

# Demystifying the Base and Novel Performances for Few-shot Class-incremental Learning

**Jaehoon Oh**

*Graduate School of Data Science, KAIST, Korea*

JHOON.OH@KAIST.AC.KR

**Se-Young Yun**

*Graduate School of Artificial Intelligence, KAIST, Korea*

YUNSEYOUNG@KAIST.AC.KR

## Abstract

Few-shot class-incremental learning (FSCIL) has addressed challenging real-world scenarios where unseen novel classes continually arrive with few samples. In these scenarios, it is required to develop a model that recognizes the novel classes without forgetting prior knowledge. In other words, FSCIL aims to maintain the base performance and improve the novel performance simultaneously. However, there is little study to investigate the two performances separately. In this paper, we first decompose the entire model into four types of parameters and demonstrate that the tendency of the two performances varies greatly with the updated parameters when the novel classes appear. Based on the analysis, we propose a simple method for FSCIL, coined as NoNPC, which uses normalized prototype classifiers without further training for incremental novel classes. It is shown that our straightforward method has comparable performance with the sophisticated state-of-the-art algorithms.

**Keywords:** Few-shot class-incremental learning, parameter decomposition, prototype classifier

## 1. Introduction

Deep learning has achieved considerable success under the IID (independent and identically distribution) and stationary assumption. However, in real-world scenarios, dynamic and open environments are more natural, discouraging practitioners from developing deep models in many applications. To address this concern, class-incremental learning (CIL) has gained much attention (Belouadah and Popescu, 2019; Masana et al., 2020; Zhu et al., 2021; Shim et al., 2021; Mai et al., 2022), where the unseen *novel* classes continually appear. CIL aims to learn novel classes without forgetting the past knowledge, where previously seen data are not available due to privacy (Mai et al., 2022; Joseph et al., 2022) and memory (Fini et al., 2020) issues.

Conventional CIL methods have been studied based on a large amount of data for the novel classes. However, when only a small amount of training data for the novel classes is available, they are known to suffer from severe performance degradation (Shi et al., 2021). To address this challenging scenario, *few-shot class-incremental learning* (FSCIL) (Tao et al., 2020; Zhang et al., 2021; Shi et al., 2021; Zhou et al., 2022) has emerged. FSCIL learns base classes consisting of a huge amount of data (i.e., *many-shot*), and then learns novel classes incrementally with very few data (i.e., *few-shot*). In this regime, FSCIL

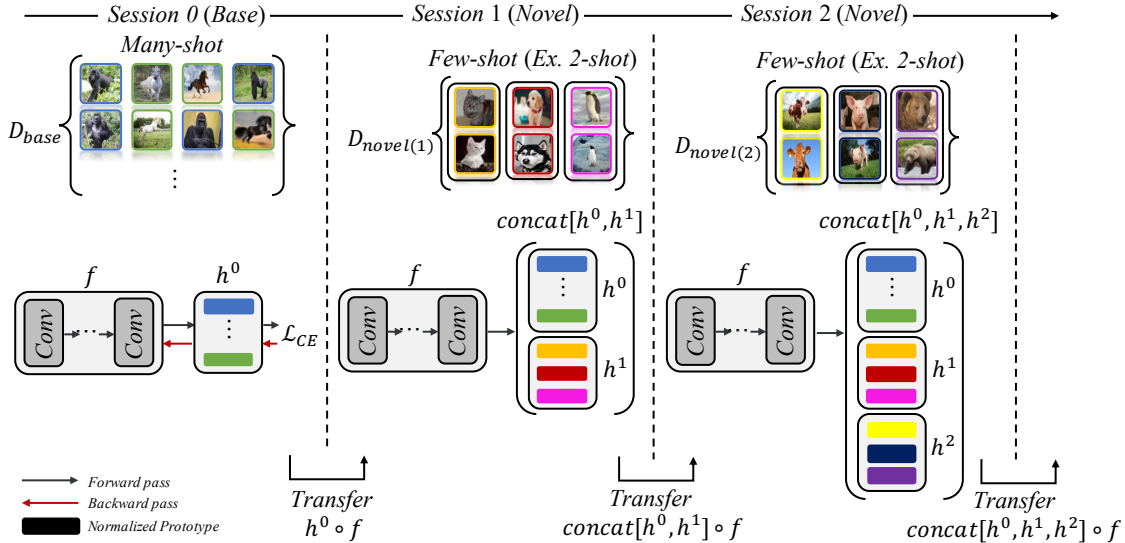


Figure 1: Overview of NoNPC (**No** Training with **N**ormalized **P**rototype **C**lassifiers).  $f$  indicates a feature extractor.  $h^0$  and  $h^i$  indicate classifiers for base and novel classes for session  $i$  ( $i > 0$ ), respectively. Our algorithm does not learn a model during novel sessions.

aims to recognize novel classes from few novel data that arrive sequentially without forgetting the knowledge of base classes. Namely, it is required to maintain the base performance and improve the novel performance simultaneously.

Recent methods for FSCIL have improved the *weighted* performance, which weighs the *base* and *novel* performances based on the number of base and novel classes. However, the improved weighted performance does not mean that *base* and *novel* performances are both improved. In this paper, we analyze which components in a model are responsible for the base and novel performances by parameter decomposition. Based on this analysis, we propose a simple yet effective method, **NoNPC**, which uses normalized prototype classifiers without training when novel classes appear. This simple algorithm achieves comparable performance to SOTA algorithms.

