹ and Xiaobin Chang^{1,2,3}✉

¹ School of Artificial Intelligence, Sun Yat-sen University, China

² Guangdong Key Laboratory of Big Data Analysis and Processing, Guangzhou 510006, P.R.China

³ Key Laboratory of Machine Intelligence and Advanced Computing, Ministry of Education, China

chenxw83@mail2.sysu.edu.cn, changxb3@mail.sysu.edu.cn

Abstract. Class incremental learning (CIL) aims to recognize both the old and new classes along the increment tasks. Deep neural networks in CIL suffer from catastrophic forgetting and some approaches rely on saving exemplars from previous tasks, known as the exemplar-based setting, to alleviate this problem. On the contrary, this paper focuses on the Exemplar-Free setting with no old class sample preserved. Balancing the plasticity and stability in deep feature learning with only supervision from new classes is more challenging. Most existing Exemplar-Free CIL methods report the overall performance only and lack further analysis. In this work, different methods are examined with complementary metrics in greater detail. Moreover, we propose a simple CIL method, Rotation Augmented Distillation (RAD), which achieves one of the top-tier performances under the Exemplar-Free setting. Detailed analysis shows our RAD benefits from the superior balance between plasticity and stability. Finally, more challenging exemplar-free settings with fewer initial classes are undertaken for further demonstrations and comparisons among the state-of-the-art methods.

Keywords: Class Incremental Learning · Catastrophic Forgetting · Exemplar Free.

1 Introduction

AI agents, e.g., deep neural networks (DNNs), deployed in the real world face an ever-changing environment, i.e., with new concepts and categories continually emerging [3,4,5]. Therefore, class incremental learning (CIL) [6,7,8] has attracted much attention as it aims to equip deep models with the capacity to continuously handle new categories. However, learning to discriminate sets of disjoint classes sequentially with deep models is challenging, as they are prone to entirely forgetting the previous knowledge, thus resulting in severe performance

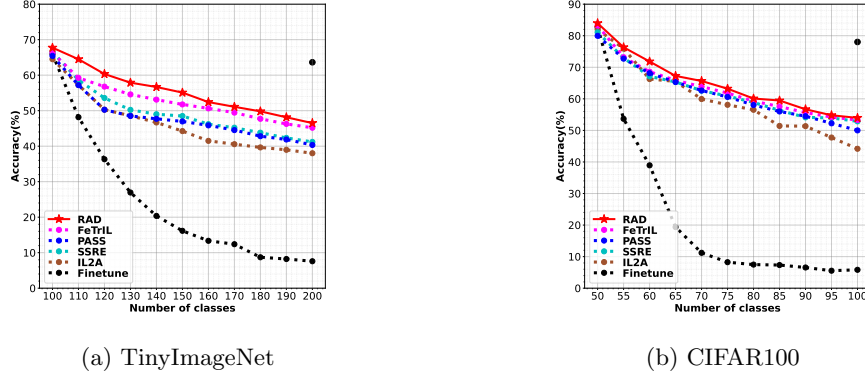


Fig. 1: Incremental Accuracy of TinyImageNet and CIFAR100 with 10 incremental steps. The top-1 accuracy (%) after learning each task is shown. Existing SOTA methods achieve similar performance. The proposed Rotation Augmented Distillation (RAD) achieves the SOTA performance as well. The black dot in the upper right corner indicates the upper bound that the model trained on all the data. The black dotted line indicates the lower bound, a simple finetune method.

degradation of old tasks. It is known as catastrophic forgetting [9]. To maintain the old knowledge, a small number of exemplars from the previous tasks can be stored in the memory buffer, known as the exemplar-based methods [8,6,10]. However, exemplars from previous tasks can be unavailable due to constraints, e.g., user privacy or device limitations.

To this end, this paper focuses on exemplar-free class incremental learning (EFCIL) setting [11,12,13,15,17]. Model distillation [16] plays an important role in preserving past knowledge by existing methods [11,12,13,14]. Both the feature extractor and new classifier are learned with the classification loss on the new data and knowledge distillation from the previous model. The variations of deep features during the new task training are penalized by distillation loss. Other methods [15,18] learn a feature extractor in the initial task only and then fixed it for increment tasks. New classifiers are learned on the extracted features of new tasks while the old class statistics, e.g., prototype vectors and covariance matrices, are also computed and preserved. The state-of-the-art (SOTA) methods [11,12,13,15] follow either the two paradigms mentioned above. As shown in Fig. 1, these SOTA methods report similar overall performance, i.e., average incremental accuracy. However, further detailed analysis is not conducted within these methods. In this work, we further incorporate the measures of Forgetting and Intransigence as in [34]. The performance of SOTA EFCIL methods [11,12,15] under various settings is reproduced for more detailed comparisons and analysis.

