

Heterogeneous Forgetting Compensation for Class-Incremental Learning

Jiahua Dong^{1, 2, 3*}, Wenqi Liang^{1, 2, 3*}, Yang Cong^{4†}, Gan Sun^{1, 2}

¹State Key Laboratory of Robotics, Shenyang Institute of Automation,
Chinese Academy of Sciences, Shenyang, 110016, China.[‡]

²Institutes for Robotics and Intelligent Manufacturing,
Chinese Academy of Sciences, Shenyang, 110169, China.

³University of Chinese Academy of Sciences, Beijing, 100049, China.

⁴South China University of Technology, Guangzhou, 510640, China.

{dongjiahua1995, liangwenqi0123, congyang81, sungan1412}@gmail.com

Abstract

Class-incremental learning (CIL) has achieved remarkable successes in learning new classes consecutively while overcoming catastrophic forgetting on old categories. However, most existing CIL methods unreasonably assume that all old categories have the same forgetting pace, and neglect negative influence of forgetting heterogeneity among different old classes on forgetting compensation. To surmount the above challenges, we develop a novel **H**eterogeneous **F**orgetting **C**ompensation (**HFC**) model, which can resolve heterogeneous forgetting of easy-to-forget and hard-to-forget old categories from both representation and gradient aspects. Specifically, we design a task-semantic aggregation block to alleviate heterogeneous forgetting from representation aspect. It aggregates local category information within each task to learn task-shared global representations. Moreover, we develop two novel plug-and-play losses: a gradient-balanced forgetting compensation loss and a gradient-balanced relation distillation loss to alleviate forgetting from gradient aspect. They consider gradient-balanced compensation to rectify forgetting heterogeneity of old categories and heterogeneous relation consistency. Experiments on several representative datasets illustrate effectiveness of our HFC model. The code is available at <https://github.com/JiahuaDong/HFC>.

1. Introduction

Class-incremental learning (CIL) [40, 12, 5] has attracted appealing attentions recently by accumulating previous learned experience to learn new classes incrementally. It

*Equal contributions. †The corresponding author is Prof. Yang Cong.

‡This work was supported in part by the National Nature Science Foundation of China under Grant 62127807, 62273333 and 62133005.




Figure 1. Forgetting heterogeneity versus accuracy on CIFAR-100 [33] when the backbone is ViT-Base and number of tasks is 5.

plays an indispensable role in developing a large number of intelligent learning systems, such as autonomous driving [38] and automated surveillance [26]. When learning new classes continuously under the settings of limited memory to replay all previous data of old classes [45], these CIL systems heavily suffer from forgetting on old classes caused by severe class imbalance between old and new categories [40, 31, 14]. To surmount catastrophic forgetting, a growing amount of CIL methods mainly perform knowledge distillation [24] to preserve past experience [34, 12, 35]; introduce a rehearsal strategy to replay part of old data [40, 17]; or dynamically expand network architectures [28, 41, 54].

However, most existing CIL methods [40, 45, 54, 18] neglect heterogeneous forgetting speeds of different old classes. They unrealistically assume that all old classes suffer from the same degree of catastrophic forgetting, and compensate forgetting for each old classes equally and independently. Such impracticable assumption enforces existing CIL models [40, 56, 61] to suffer from imbalanced gradient optimization among different old classes, thus favoring more forgetting compensation for hard-to-forget old classes while ne-

glecting those easy-to-forget old classes (*i.e.*, heterogeneous forgetting). More importantly, such forgetting heterogeneity can significantly worsen the forgetting on hard-to-forget old categories, when new streaming classes becomes part of old categories continually. For example, some old classes with various modalities and appearances (*e.g.*, `person` and `car`) in autonomous driving [38] are more difficult to explore task-shared representations across different incremental tasks, when compared with other hard-to-forget old classes (*e.g.*, `road` and `traffic sign`) with easily-distinguished visual properties. This phenomenon causes imbalanced gradient propagation between easy-to-forget and hard-to-forget old categories, thus exacerbating heterogeneous forgetting speeds among different old classes. When autonomous vehicles [38] learn new classes continually, such forgetting heterogeneity aggravates forgetting on those hard-to-forget old classes (*e.g.*, `traffic sign`) to some extent.