## 2. Related Work

**Few-shot Class-incremental Learning (FSCIL).** FSCIL is a combination of class-incremental learning (CIL) and few-shot learning (FSL), aiming to incrementally update a classifier with limited data from novel classes to discriminate all classes seen before. The incremental learning procedure consists of a base session followed by novel sessions. Recent FSCIL methods (Tao et al., 2020; Mazumder et al., 2021; Shi et al., 2021; Zhang et al., 2021; Zhou et al., 2022) can be categorized by which session they focus on, i.e., base or novel session. The former prepares to learn novel classes in the base session by revising the standard training process in the base session (Shi et al., 2021; Zhou et al., 2022) or training an additional network for the incoming novel classes (Zhang et al., 2021). On the other hand, the latter devise how to fine-tune a model incrementally when they encounter unseen classes in the novel sessions by imposing regularization on fine-tuning in the novel

Table 1: Notations of decomposed parameters and models according to the update parts for the current novel session  $i$  ( $i > 0$ ).

Notation	Description	Model	Description
$f$	Feature extractor	M1	No update
$h^0$	Classifier for base classes	M2	Update $h^i$
$h^{1:i-1}$	Classifiers for novel classes before session $i$ ( <i>i.e.</i> , $\{h^1, \dots, h^{i-1}\}$ )	M3	Update $h^{1:i-1}$ and $h^i$
$h^i$	Classifier for novel classes in session $i$	M4	Update $h^0, h^{1:i-1}$ and $h^i$
		M5	Update $f, h^0, h^{1:i-1}$ and $h^i$

session (Mazumder et al., 2021; Hersche et al., 2022). Previous studies focus on improving the weighted performance, while we analyze the *base* and *novel* performances separately and propose a simple yet effective method without additional costs.

**Parameter Decomposition.** Parameter decomposition is widely used technique for in-depth analysis or improved algorithms in many fields such as long-tailed distribution (Kang et al., 2020; Yu et al., 2020), federated learning (Arivazhagan et al., 2019; Collins et al., 2021; Oh et al., 2022), meta-learning (Lee and Choi, 2018; Raghu et al., 2020; Oh et al., 2021), and continual learning (Shi et al., 2021; Davari et al., 2022). Most of the prior works decompose the entire model into two parts, feature extractors and classifiers. Feature extractors can be transferred well, while classifiers are easily distorted under non-IID environments, meaning that they are susceptible to bias (Kang et al., 2020; Oh et al., 2022). We decompose entire parameters into four parts under FSCIL environments: extractors and three types of classifiers, described in Table 1, to investigate the base and novel performances for FSCIL.

### 3. Problem Setup

In this section, we formally summarize the problem setup. The procedure of FSCIL includes continuous sessions, where each session consists of training and evaluation. For the first session that we call base session (defined as session 0), a model  $h^0 \circ f$  is trained using  $D_{base}$  consisting of base classes, where  $f$  is a feature extractor and  $h^0$  is a classifier for the base classes. Note that during the base session, we can use abundant data from base classes. Let  $x$  be an input and  $d$  be the output dimension of  $f$ , then  $f(x) \in \mathbb{R}^d$  and  $h^0 \in \mathbb{R}^{K \times d}$  where  $K$  is the number of base classes. We define  $h_j^0$  as the  $j$ -th row vector of  $h^0$ .

For the subsequent session  $i$  ( $i \in \{1, \dots, S\}$ ) that we call *novel* sessions, an extended model ( $concat[h^0, h^1, \dots, h^i] \circ f$ ) is trained using  $D_{novel(i)}$  consisting of novel classes, where  $h^i$  is a classifier for novel classes in session  $i$ . Unlike  $D_{base}$ ,  $D_{novel(i)}$  consists of few samples, in general  $nk$  samples, where  $n$  is the number of incremental novel classes and  $k$  is the number of samples per class.  $k$  is 5 in our experiments. The incrementally extended model ( $concat[h, h^1, \dots, h^i] \circ f$ ) is evaluated for both base and incremental novel classes until session  $i$ . After  $S$  sessions, the final classifier ( $concat[h, h^1, \dots, h^S]$ ) is in  $\mathbb{R}^{(K+nS) \times d}$ .

For evaluation, we use three metrics: base, novel, and weighted performances. The base and novel performance indicate accuracy on base and novel classes, respectively. These performances are weighted based on the number of classes for calculating the weighted performance. All results are averaged by five runs. The detailed implementation is described in Appendix A.

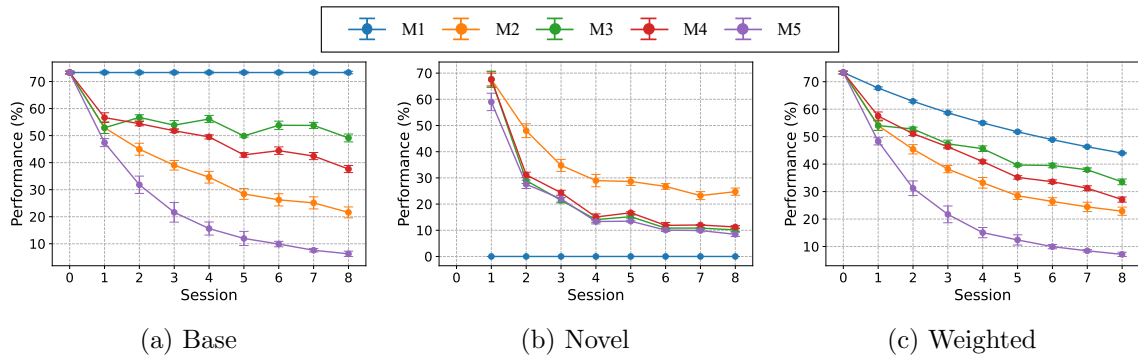


Figure 2: Base, novel, and weighted performances according to the update parts on CIFAR100. Models are described in Table 1.