As EFCIL is a challenging task, the existing experimental setting is with half of the dataset in the initial task. Therefore, a model trained on such an initial task

may be strong enough for the following incremental tasks. A baseline method, named Feat*, is proposed for such demonstration and comparison. Specifically, a deep feature extractor is trained with the initial data and frozen. NME classifier is then used to discriminate different classes across incremental steps. However, Feat* achieves the SOTA level performance under the existing setting. In this work, we propose a simple yet effective method, Rotation Augmented Distillation (RAD), to enable the continuous training of the entire deep model along the increment tasks. On the one hand, the data augmentation used provides the plasticity by introducing the varied training samples during training. On the other hand, the knowledge distillation can achieve the stability by alleviating the forgetting of past knowledge. To alleviate the bias of the strong initial model, a more challenging setting with much fewer data in the initial task is introduced. Nevertheless, our RAD still achieves one of the best results among the SOTA methods due to its superior balance of stability and plasticity, as revealed by the detailed analysis.

The main contributions of this work are summarized as follows:

1. We provide a more detailed analysis of the EFCIL methods, rather than the overall results only in existing works. Specifically, the complementary metrics, forgetting and intransigence, are also used to evaluate SOTA methods. Detailed comparisons and analyses are thus enabled.
2. A simple and intuitive method, Rotation Augmented Distillation (RAD), is designed to better alleviate the plasticity-stability dilemma. Its effectiveness is demonstrated by its superior performance under various EFCIL settings.
3. A new challenging setting is provided to alleviate the bias brought by the strong initial model. Detailed comparisons and analyses are also conducted.

2 Related Work

Class incremental learning aims for a well-performing deep learner that can sequentially learn from the streaming data. Its main challenge is Catastrophic Forgetting [9] which depicts the deep models prone to entirely forgetting the knowledge learned from previous tasks. Different strategies have been proposed to handle this issue. Regularization strategies such as elastic weight consolidation (EWC) [19] use different metrics to identify and penalize the changes of important parameters of the original network when learning a new task. Rehearsal strategies [1,8,26,27] are widely adopted as well. The model is permitted to access data from previous tasks partially by maintaining a relatively small memory, enabling it to directly recall the knowledge from the previous data and mitigate forgetting. iCaRL [8] stores a subset of samples per class by selecting the good approximations to class means in the feature space. However, access to the exemplars of old tasks is not guaranteed and could be limited to data security and privacy constraints [25].

The exemplar-free class incremental learning (EFCIL) [2,11,12,13,15,17,18] is a challenging setting, where no data sample of old tasks can be directly stored. Existing EFCIL methods either use regularization to update the deep model for

each incremental step [19,20] or adapt distillation to preserve past knowledge by penalizing variations for past classes during model updates [11,12,13]. Moreover, the prototypes of old tasks can be exploited in conjunction with distillation to improve overall performance, as shown in [11,12,13]. Specifically, PASS [12] proposes prototype augmentation in the deep feature space to improve the discrimination of classes learned in different increment tasks. A prototype selection mechanism is proposed in SSRE [11] to better enhance the discrimination between old and new classes. Feature generation for past classes is introduced in IL2A [13] by leveraging information about the class distribution. However, IL2A is inefficient to scale up for a large number of classes since a covariance matrix needs to be stored for each class. Inspired by the transfer learning scheme [28], independent shallow classifiers can be learned based on a fixed feature extractor, as in DeeSIL [18]. A pseudo-features generator is further exploited in FeTrIL [15] to create representations of past classes to improve classifier learning across tasks. The proposed Rotation Augmented Distillation (RAD) method is end-to-end trainable on the full model and simply based on a distillation strategy [16] and rotation data augmentation [21]. Detailed comparisons show that RAD benefits from the superior balance between stability and plasticity in EFCIL model training.