Inspired by the above practical scenario, we investigate that different old classes have significant forgetting heterogeneity in this paper, as shown in Fig. 1. This heterogeneous forgetting might heavily weaken forgetting compensation on hard-to-forget old classes, as new classes become a subset of old classes consecutively. In summary, the challenges to tackle heterogeneous forgetting lie in two major aspects:

- **Representation Aspect:** Some easy-to-forget old classes (*e.g.*, `car` and `bus`) with diverse appearances are more difficult for existing CIL methods [45, 54, 12, 3] to learn intrinsic task-shared representations, and thus are significantly easier to be forgotten than hard-to-forget old classes (*e.g.*, `road`) with distinctive attributes. Thus, exploring task-shared representations is essential to address heterogeneous forgetting among different old classes.
- **Gradient Aspect:** To learn complex visual characterizations for easy-to-forget old classes with various modalities and appearances, existing CIL methods [40, 4, 49] are required to allocate more network architectures for gradient updating. It can result in imbalanced gradient propagation between easy-to-forget and hard-to-forget old categories, thus aggravating forgetting heterogeneity of old classes.

To overcome the above challenges, we propose a novel **Heterogeneous Forgetting Compensation (HFC)** model, which is an earlier exploration to address heterogeneous forgetting from both representation and gradient perspectives in the CIL field [40, 4]. Specifically, we propose a task-semantic aggregation (TSA) block to alleviate heterogeneous forgetting from representation perspective. It can explore task-shared global representations for each class via aggregating long-range local category information within each task. Meanwhile, to tackle heterogeneous forgetting from gradient perspective, we develop two novel plug-and-play losses: a gradient-balanced forgetting compensation (GFC) loss and a gradient-balanced relation distillation (GRD) loss. The GFC loss can rectify heterogeneous forgetting speeds

of easy-to-forget and hard-to-forget old classes, while normalize different learning paces of new categories to achieve gradient-balanced propagation. Besides, the GRD loss can distill heterogeneous relation consistency brought by forgetting heterogeneity among different old classes. Experiments show our model has large improvements on several representative datasets, compared with baseline methods [40, 36]. More importantly, we apply two plug-and-play losses into existing distillation-based CIL methods [40, 17, 4] to significantly improve their performance from gradient aspect. The novel contributions of this paper are presented as follows:

- We develop a novel Heterogeneous Forgetting Compensation (HFC) model to address different forgetting speeds of easy-to-forget and hard-to-forget old classes. To our best knowledge, this paper is the first exploration to tackle forgetting heterogeneity among old categories from representation and gradient aspects in the CIL field.
- We design a task-semantic aggregation (TSA) block to alleviate heterogeneous forgetting from representation aspect. It can explore robust task-shared representations for each class via aggregating local category information.
- We propose two novel plug-and-play losses: a gradient-balanced forgetting compensation (GFC) loss and a gradient-balanced relation distillation (GRD) loss to surmount forgetting heterogeneity from gradient aspect. They can balance different forgetting of old classes and heterogeneous category-relation consistency to improve performance when applying them into existing methods.

2. Related Work

We discuss some class-incremental learning (CIL) methods [60, 37, 18], and divide them into three categories:

Knowledge Distillation: After Li *et al.* [34] apply knowledge distillation [24] to continual learning, [40, 39, 59] follow them to transfer past experience from old model to new model. iCaRL [40] and EEIL [6] perform knowledge distillation in the output space. [44] aims to preserve consistent topology feature space for old and new tasks. [16] proposes feature-level knowledge distillation when applying continual learning [37] into semantic segmentation task [53, 52, 11].

Rehearsal Strategy: For replaying past experience, lots of CIL methods [49, 8, 42] allocate a memory to store exemplars of old classes or synthesize samples via generative adversarial model [21, 58]. For saving large memory overhead, [29] proposes to store low-dimension features rather than raw samples. [51] proposes a “white box” framework deviated from rate reduction to preserve past experience.

Dynamic Architecture: Many works [28, 54, 1] allocate dynamical networks for new classes as the growing of learned classes. Yoon *et al.* [55] dynamically expand architecture capacity via selective retraining. [54, 18] address the problem of relying on task index when performing dynamic

architectures. However, the above-mentioned CIL methods [40, 27] neglect negative influence of heterogeneous forgetting of old classes on forgetting compensation.