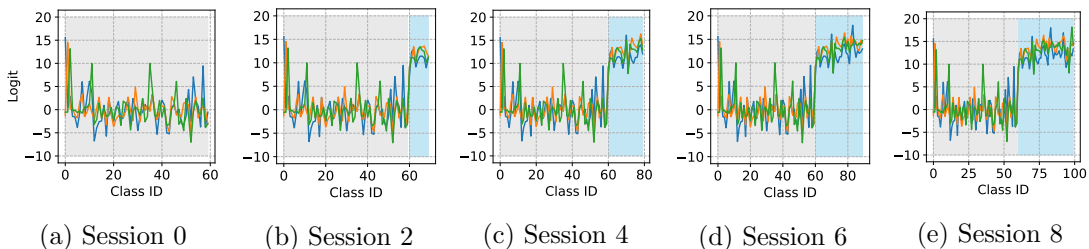


Figure 3: Logit distributions of base samples according to the session on CIFAR100 when training only  $h^i$ . Gray and blue backgrounds indicate the base and novel classes, respectively. Each line is the average of logits of samples belonging to class 0, 1, and 2.

#### 4. Analysis: Parameter Decomposition during Novel Sessions

To investigate which parameters are relevant to each performance, we decompose the entire model into four update parts described in Section 3, during novel sessions:  $f$ ,  $h^0$ ,  $h^{1:i-1}$ , and  $h^i$ . Figure 2 describes the base, novel, and weighted performances according to the decomposed update parts on CIFAR100. Following Tao et al. (2020), for the cases where an encoder  $f$  is not updated, the running statistics of batch normalization layers are also fixed based on  $D_{base}$ . This result provides interesting observations as follows<sup>1</sup>:

- No training during novel sessions (M1 in Figure 2) undoubtedly is the best on the base performance, whereas the worst on the novel performance. The weighted performance, which is the main evaluation measurement for FSCIL, can be misleading by the base performance.
- Training including an extractor  $f$  (M5 in Figure 2) significantly deteriorates the base performance, which is in line with Jie et al. (2022). Moreover, this training scheme has a similar novel performance to the models with better base performance (M3 and M4 in Figure 2). Therefore, simply updating  $f$  is not an appealing strategy.
- Training only the current classifier  $h^i$  (M2 in Figure 2) is the best strategy for improving the novel performance; however, this training scheme degrades the base performance even with the fixed extractor  $f$ .

1. The same results are observed on CUB200 and miniImageNet, reported in Appendix B.

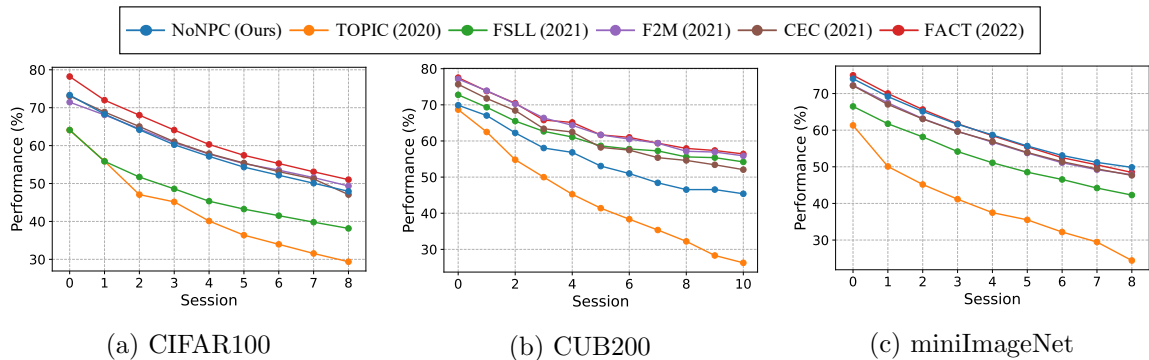


Figure 4: Weighted performance comparison. CEC (Zhang et al., 2021) and FacT (Zhou et al., 2022) are reproduced and other algorithms are from Shi et al. (2021). The base and novel performances comparisons are reported in Appendix D.

We further investigate why the third observation occurs. Under M2, after the base session, every logit on base classes of samples  $x$  belonging to the base classes are fixed to  $h^0 f(x)$  because  $f$  and  $h^0$  are frozen. However, the predicted probability  $p_j$  corresponding to the base class  $j$  strictly decreases as novel classes increase as follows:

$$p_j = \frac{\exp(z_j)}{\sum_{k=1}^K \exp(z_k)} \rightarrow \frac{\exp(z_j)}{\sum_{k=1}^{K+nS} \exp(z_k)} \quad (1)$$

In this situation, we describe the logit distributions of base samples, depicted in Figure 3. Training only  $h^i$  leads the logit values of incremental novel classes to exceed the top logit value of base classes, regardless of which class comes in. The logit values for the novel classes are large overall by making the norm of  $h_k^i$  ( $i > 0$  and  $k \in \{1, \dots, n\}$ ) larger than the norms of row vectors of  $h^0$ . The logit distribution of base samples on other datasets are provided in Appendix C.