3 Methodology

3.1 Preliminary

In class incremental learning, a model is learned from an initial task (0) a sequence of T incremental tasks, where each task t has a set of n_t different classes $C_t = \{c_{t,1}, \dots, c_{t,n_t}\}$. The classes in different tasks are disjoint, $C_i \cap C_j = \emptyset, i \neq j, i, j \in \{0, \dots, T\}$. The training data of task t is denoted as D_t . D_t consists of data tuples (x, y) where x is an image and y is its corresponding ground-truth class label. A deep classification model Φ consists of two modules, the feature extractor F and the classifier head h . F is parameterized with θ and the feature representation of image x is obtained via $F(x) \in \mathbb{R}^d$. The classifier head h is parameterized with ω .

Under the exemplar-free class incremental learning (EFCIL) setting, a model can only access D_t when training on task t . At the initial task, the classification model Φ_0 is trained under the full supervision of D_0 and resulting in F_0 and h_0 . At the incremental task $t, t \in \{1, \dots, T\}$, Φ_t is partially initialized with θ_{t-1} and trained with D_t . The corresponding feature extractor F_t and classifier head h_t are learned. The overall classifier H_t is an aggregation of a set of task-specific classifiers $h_i, i = \{1, \dots, t\}$, $H_t = \{h_0, h_2, \dots, h_t\}$. During testing, data samples are from all observed classes so far with balanced distributions.

3.2 A baseline method: Feat*

The feature extractor learned at the initial task is frozen and denoted as F_0^* . The corresponding classifier h_0 is abandoned. F_0^* is used as the feature extractor across all tasks and no further training on the deep classifier is needed. This

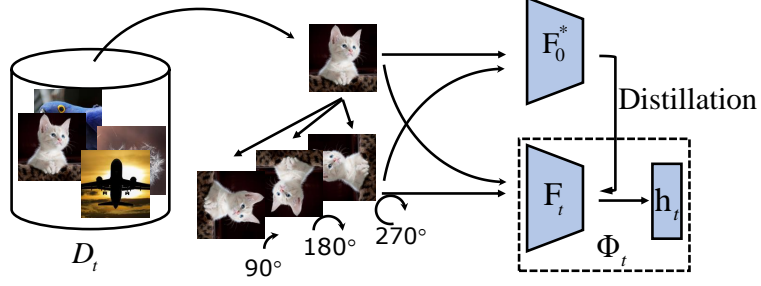


Fig. 2: Illustrations of the proposed Rotation Augmented Distillation (RAD) method for exemplar-free class incremental learning at task t . * indicates the corresponding module is frozen at training.

method is denoted as Feat*. F_0^* forwards all training samples in $D_t, t \in \{0, \dots, T\}$ once for the class-specific mean feature vectors and such prototypes are preserved. NME classifier can then be applied to the seen prototypes so far during testing. Different from Feat*, a holistic classifier H_t for the seen classes till task t is trained in FeTrIL based on the pseudo-feature generation. During testing of FeTrIL, H_t rather than NME classifier is exploited.

3.3 Rotation Augmented Distillation

The deep classification model Φ_t can be end-to-end learned with two simple techniques, rotation data augmentation and knowledge distillation, as illustrated in Fig. 2. Specifically, for each class, we rotate each training sample x with 90, 180 and 270 degrees [21] and obtained the augmented sample x' :

$$x' = \text{rotate}(x, \delta), \delta \in \{90, 180, 270\}. \quad (1)$$

Each augmented sample x' subject to a rotated degree is assigned a new label y' , extending the original K -class problem to a new $4K$ -class problem. The augmented dataset of task t is denoted as D'_t .

Both the feature extractor F_t and the classifier H_t of Φ_t are jointly optimized in RAD. Based on dataset D_t and the augmented one D'_t , the cross-entropy loss is computed,

$$\mathcal{L}_c = \sum_{(x,y) \in D_t} \text{CE}(\Phi_t(x), y) + \sum_{(x',y') \in D'_t} \text{CE}(\Phi_t(x'), y'). \quad (2)$$