3. The Proposed Model

3.1. Problem Definition and Overview

Problem Definition: Following class-incremental learning (CIL) methods [40, 50, 17, 18, 54], we set a series of consecutive learning tasks as $\mathcal{T} = \{\mathcal{T}^t\}_{t=1}^T$, where T indicates the total task quantity. The t -th task $\mathcal{T}^t = \{\mathbf{x}_i^t, \mathbf{y}_i^t\}_{i=1}^{N^t}$ is composed of N^t pairs of image \mathbf{x}_i^t and one-hot groundtruth $\mathbf{y}_i^t \in \mathcal{Y}^t$, and \mathcal{Y}^t denotes label space of the t -th incremental task containing K^t new categories. The label spaces between any two incremental tasks have no overlap: $\mathcal{Y}^t \cap (\cup_{j=1}^{t-1} \mathcal{Y}^j) = \emptyset$. Namely, K^t new classes belonging to the t -th task \mathcal{T}^t are different from $K^o = \sum_{i=1}^{t-1} K^i \subset \cup_{j=1}^{t-1} \mathcal{Y}^j$ old classes learned from previous tasks $\{\mathcal{T}^i\}_{i=1}^{t-1}$. As introduced in CIL baselines [40, 50, 54], we set a fixed exemplar memory \mathcal{M} , and store only few images (*i.e.*, $\frac{|\mathcal{M}|}{K^o}$) for each old class in the t -th task \mathcal{T}^t , where \mathcal{M} satisfies $\frac{N^t}{K^t} \gg \frac{|\mathcal{M}|}{K^o}$. In the t -th task, we aim to identify both K^t new classes and K^o old classes via optimizing the model on \mathcal{T}^t and \mathcal{M} .

Overview: The overview of our HFC model to surmount forgetting heterogeneity of different old categories from both representation and gradient aspects is shown in Fig. 2. Denote the proposed HFC model learned at the $(t-1)$ -th and t -th tasks as old and new models (*i.e.*, Θ^{t-1} and Θ^t), where Θ^t is inherited from Θ^{t-1} and only expands the number of output neurons in the classifier to identify K^t new classes. In the t -th incremental task, given an image $\mathbf{x}_i^t \in \mathcal{T}^t \cup \mathcal{M}$, we forward it into Θ^{t-1} and Θ^t to extract task-shared global representation via task-semantic aggregation (TSA) block, which alleviates heterogeneous forgetting from representation aspect (Section 3.2). After obtaining probabilities $\mathbf{P}^{t-1}(\mathbf{x}_i^t, \Theta^{t-1}) \in \mathbb{R}^{K^o}$ and $\mathbf{P}^t(\mathbf{x}_i^t, \Theta^t) \in \mathbb{R}^{K^o+K^t}$ predicted via Θ^{t-1} and Θ^t , we develop two plug-and-play losses: a gradient-balanced forgetting compensation (GFC) loss \mathcal{L}_{FC} and a gradient-balanced relation distillation (GRD) loss \mathcal{L}_{RD} to surmount heterogeneous forgetting from semantic aspect (Sections 3.3 and 3.4). \mathcal{L}_{FC} and \mathcal{L}_{RD} can normalize heterogeneous forgetting of old classes and distill heterogeneous relation consistency between Θ^{t-1} and Θ^t .

3.2. Task-Semantic Aggregation Block

To explore long-range semantic dependencies across different tasks, Douillard *et al.* [18] introduce Vision Transformer (ViT) [15, 46, 22] into class-incremental learning (CIL). However, it significantly increases the memory overhead to store task tokens and task-specific decoders, when learning large-scale new classes. Besides, it cannot aggregate local category information to explore task-shared global

representations, resulting in heterogeneous forgetting from representation aspect. To tackle these issues, we develop a task-semantic aggregation (TSA) block to extract task-shared global representations, which is effective to tackle heterogeneous forgetting from representation aspect. More importantly, the TSA block cannot dynamically increase memory cost with the arrival of large-scale new categories.