## 5. NoNPC: No Training with Normalized Prototype Classifiers

We propose a simple yet powerful method called **NoNPC**, which means **No** training with **Normalized Prototype Classifiers**, inspired by observations in Section 4: combining (1) M1 to maintain the base performance and (2) M2 to improve the novel performance. Note that M1 does not work on the novel classes at all, while M2 deteriorates the base performance. To solve this conundrum, we use non-parametric prototype classifiers without training. The prototype  $c_j^i$  for the novel class  $j$  in the session  $i$  ( $i > 0$ ) are defined as:

$$c_j^i = \frac{1}{k} \sum_{x \text{ belongs to class } j} f(x) \quad (2)$$

However,  $c_j^i$  is unstable because  $k$  is significantly small (e.g., 1 or 5). To mitigate this issue, we transform prototype  $c_j^i$  through L2-normalization, following Wang et al. (2019), in which it is shown that feature transformation improves performance. Finally, we use the normalized prototype as a classifier, i.e.,  $h_j^i = \frac{c_j^i}{\|c_j^i\|_2}$ . Furthermore, to match the character-

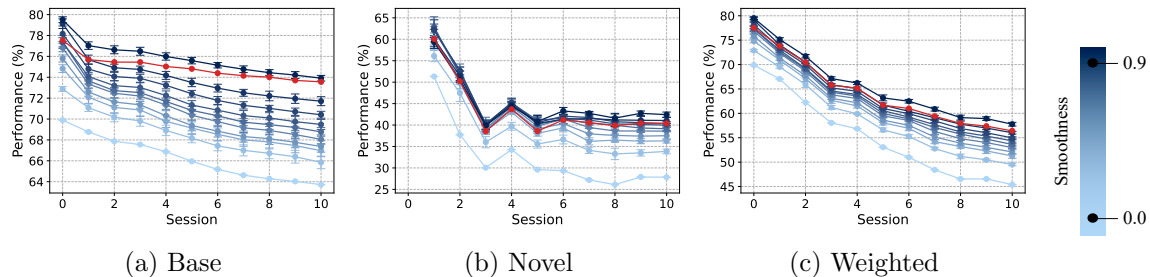


Figure 5: Base, novel, and weighted performances according to the degree of label smoothing on CUB200. The most light and dark blue colors indicate smoothness of 0.0 and 0.9, respectively. The value of 0.0 means label smoothing is not used and the value increases by 0.1 from 0.0 to 0.9. The red line indicates the performance of FACT (Zhou et al., 2022).

istics between  $h^0$  and  $h^i$  ( $i > 0$ ), we replace the base classifier  $h^0$  optimized via stochastic gradient descent with a normalized prototype classifier in the same way after the base session. The ablation studies related to prototype normalization are described in Appendix E.

Figure 4 describes the weighted performance comparison on CIFAR100, CUB200, and miniImageNet. It is observed that our simple method has the best performance on miniImageNet and comparable performance to recent algorithms on CIFAR100. However, on CUB dataset, our algorithm does not achieve the desired performance.

### 5.1 Label Smoothing on Fine-grained Dataset

CUB is one of the fine-grained datasets, which implies that the distance in the representation spaces between different classes can be closer. However, for both base and novel performances, it seems crucial that representations within the same class are clustered tightly and representations between different classes maintain the distance, before the novel classes increase. Therefore, we believe that label smoothing is appropriate for fine-grained datasets during the base session. This is because label smoothing helps tight clustering, making equidistance between different classes (Müller et al., 2019).

Figure 5 describes performances according to the degree of label smoothing on CUB200. Surprisingly, both base and novel performances consistently increase until the smoothness reaches 0.9. In addition, this outperforms FACT (Zhou et al., 2022). The results on other datasets are reported in Appendix F.

## 6. Conclusion

In this paper, we decomposed the entire network into four partial parameters to investigate the relationship between the update parts and the two performances. Based on observations, we proposed a simple yet powerful method called NoNPC, which does not learn a model for the novel sessions and only makes inferences with normalized prototype classifiers. This straightforward method achieved comparable performance with the state-of-the-art algorithms. Furthermore, we showed that the label smoothing technique on the fine-grained dataset boosts the performance. We hope that our NoNPC will be used as a baseline in future studies on FSCIL.

## Appendix A. Implementation Details

### A.1 Dataset Details

Following the benchmark setting for FSCIL (Tao et al., 2020), we evaluate the base, novel, and weighted performances on CIFAR100 (Krizhevsky et al., 2009), CUB200-2011 (Wah et al., 2011), and miniImageNet (Russakovsky et al., 2015). The number of novel classes is the product of the number of incremental novel classes and novel sessions.

Dataset	# of Base classes	# of incremental novel classes	# of novel sessions
CIFAR100	60	5	8
CUB200	100	10	10
miniImageNet	60	5	8

Table 2: Class incremental setup.

### A.2 Training Details

Our algorithm requires the base training only.

Configuration	
Models	ResNet20 (for CIFAR100) and ResNet18 (for CUB200 and miniImageNet)
Epochs	200
Optimizer	SGD with momentum 0.9 (nesterov=True)
Batch size	256
Learning rate	0.1 with milestone scheduler at 120 and 160 epochs (gamma: 0.1)
Weight decay	5e-4 (for CIFAR100 and miniImageNet) and 5e-5 (for CUB200)

Table 3: Training setup.