To alleviate the mismatch between the saved old prototypes and the feature extractor, the knowledge distillation [16] is employed to regularize learning of the feature extractor. Specifically, we restrain the feature extractor by matching the features of new data extracted by the current model with that of the initial

model F_0^* :

$$\mathcal{L}_{distil} = \sum_{(x,y) \in D_t} \text{KL}(F_t(x; \theta_t) || F_0^*(x; \theta_0)) + \sum_{(x',y') \in D'_t} \text{KL}(F_t(x'; \theta_t) || F_0^*(x'; \theta_0)). \quad (3)$$

The total loss of RAD comprises two terms,

$$\mathcal{L}_{all} = \alpha \mathcal{L}_c + \beta \mathcal{L}_{distil}, \quad (4)$$

with α and β are balancing hyper-parameters. Both of them are set to 1 across all experiments. The learning objective of RAD becomes $\min_{\theta_t, \omega_t} \mathcal{L}_{all}$.

4 Experiments

Datasets The exemplar-free class incremental learning (EFCIL) is conducted on three datasets. Two large-scale datasets are TinyImageNet [29] and ImageNet100 [30]. The medium-size dataset used is CIFAR100 [31].

Protocols Two EFCIL settings are followed in our experiments. (1) Conventional EFCIL setting. The conventional setting usually with half of the data as initial tasks. TinyImageNet contains images from 200 classes in total and the initial task (0) includes half of the dataset, i.e., images from 100 classes, denoted as **B100**. The data of the remaining classes are split into T tasks. For example, $T=5$ corresponds to 5 incremental steps and each step contains the data of 20 classes. Therefore, this EFCIL setting of TinyImageNet is denoted as **B100 5 steps**. **B100 10 steps** and **B100 20 steps** are the other two EFCIL protocols of TinyImageNet. Both CIFAR100 and ImageNet100 are with 100 classes. They have the following three EFCIL protocols: **B50 5 steps**, **B50 10 steps** and **B40 20 steps**. (2) Challenging EFCIL setting. Training a method with half of the data can result in a strong initial model, which can introduce a non-negligible bias towards the following incremental learning performance. Therefore, a more challenging EFCIL setting is proposed by reducing the data in the initial task to half or even less. Specifically, the challenging EFCIL settings of TinyImageNet are **B50 5 steps**, **B50 10 steps** and **B50 25 steps**.

Implementation details The proposed Rotation Augmented Distillation (RAD) and the reproduced SOTA EFCIL methods [11,12,15] are implemented with PyTorch [32] and trained and tested on NVIDIA GTX 3080Ti GPU. Other results are from [15]. The initial model is trained for 200 epochs, and the learning rate is 0.1 and gradually reduces to zero with a cosine annealing scheduler. For incremental tasks, we train models for 30 epochs, and the initial learning rate is 0.001 with a cosine annealing scheduler. For a fair comparison, we adopt the default ConvNet backbone: an 18-layer ResNet [33]. The reproduced results of SOTA methods are with the same augmentation and assigned to incremental tasks using the same random seed as ours.

Evaluation metrics Three complementary metrics are used throughout the experiments. **Overall performance** is typically evaluated by average incremental accuracy [8]. After each batch of classes, the evaluation result is the

Table 1: Results of different EFCIL methods under the conventional setting. The overall performance, average incremental accuracy (%), is reported. Best results in red, second best in blue.

Methods	TinyImageNet			ImageNet100			CIFAR100		
	5 steps	10 steps	20 steps	5 steps	10 steps	20 steps	5 steps	10 steps	20 steps
Finetune	28.8	24.0	21.6	31.5	25.7	20.2	27.8	22.4	19.5
EWC [19]	18.8	15.8	12.4	-	20.4	-	24.5	21.2	15.9
LwF-MC [8]	29.1	23.1	17.4	-	31.2	-	45.9	27.4	20.1
DeeSIL [18]	49.8	43.9	34.1	67.9	60.1	50.5	60.0	50.6	38.1
LUCIR [24]	41.7	28.1	18.9	56.8	41.4	28.5	51.2	41.1	25.2
MUC [22]	32.6	26.6	21.9	-	35.1	-	49.4	30.2	21.3
SDC [17]	-	-	-	-	61.2	-	56.8	57.0	58.9
ABD [23]	-	-	-	-	-	-	63.8	62.5	57.4
IL2A [13]	47.3	44.7	40.0	-	-	-	66.0	60.3	57.9
PASS [12]	49.6	47.3	42.1	64.4	61.8	51.3	63.5	61.8	58.1
SSRE [11]	50.4	48.9	48.2	65.4	62.1	58.8	65.9	65.0	61.7
FeTrIL [15]	54.8	53.1	52.2	72.2	71.2	67.1	66.3	65.2	61.5
Feat*	53.3	53.0	52.8	70.1	69.9	65.1	63.9	63.6	59.8
RAD	55.9	55.6	55.2	72.4	71.8	67.4	66.5	65.3	61.9