Specifically, as presented in Fig. 2, the proposed model $\Theta^t = \Theta_c^t \circ \Theta_f^t$ at the t -th task is composed of a feature extractor Θ_f^t to extract task-shared global representations, and a task classifier Θ_c^t to classify K^o old classes and K^t new classes. In the t -th task, the network parameters of Θ_f^{t-1} and Θ_c^{t-1} in Θ^{t-1} learned at the $(t-1)$ -th task are frozen to preserve semantic knowledge for old classes. Different from [18] that uses vanilla ViT [15, 30, 57, 7] to extract features, we employ l_m multi-head self-attention (MSA) blocks [20] and l_t task-semantic aggregation (TSA) block as feature extractor Θ_f^t to obtain task-shared global representations.

• **Multi-Head Self-Attention (MSA):** Given an image $\mathbf{x}_i^t \in \mathcal{T}^t \cup \mathcal{M}$ in the t -th task, following ViT [15, 9], we crop it into N patches with equal dimension, flatten these patches and map them into a D -dimension feature space via a linear projection to obtain patch embedding $\mathbf{Z}_e \in \mathbb{R}^{N \times D}$. Then we concatenate \mathbf{Z}_e with a class token $\mathbf{Z}_{cls} \in \mathbb{R}^D$, and perform element-wise sum operation with a position token $\mathbf{Z}_p \in \mathbb{R}^{(N+1) \times D}$ to get $\mathbf{Z}_0 = [\mathbf{Z}_e \oplus \mathbf{Z}_{cls}] + \mathbf{Z}_p \in \mathbb{R}^{(N+1) \times D}$, where \oplus is concatenation function. As shown in Fig. 2, \mathbf{Z}_0 is fed into Θ_f^t including l_m MSA blocks to extract long-range semantics across different tasks. For the l -th ($l = 1, \dots, l_m$) MSA block, we regard $\mathbf{Z}_{l-1} \in \mathbb{R}^{(N+1) \times D}$ learned from the $(l-1)$ -th MSA block as input, and execute parallel self-attention H times. Thus, the self-attention $\mathbf{A}_l^h \in \mathbb{R}^{(N+1) \times d}$ of \mathbf{Z}_{l-1} for the h -th ($h = 1, \dots, H$) head is:

$$\mathbf{A}_l^h = \sigma\left(\frac{\mathbf{Z}_{l-1} \mathbf{W}_q^h (\mathbf{Z}_{l-1} \mathbf{W}_k^h)^\top}{\sqrt{d}}\right) (\mathbf{Z}_{l-1} \mathbf{W}_v^h), \quad (1)$$

where $\mathbf{W}_q^h, \mathbf{W}_k^h, \mathbf{W}_v^h \in \mathbb{R}^{D \times d}$ are query, key and value mapping matrices. $d = \frac{D}{H}$ represents channel dimension of each head and σ is the softmax function. Then we concatenate $\{\mathbf{A}_l^h\}_{h=1}^H$ along channel dimension and project this concatenated result via a mapping matrix $\mathbf{W}_o \in \mathbb{R}^{D \times D}$ to perform MSA: $\text{MSA}(\mathbf{Z}_{l-1}) = \mathbf{Z}_{l-1} + [\mathbf{A}_l^1 \oplus \mathbf{A}_l^2, \dots, \oplus \mathbf{A}_l^H] \mathbf{W}_o$, where \oplus is concatenation function. After obtaining $\text{MSA}(\mathbf{Z}_{l-1}) \in \mathbb{R}^{(N+1) \times D}$ ($D = dH$), we utilize a multi-layer perceptron (MLP) block to obtain output $\mathbf{Z}_l \in \mathbb{R}^{(N+1) \times D}$ of the l -th MSA block:

$$\mathbf{Z}_l = \text{MSA}(\mathbf{Z}_{l-1}) + \text{MLP}(\text{MSA}(\mathbf{Z}_{l-1})). \quad (2)$$

• **Task-Semantic Aggregation (TSA):** The patch embedding $\mathbf{Z}_{l_m} \in \mathbb{R}^{(N+1) \times D}$ encoded via the l_m -th MSA block contains rich long-range semantic dependencies shared