## Appendix B. Performance According to Decomposition

Figure 6 and 7 describe the base, novel, and weighted performances according to the decomposition on CUB200 and miniImageNet, respectively. Models are described in Table 1.

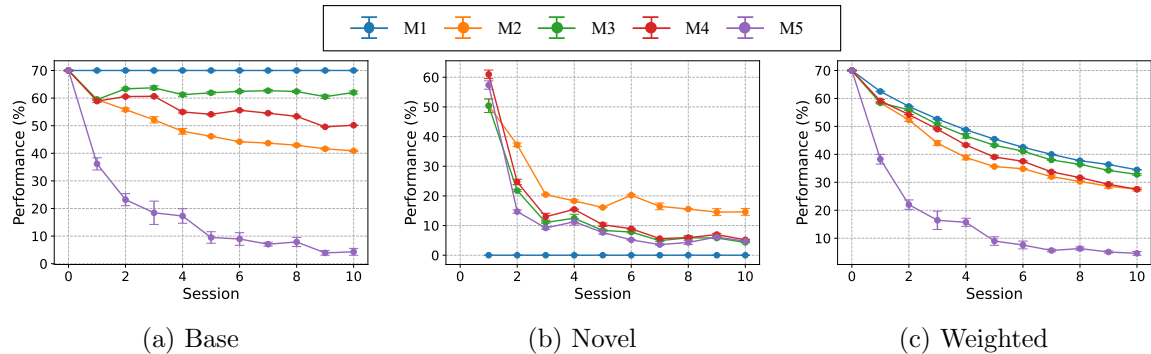


Figure 6: Performances according to the update parts on CUB200.

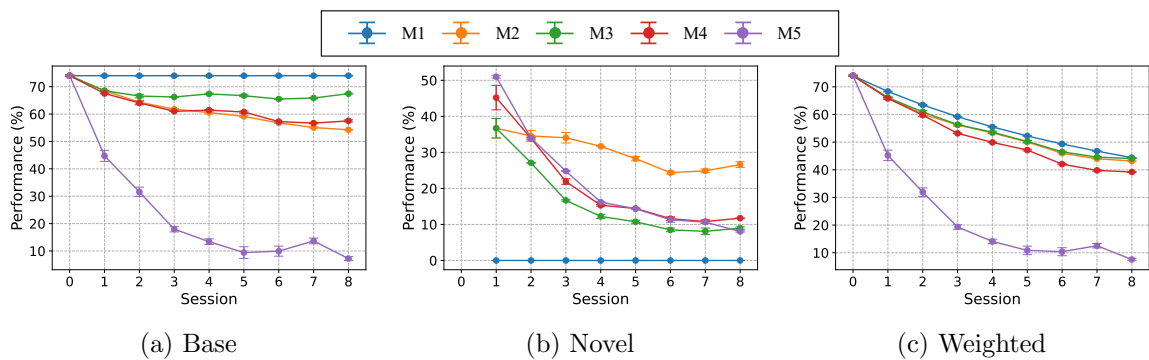


Figure 7: Performances according to the update parts on miniImageNet.

### Appendix C. Logit Distributions

Figure 8 and 9 describe the logit distributions of base samples when training only the current classifier  $h^i$  on CUB200 and miniImageNet, respectively.

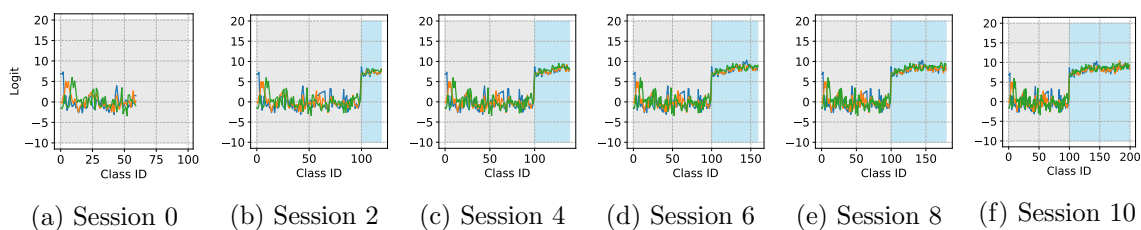


Figure 8: Logit distributions of base samples according to the session on CUB200 when training only  $h^i$ . Gray and blue backgrounds indicate base and novel classes, respectively. Each line is the average of logits of samples belonging to class 0, 1, and 2.

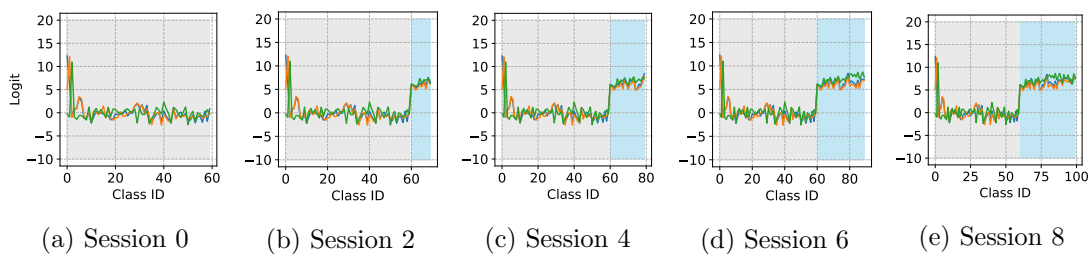


Figure 9: Logit distributions of base samples according to the session on miniImageNet when training only  $h^i$ . Gray and blue backgrounds indicate base and novel classes, respectively. Each line is the average of logits of samples belonging to class 0, 1, and 2.