classification accuracy curve. If a single number is preferable, the average of these accuracies is reported as average incremental accuracy. **Forgetting** [34] is defined to estimate the forgetting of previous tasks. The forgetting measure for the t -th task after the model has been incrementally trained up to task k as:

$$f_t^k = \max_{l \in \{0, \dots, k-1\}} (a_{l,t} - a_{k,t}), \forall t < k. \quad (5)$$

Note, $a_{m,n}$ is the accuracy of task n after training task m . The average forgetting at k -th task is then defined as $\mathbb{F}_k = \frac{1}{k} \sum_{i=0}^{k-1} f_i^k$. Lower \mathbb{F}_k implies less forgetting on previous tasks. **Intransigence** [34] measures the inability of a model to learn new tasks. We train a reference model with dataset $\cup_{t=0}^k D_t$ and measure its accuracy on the held-out set of the k -th task, denoted as a_k^- . The intransigence for the k -th task as:

$$\mathbb{I}_k = a_k^- - a_{k,k}, \quad (6)$$

where $a_{k,k}$ denotes the accuracy on the k -th task when trained up to task k in an incremental manner. Note, the lower the \mathbb{I}_k the better the model.

4.1 Detailed Analysis under conventional EFCIL setting

The overall results of different methods under conventional EFCIL settings of TinyImageNet, ImageNet100 and CIFAR100 are reported in Tab. 1. The simple baseline method, Feat*, achieves similar results compared to the SOTA methods. For example, the margins between FeTrIL and Feat* are no more than 2%. Feat* is observed to be slightly better (0.6% improvements) than FeTrIL under the TinyImageNet B100 20 steps setting. Therefore, it demonstrates that a

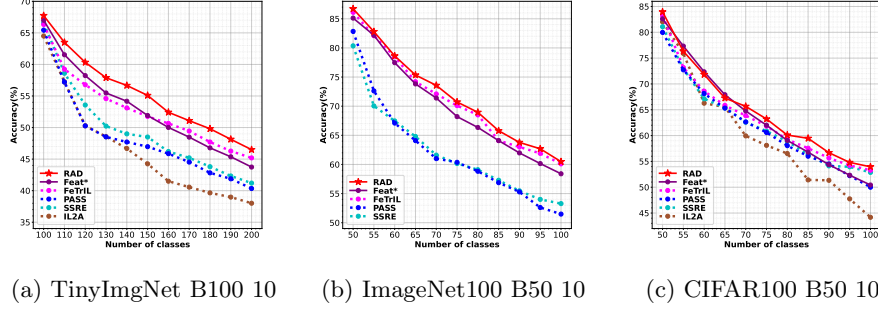


Fig. 3: Incremental Accuracy Curves of the SOTA methods. Each point represents the incremental classification accuracy (%) after model learning on each task.

Table 2: Results on TinyImageNet and ImageNet100 with intransigence and forgetting metrics. \mathbb{I} represents Intransigence (%), \mathbb{F} represents Average Forgetting (%). Best results in red, second best in blue.