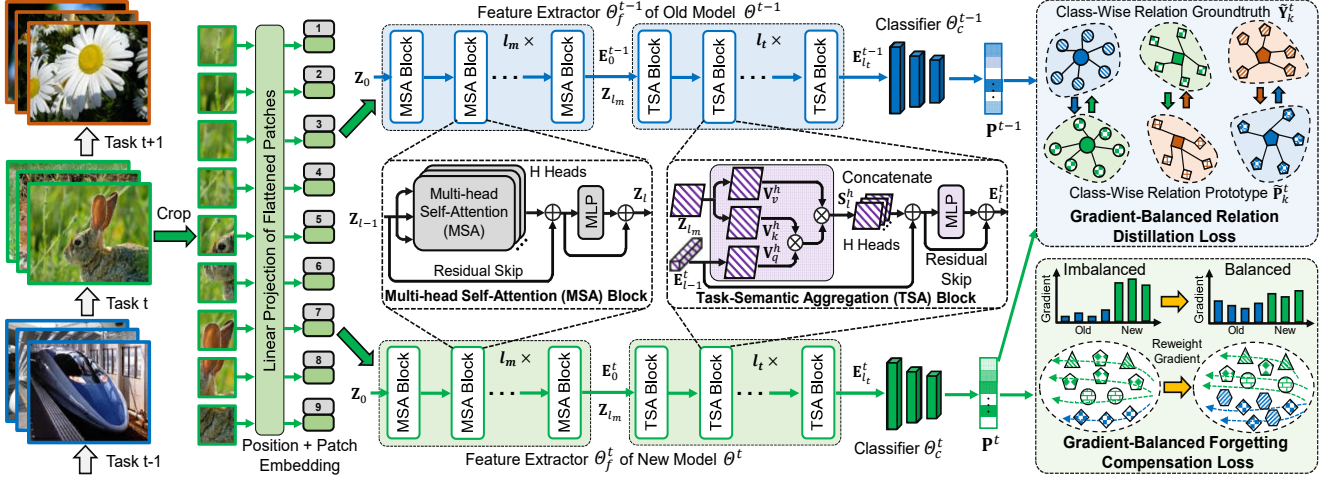


Figure 2. Illustration of the HFC model. It mainly contains a *task-semantic aggregation (TSA)* block to alleviate heterogeneous forgetting from representation aspect via exploring task-shared global representations, and two novel plug-and-play losses (i.e., a *gradient-balanced forgetting compensation loss* \mathcal{L}_{FC} and a *gradient-balanced relation distillation loss* \mathcal{L}_{RD}) to overcome forgetting heterogeneity from gradient aspect by compensating gradient-imbalanced propagation and heterogeneous class-relation consistency.

across different tasks. \mathbf{Z}_{l_m} is then forwarded into l_t task-semantic aggregation (TSA) blocks to extract task-shared global embedding. The TSA block aggregates local category context from long-range semantic dependencies to global representations shared across all incremental tasks, which is essential to alleviate heterogeneous forgetting from representation aspect. In the t -th task, as depicted in Fig. 2, we introduce a learnable task-shared embedding $\mathbf{E}_0^t \in \mathbb{R}^D$ that is initialized via \mathbf{E}_0^{t-1} learned in the $(t-1)$ -th task. Then we forward \mathbf{E}_0^t along with \mathbf{Z}_{l_m} to l_t TSA blocks. For the l -th ($l = 1, \dots, l_t$) TSA block, $\mathbf{E}_{l-1}^t \in \mathbb{R}^D$ learned at $(l-1)$ -th TSA block and \mathbf{Z}_{l_m} are used to perform task-semantic attention H times in parallel. Thus, we formulate task-semantic attention $\mathbf{S}_l^h \in \mathbb{R}^d$ for the h -th ($h = 1, \dots, H$) head as:

$$\mathbf{S}_l^h = \sigma\left(\frac{\mathbf{E}_{l-1}^t \mathbf{V}_q^h (\mathbf{Z}_{l_m} \mathbf{V}_k^h)^\top}{\sqrt{d}}\right) (\mathbf{Z}_{l_m} \mathbf{V}_v^h), \quad (3)$$