## Appendix D. Base, Novel, and Weighted Performances Comparison

Table 10, 11, and 12 describe the base, novel, and weighted performances on CIFAR100, CUB200, and miniImageNet, respectively.

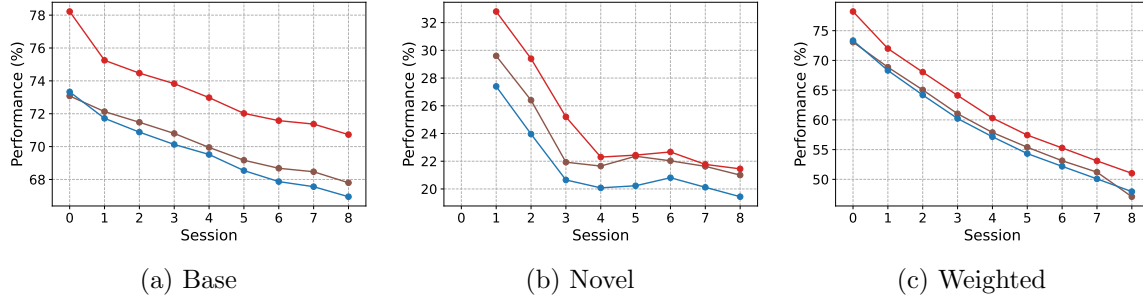


Figure 10: Performances comparison on CIFAR100. The blue, brown, and red lines indicate NoNPC (ours), CEC (Zhang et al., 2021), FACT (Zhou et al., 2022).

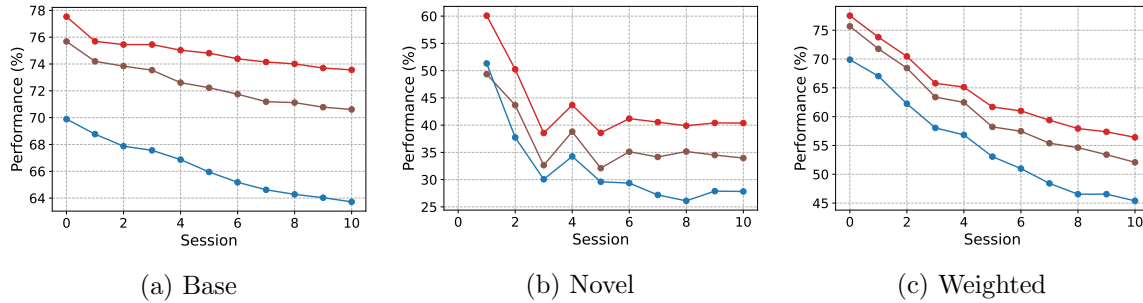


Figure 11: Performances comparison on CUB200. The blue, brown, and red lines indicate NoNPC (ours), CEC (Zhang et al., 2021), FACT (Zhou et al., 2022).

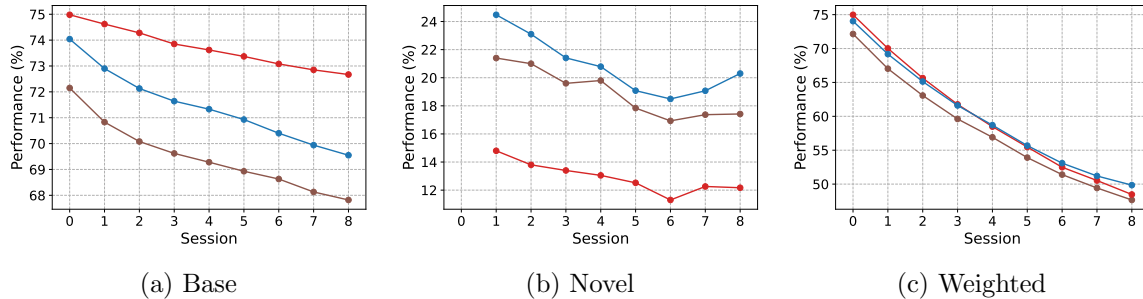


Figure 12: Performances comparison on miniImageNet. The blue, brown, and red lines indicate NoNPC (ours), CEC (Zhang et al., 2021), FACT (Zhou et al., 2022).

## Appendix E. Prototype Normalization

In this section, for ablation studies, we compare below models:

Table 4: Models for ablation studies related to prototype normalization.

Model	Base classifier	Novel classifier	Model	Base classifier	Novel classifier
M1	-	-	M3	-	NP
M2	-	P	M4 (NoNPC)	NP	NP

where ‘-’ of a base classifier and a novel classifier means the optimized classifier via stochastic gradient descent and a random initialized classifier, respectively. ‘P’ and ‘NP’ indicate a prototype classifier and a normalized prototype classifier. Note that M1 in Table 4 is the same with M1 in Table 1, while other models have nothing to do with models in Table 1.