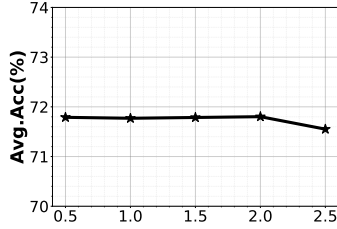
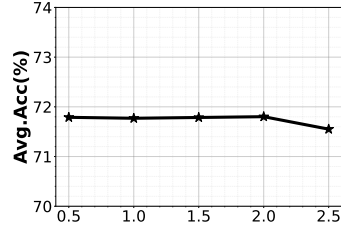
Methods	TinyImageNet						ImageNet100					
	5 steps		10 steps		20 steps		5 steps		10 steps		20 steps	
	\mathbb{I} (\downarrow)	\mathbb{F} (\downarrow)	\mathbb{I} (\downarrow)	\mathbb{F} (\downarrow)	\mathbb{I} (\downarrow)	\mathbb{F} (\downarrow)	\mathbb{I} (\downarrow)	\mathbb{F} (\downarrow)	\mathbb{I} (\downarrow)	\mathbb{F} (\downarrow)	\mathbb{I} (\downarrow)	\mathbb{F} (\downarrow)
PASS [12]	22.5	15.4	24.3	20.6	29.6	25.2	27.0	19.3	32.0	25.7	42.4	31.6
SSRE [11]	21.3	16.1	23.2	21.1	25.4	24.3	25.7	24.9	30.2	30.5	34.0	35.0
FeTrIL [15]	18.9	12.6	19.4	11.8	21.0	12.9	21.7	14.7	23.4	15.4	27.7	18.1
Feat*	20.9	8.9	20.9	9.0	20.9	8.9	25.1	8.9	25.1	9.8	31.9	10.4
RAD	18.0	8.7	18.1	9.6	18.4	10.2	21.5	9.7	23.0	11.4	27.6	11.0

model trained with half of the dataset at the initial task may already equip sufficient classification capability for the subsequent incremental tasks even without further learning. This conclusion is frequently discussed in the previous sections of this paper and inspires the introduction of the challenging EFCIL setting. The proposed RAD consistently achieves the best results under conventional EFCIL settings of all datasets, as shown in Tab. 1. For example, RAD is better than FeTrIL on the largest dataset, TinyImageNet. The more incremental steps the larger improvements. Such improvement increased from 1.1% under 5 steps to 2.5% under 10 steps and reached 3.0% under 20 incremental steps. Detailed comparisons among different methods along the incremental learning procedure are also illustrated in Fig. 3. The curve of a superior method is usually above the inferior ones. For example, the superiority of RAD over FeTrIL can be demonstrated by their curves in Fig. 3(a).

To provide more detailed analyses on the SOTA EFCIL methods, their performance on another two metrics, forgetting and intransigence, are reported in Tab. 2. The proposed baseline Feat* achieves the lowest \mathbb{F} under most cases. It shows that Feat* suffers the least from forgetting since the deep model is frozen

Table 3: Ablative Study of the proposed RAD. Results of TinyImageNet B100 10 steps and ImageNet100 B50 10 steps are reported.

Components		TinyImageNet			ImageNet100		
Rotation	Distillation	Avg (\uparrow)	\mathbb{I} (\downarrow)	\mathbb{F} (\downarrow)	Avg (\uparrow)	\mathbb{I} (\downarrow)	\mathbb{F} (\downarrow)
		24.0	57.0	76.1	25.7	77.3	90.2
✓		11.2	60.1	74.3	13.7	78.7	87.7
	✓	52.9	21.1	11.6	70.0	24.2	26.7
✓	✓	55.6	18.1	9.6	71.8	23.0	11.4

(a) α (b) β Fig. 4: Impacts of varying α and β on overall results of ImageNet100 B50 10 steps.

at the initial state and does not fine-tune on the following tasks. Moreover, the \mathbb{I} of Feat* still achieves mid-level performance. It further demonstrates that the initial model trained on half of the dataset can provide sufficient discriminative power for the rest classes, as described above. FeTrIL achieves the best overall performance among the existing SOTA methods, as shown in Tab. 1. Therefore, both \mathbb{I} and \mathbb{F} of FeTrIL are lower than those of its counterpart SOTA methods, as shown in Tab. 2. Similarly, the proposed RAD achieves the best overall performance on various datasets. This is consistent with the results in Tab. 2. The \mathbb{I} and \mathbb{F} of RAD are either the best or second best results. Specifically, RAD is only slightly worse than Feat* on the \mathbb{F} metric. This is compensated by the best \mathbb{I} performance of RAD. Moreover, RAD is consistently better than FeTrIL on both metrics. Therefore, the superiority of our RAD relies on the more balanced performance between plasticity and stability.

Ablation study Contributions of different components in the proposed RAD are illustrated in Tab. 3. First, applying data rotation in model finetuning harms the overall performance, as the augmented images seem to alleviate the forgetting issue (lower \mathbb{F}) but make less discriminative new knowledge learned (higher \mathbb{I}). Second, model distillation clearly boosts the overall performance. Distillation helps the model learning not only to defy forgetting the old knowledge but also to better capture the discriminative information from new tasks, as its \mathbb{F} and \mathbb{I} are relatively low. Finally, RAD combines these two techniques. Detailed analysis shows that model distillation benefits from rotation data augmentation and then achieves a superior balance between plasticity and stability.