where $\mathbf{V}_q^h, \mathbf{V}_k^h, \mathbf{V}_v^h \in \mathbb{R}^{D \times d}$ are projection matrices of query, key and value. After utilizing Eq. (3) to obtain $\{\mathbf{S}_l^h\}_{h=1}^H$, we concatenate them along channel dimension and project this fused result via $\mathbf{V}_o \in \mathbb{R}^{D \times D}$ to execute H -head task-semantic attention: $\text{TSA}(\mathbf{E}_{l-1}^t) = [\mathbf{S}_l^1 \oplus \mathbf{S}_l^2, \dots, \oplus \mathbf{S}_l^H] \mathbf{V}_o \in \mathbb{R}^D$. $\text{TSA}(\mathbf{E}_{l-1}^t)$ is encoded via a MLP block to get output $\mathbf{E}_l^t \in \mathbb{R}^D$ of the l -th TSA block:

$$\mathbf{E}_l^t = \text{TSA}(\mathbf{E}_{l-1}^t) + \text{MLP}(\text{TSA}(\mathbf{E}_{l-1}^t)). \quad (4)$$

As aforementioned, we obtain the task-shared global representation $\mathbf{E}_{l_t}^t \in \mathbb{R}^D$ via the l_t -th TSA block in feature extractor Θ_f^t . Obviously, the proposed TSA blocks can aggregate local category context of the t -th task from long-range semantic dependencies to task-shared global representation $\mathbf{E}_{l_t}^t$. It is essential to alleviate heterogeneous forgetting among different old classes from representation aspect, by

exploring task-shared semantic context across different tasks. Note that we employ layer norm before TSA, MSA and MLP blocks, while we omit it for simplicity in this paper.

• **Task Classifier:** The task-shared global representation $\mathbf{E}_{l_t}^t$ obtained via Θ_f^t is fed into task classifier Θ_c^t to predict K^o old classes and K^t new classes. Given a mini-batch $\{\mathbf{x}_i^t, \mathbf{y}_i^t\}_{i=1}^b \in \mathcal{T}^t \cup \mathcal{M}$, the classification loss \mathcal{L}_{CE} is:

$$\mathcal{L}_{CE} = \frac{1}{b} \sum_{i=1}^b \mathcal{D}_{CE}(\mathbf{P}^t(\mathbf{x}_i^t, \Theta^t), \mathbf{y}_i^t), \quad (5)$$

where $\mathbf{P}^t(\mathbf{x}_i^t, \Theta^t) \in \mathbb{R}^{K^o + K^t}$ is softmax probability predicted by Θ^t in the t -th learning task. $\mathcal{D}_{CE}(\cdot, \cdot)$ indicates traditional cross-entropy loss, and b is the batch size.

However, the severe class imbalance (i.e., $\frac{N^t}{K^t} \gg \frac{|M|}{K^o}$) among old and new categories in the t -th task enforces the prediction $\mathbf{P}^t(\mathbf{x}_i^t, \Theta^t)$ in Eq. (5) to suffer from catastrophic forgetting [40] on old categories. Moreover, the gradient optimization of Eq. (5) neglects heterogeneous forgetting speeds of easy-to-forget old classes with various appearances and hard-to-forget old classes with easily-distinguished attributes. To address heterogeneous forgetting from gradient aspect, we propose two novel plug-and-play losses: a gradient-balanced forgetting compensation (GFC) loss \mathcal{L}_{FC} in Section 3.3 and a gradient-balanced relation distillation (GRD) loss \mathcal{L}_{RD} in Section 3.4.

3.3. Gradient-Balanced Forgetting Compensation

As incremental tasks arrive consecutively, easy-to-forget and hard-to-forget old categories change dynamically, thus aggravating difficulty to tackle heterogeneous forgetting from gradient aspect. In light of this, we develop a gradient-balanced forgetting compensation (GFC) loss \mathcal{L}_{FC} . It adaptively balances large forgetting heterogeneity of old classes

via performing balanced gradient propagation. Specifically, we perform task-adaptive gradient normalization for different classes, and then reweight classification loss \mathcal{L}_{CE} in Eq. (5) to compensate imbalanced gradients. For the t -th incremental task, as claimed in [48, 13], the gradient Γ_i^t of an given image $(\mathbf{x}_i^t, \mathbf{y}_i^t) \in \mathcal{T}^t \cup \mathcal{M}$ with respect to the k -th ($k = \arg \max \mathbf{y}_i^t$) neuron \mathcal{N}_k^t of classifier Θ_c^t in Θ^t is:

$$\Gamma_i^t = \frac{\partial \mathcal{D}_{\text{CE}}(\mathbf{P}^t(\mathbf{x}_i^t, \Theta^t), \mathbf{y}_i^t)}{\partial \mathcal{N}_k^t} = \mathbf{P}^t(\mathbf{x}_i^t, \Theta^t)_k - 1, \quad (6)$$

where $\mathbf{P}^t(\mathbf{x}_i^t, \Theta^t)_k$ denotes probability of the k -th class.