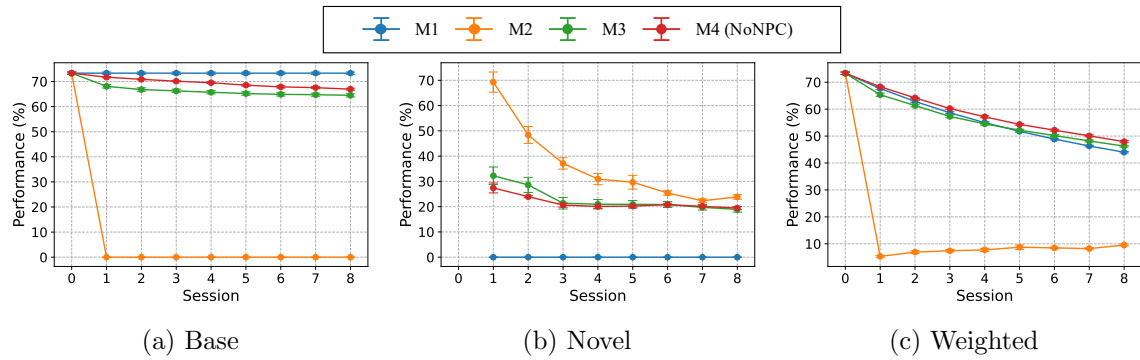


Figure 13: Ablation study according to prototype normalization on CIFAR100.

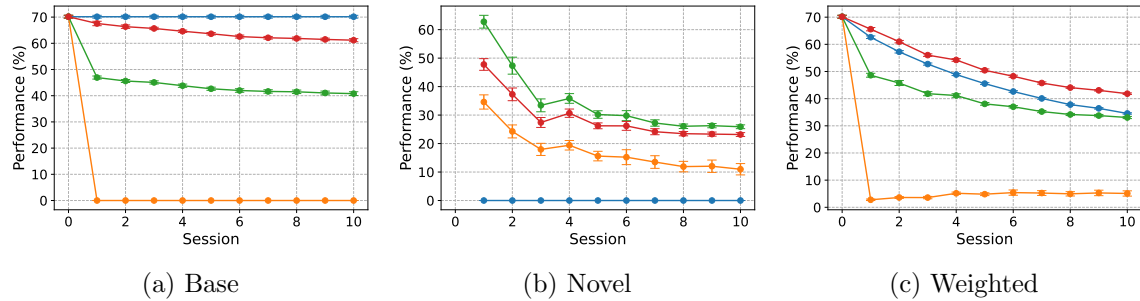


Figure 14: Ablation study according to prototype normalization on CUB200.

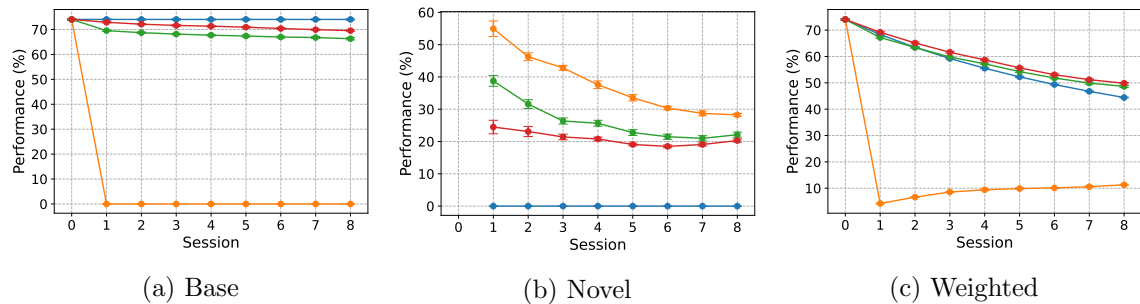


Figure 15: Ablation study according to prototype normalization on miniImageNet.

## Appendix F. Label Smoothing during the Base Session

Figure 16 and 17 describe the base, novel, and weighted performances according to the degree of label smoothing on CIFAR100 and miniImageNet, respectively. The tendency for the novel classes is opposite to the tendency on CUB200. For coarse-grained datasets, small smoothness can increase the novel performance rather than large smoothness.

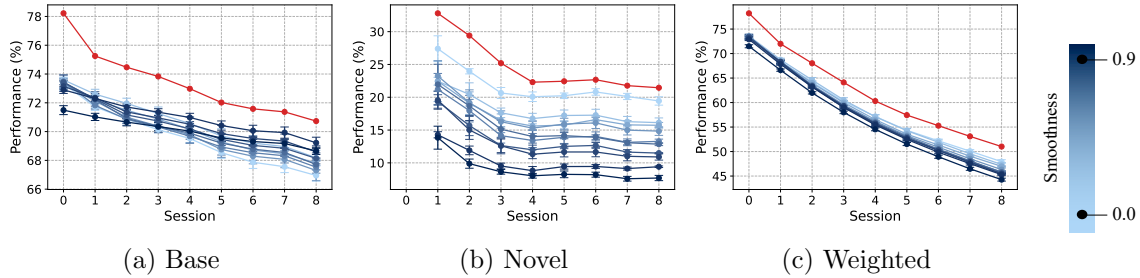


Figure 16: Base, novel, and weighted performances according to the degree of label smoothing on CIFAR100. The most light and dark blue colors indicate smoothness of 0.0 and 0.9, respectively. The value of 0.0 means label smoothing is not used. The red line indicates the performance of FACT (Zhou et al., 2022).