Table 4: Results of TinyImageNet under challenging EFCIL setting with average incremental accuracy (Avg.), intransigence (\mathbb{I}) and forgetting metrics (\mathbb{F}) reported. Best results in red, second best in blue.

Methods	5 steps			10 steps			25 steps		
	Avg. (\uparrow)	\mathbb{I} (\downarrow)	\mathbb{F} (\downarrow)	Avg. (\uparrow)	\mathbb{I} (\downarrow)	\mathbb{F} (\downarrow)	Avg. (\uparrow)	\mathbb{I} (\downarrow)	\mathbb{F} (\downarrow)
PASS [12]	45.1	30.3	15.5	42.5	32.0	20.2	39.7	34.6	26.5
SSRE [11]	43.6	32.2	12.2	41.1	33.7	14.8	40.1	33.9	21.5
FeTrIL [15]	47.9	29.7	13.5	46.5	30.4	13.4	45.1	30.8	13.9
Feat*	44.1	34.6	10.1	43.3	34.6	9.9	42.7	34.6	10.1
RAD	48.6	30.2	11.6	47.7	30.3	11.5	47.0	30.2	12.3

Impacts of balancing hyper-parameters There are two balancing hyper-parameters, α and β . They are set to 1.0 by default. As shown in Fig. 4, Our method is not sensitive to such changes.

4.2 Detailed Analysis under challenging EFCIL setting

The results of challenging EFCIL setting are reported in Tab. 4. Comparing the overall performance, the average incremental accuracy, of the methods in Tab. 4 with those in Tab. 1, all of them experience significant drops, usually more than 6%, due to much less data being provided in the initial task under the new setting. Specifically, the baseline Feat* suffers from the worst performance degradation, from 9.2% to 10.1%, under different incremental steps. This is because Feat* is a simple CIL method that relies on the initial task training only. With less initial data provided, the corresponding model can be much less generalizable on the new data from unseen classes. Under the challenging EFCIL setting, the proposed RAD still achieves the best overall performance against the existing SOTA methods. With the \mathbb{I} and \mathbb{F} of different methods compared in Tab. 4, RAD consistently achieve either the best or second best performance on these metrics. It demonstrates that RAD can better alleviate the plasticity and stability dilemma in EFCIL than existing methods.

5 Conclusion

We provide a more detailed analysis and comparison of different exemplar-free class incremental learning (EFCIL) methods than many existing works. Besides the overall performance, i.e., the average incremental accuracy, two complementary metrics, the Forgetting and the Intransigence, are included to measure the EFCIL methods from the perspectives of stability and plasticity. A simple yet effective EFCIL method, Rotation Augmented Distillation (RAD), has been proposed. RAD consistently achieves one of the state-of-the-art overall performances under various EFCIL settings. Based on the detailed analysis, we find that the superiority of RAD comes from the more balanced performance between plasticity and stability. Moreover, a challenging EFCIL new setting with much fewer

data in initial tasks is proposed. It aims to alleviate the bias of a strong initial pre-trained model and stand out the incremental learning performance.

Acknowledgement This research is supported by the National Science Foundation for Young Scientists of China (No. 62106289).