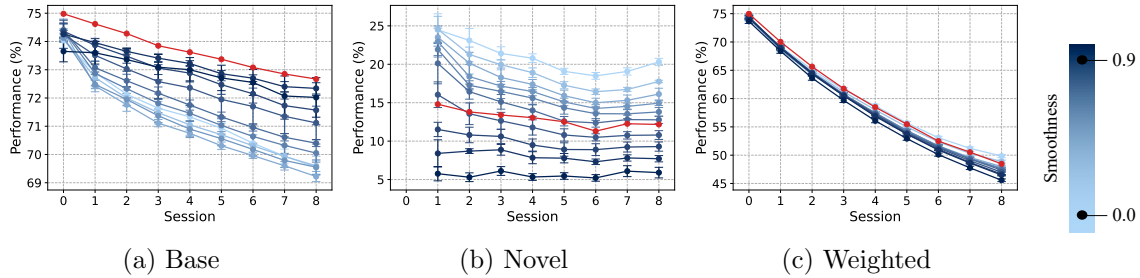


Figure 17: Base, novel, and weighted performances according to the degree of label smoothing on miniImageNet. The most light and dark blue colors indicate smoothness of 0.0 and 0.9, respectively. The value of 0.0 means label smoothing is not used. The red line indicates the performance of FACT (Zhou et al., 2022).

## References

- Manoj Ghuhana Arivazhagan, Vinay Aggarwal, Aaditya Kumar Singh, and Sunav Choudhary. Federated learning with personalization layers. *arXiv preprint arXiv:1912.00818*, 2019.
- Eden Belouadah and Adrian Popescu. Il2m: Class incremental learning with dual memory. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- Liam Collins, Hamed Hassani, Aryan Mokhtari, and Sanjay Shakkottai. Exploiting shared representations for personalized federated learning. In *International Conference on Machine Learning*, pages 2089–2099. PMLR, 2021.
- MohammadReza Davari, Nader Asadi, Sudhir Mudur, Rahaf Aljundi, and Eugene Belilovsky. Probing representation forgetting in supervised and unsupervised continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- Enrico Fini, Stéphane Lathuiliere, Enver Sangineto, Moin Nabi, and Elisa Ricci. Online continual learning under extreme memory constraints. In *European Conference on Computer Vision*, pages 720–735. Springer, 2020.
- Michael Hersche, Geethan Karunaratne, Giovanni Cherubini, Luca Benini, Abu Sebastian, and Abbas Rahimi. Constrained few-shot class-incremental learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- Shibo Jie, Zhi-Hong Deng, and Ziheng Li. Alleviating representational shift for continual fine-tuning. *arXiv preprint arXiv:2204.10535*, 2022.
- KJ Joseph, Salman Khan, Fahad Shahbaz Khan, Rao Muhammad Anwer, and Vineeth N Balasubramanian. Energy-based latent aligner for incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=r1gRTCvFvB>.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Yoonho Lee and Seungjin Choi. Gradient-based meta-learning with learned layerwise metric and subspace. In *International Conference on Machine Learning*, pages 2927–2936. PMLR, 2018.
- Zheda Mai, Ruiwen Li, Jihwan Jeong, David Quispe, Hyunwoo Kim, and Scott Sanner. Online continual learning in image classification: An empirical survey. *Neurocomputing*, 469:28–51, 2022.

- Marc Masana, Xialei Liu, Bartłomiej Twardowski, Mikel Menta, Andrew D Bagdanov, and Joost van de Weijer. Class-incremental learning: survey and performance evaluation on image classification. *arXiv preprint arXiv:2010.15277*, 2020.
- Pratik Mazumder, Pravendra Singh, and Piyush Rai. Few-shot lifelong learning. 2021.
- Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? *Advances in neural information processing systems*, 32, 2019.
- Jaehoon Oh, Hyungjun Yoo, ChangHwan Kim, and Se-Young Yun. Boil: Towards representation change for few-shot learning. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=umIdUL8rMH>.
- Jaehoon Oh, SangMook Kim, and Se-Young Yun. FedBABU: Toward enhanced representation for federated image classification. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=HuaYQfggn5u>.
- Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. Rapid learning or feature reuse? towards understanding the effectiveness of maml. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rkgMkCEtPB>.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- Guangyuan Shi, Jiaxin Chen, Wenlong Zhang, Li-Ming Zhan, and Xiao-Ming Wu. Overcoming catastrophic forgetting in incremental few-shot learning by finding flat minima. In *Advances in Neural Information Processing Systems*, volume 34, 2021.
- Dongsub Shim, Zheda Mai, Jihwan Jeong, Scott Sanner, Hyunwoo Kim, and Jongseong Jang. Online class-incremental continual learning with adversarial shapley value. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 9630–9638, 2021.
- Xiaoyu Tao, Xiaopeng Hong, Xinyuan Chang, Songlin Dong, Xing Wei, and Yihong Gong. Few-shot class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.
- Yan Wang, Wei-Lun Chao, Kilian Q Weinberger, and Laurens van der Maaten. Simpleshot: Revisiting nearest-neighbor classification for few-shot learning. *arXiv preprint arXiv:1911.04623*, 2019.
- Haiyang Yu, Ningyu Zhang, Shumin Deng, Zonggang Yuan, Yantao Jia, and Huajun Chen. The devil is the classifier: Investigating long tail relation classification with decoupling analysis. *arXiv preprint arXiv:2009.07022*, 2020.

Chi Zhang, Nan Song, Guosheng Lin, Yun Zheng, Pan Pan, and Yinghui Xu. Few-shot incremental learning with continually evolved classifiers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12455–12464, June 2021.

Da-Wei Zhou, Fu-Yun Wang, Han-Jia Ye, Liang Ma, Shiliang Pu, and De-Chuan Zhan. Forward compatible few-shot class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.

Fei Zhu, Zhen Cheng, Xu-yao Zhang, and Cheng-lin Liu. Class-incremental learning via dual augmentation. *Advances in Neural Information Processing Systems*, 34, 2021.