References

1. Zhong, Changhong, et al. Discriminative Distillation to Reduce Class Confusion in Continual Learning. In Proceedings of the Pattern Recognition and Computer Vision (PRCV), 2022. 3
2. Huang, Tianjun, Weiwei Qu, and Jianguo Zhang. Continual representation learning via auto-weighted latent embeddings on person reid. In Proceedings of the Pattern Recognition and Computer Vision (PRCV), 2021. 3
3. Shaheen K, Hanif M A, Hasan O, et al. Continual learning for real-world autonomous systems: Algorithms, challenges and frameworks. Journal of Intelligent & Robotic Systems, 2022. 1
4. Geiger A, Lenz P, Urtasun R. Are we ready for autonomous driving? the kitti vision benchmark suite. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition(CVPR), 2012. 1
5. Li K, Chen K, Wang H, et al. Coda: A real-world road corner case dataset for object detection in autonomous driving. In Proceedings of the European Conference on Computer Vision (ECCV), 2022. 1
6. Rolnick D, Ahuja A, Schwarz J, et al. Experience replay for continual learning. In Advances in Neural Information Processing Systems (NeurIPS), 2019. 1, 2
7. Masana M, Liu X, Twardowski B, et al. Class-incremental learning: survey and performance evaluation on image classification. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 2022. 1
8. Rebuffi S A, Kolesnikov A, Sperl G, et al. icarl: Incremental classifier and representation learning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017. 1, 2, 3, 6, 7
9. French, Robert M. Catastrophic forgetting in connectionist networks. Trends in cognitive sciences, 1999. 2, 3
10. Zhou D W, Yang Y, Zhan D C. Learning to classify with incremental new class. IEEE Transactions on Neural Networks and Learning Systems, 2021. 2
11. Zhu, Kai, et al. Self-sustaining representation expansion for non-exemplar class-incremental learning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2022. 2, 3, 4, 6, 7, 8, 10
12. Zhu, Fei, et al. Prototype augmentation and self-supervision for incremental learning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2021. 2, 3, 4, 6, 7, 8, 10
13. Zhu, Fei, et al. Class-incremental learning via dual augmentation. In Advances in Neural Information Processing Systems (NeurIPS), 2021. 2, 3, 4, 7
14. Xu, Qi, et al. Constructing deep spiking neural networks from artificial neural networks with knowledge distillation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2023. 2
15. Petit, Grégoire, et al. Fetrl: Feature translation for exemplar-free class-incremental learning. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision(WACV), 2023. 2, 3, 4, 6, 7, 8, 10

16. Hinton, Geoffrey, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531, 2015. 2, 4, 5
17. Yu, Lu, et al. Semantic drift compensation for class-incremental learning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2020. 2, 3, 7
18. Belouadah, Eden, and Adrian Popescu. DeeSIL: Deep-Shallow Incremental Learning. In Proceedings of the European Conference on Computer Vision (ECCV) Workshop, 2018. 2, 3, 4, 7
19. Kirkpatrick, James, et al. Overcoming catastrophic forgetting in neural networks. Proceedings of the National Academy of Sciences(PNAS), 2017. 3, 4, 7
20. Dhar, Prithviraj, et al. Learning without memorizing. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019. 4
21. Lee, Hankook, Sung Ju Hwang, and Jinwoo Shin. Self-supervised label augmentation via input transformations. In International Conference on Machine Learning (ICML), 2020. 4, 5
22. Liu, Yu, et al. More classifiers, less forgetting: A generic multi-classifier paradigm for incremental learning. In Proceedings of the European Conference on Computer Vision (ECCV), 2020. 7
23. Smith, James, et al. Always be dreaming: A new approach for data-free class-incremental learning. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2021. 7
24. Hou, Saihui, et al. Learning a unified classifier incrementally via rebalancing. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019. 7
25. Wang, Yabin, Zhiwu Huang, and Xiaopeng Hong. S-Prompts Learning with Pre-trained Transformers: An Occam’s Razor for Domain Incremental Learning. In Advances in Neural Information Processing Systems (NeurIPS), 2022. 3
26. Wu, Yue, et al. Large scale incremental learning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019. 3
27. Zhao, Bowen, et al. Maintaining discrimination and fairness in class incremental learning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2020. 3
28. Ginsca, Alexandru Lucian, et al. Large-scale image mining with flickr groups. In International Conference on MultiMedia Modeling, 2015. 4
29. Le, Ya, and Xuan Yang. Tiny imagenet visual recognition challenge. CS 231N 7.7 (2015): 3. 6
30. Russakovsky, Olga, et al. Imagenet large scale visual recognition challenge. International Journal of Computer Vision(IJCV), 2015. 6
31. Krizhevsky, Alex, and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical Report TR 2009, University of Toronto, Toronto, 2009. 6
32. Grossberg, Stephen T. Studies of mind and brain: Neural principles of learning, perception, development, cognition, and motor control. Springer Science & Business Media, 2012. 6
33. He, Kaiming, et al. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016. 6
34. Chaudhry, Arslan, et al. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In Proceedings of the European Conference on Computer Vision (ECCV), 2018. 2, 7