To compensate heterogeneous forgetting of easy-to-forget and hard-to-forget old categories, we compute task-adaptive gradient means for different tasks. Consequently, given a mini-batch $\{\mathbf{x}_i^t, \mathbf{y}_i^t\}_{i=1}^b$ in the t -th learning task, the task-adaptive gradient mean Γ_η for the categories learned in the η -th ($1 \leq \eta \leq t$) incremental task is formulated as:

$$\Gamma_\eta = \frac{1}{\sum_{i=1}^b \mathbb{I}_{\mathbf{y}_i^t \in \mathcal{Y}^\eta}} \sum_{i=1}^b |\Gamma_i^t| \cdot \mathbb{I}_{\mathbf{y}_i^t \in \mathcal{Y}^\eta}. \quad (7)$$

The forgetting heterogeneity of old categories learned from $t-1$ old tasks and learning speeds of new categories in the t -th task can be effectively measured via $\{\Gamma_\eta\}_{\eta=1}^t$ in Eq. (7).

When easy-to-forget and hard-to-forget old classes are varied dynamically as incremental tasks, some noisy predictions on easy-to-forget old classes may result in large deviation to measure their forgetting heterogeneity [43, 19]. To tackle this issue, $\{\Gamma_\eta\}_{\eta=1}^t$ in Eq. (7) are expected to be sharper and more distinguishable adaptively [13], as learning tasks arrive continually. Thus, we rewrite Eq. (7) as follows:

$$\Gamma_\eta^s = \frac{\sum_{i=1}^b \log(|\Gamma_i^t|^{\frac{K^o}{K^o+K^t}} + 1) \cdot \mathbb{I}_{\mathbf{y}_i^t \in \mathcal{Y}^\eta}}{\sum_{i=1}^b \mathbb{I}_{\mathbf{y}_i^t \in \mathcal{Y}^\eta}}. \quad (8)$$

The sharper task-adaptive gradient means $\{\Gamma_\eta^s\}_{\eta=1}^t$ computed via Eq. (8) are employed to reweight classification loss \mathcal{L}_{CE} in Eq. (5). Then the gradient-balanced forgetting compensation (GFC) loss \mathcal{L}_{FC} is written as follows:

$$\mathcal{L}_{\text{FC}} = \frac{1}{b} \sum_{i=1}^b \frac{\log(|\Gamma_i^t|^{\frac{K^o}{K^o+K^t}} + 1)}{\sum_{\eta=1}^t \Gamma_\eta^s \cdot \mathbb{I}_{\mathbf{y}_i^t \in \mathcal{Y}^\eta}} \mathcal{D}_{\text{CE}}(\mathbf{P}^t(\mathbf{x}_i^t, \Theta^t), \mathbf{y}_i^t). \quad (9)$$

Obviously, \mathcal{L}_{FC} in Eq. (9) encourages the optimization of Θ^t to perform gradient-balanced propagation for easy-to-forget, hard-to-forget old classes and new classes learned in different tasks. It tackles heterogeneous forgetting on old classes from gradient aspect, via normalizing forgetting paces of different old classes adaptively with $\{\Gamma_\eta^s\}_{\eta=1}^t$.

3.4. Gradient-Balanced Relation Distillation

The relationships among old and new categories remain constant in semantic space, regardless of availability of old-class samples. In light of this, distilling inter-class relations

from old model Θ^{t-1} to new model Θ^t can address forgetting on old classes [6, 25]. However, most knowledge distillation methods used in CIL [40, 2, 50] only utilize prediction of an individual sample to perform semantic consistency of old classes between Θ^{t-1} and Θ^t . They neglect underlying relations among old and new categories to tackle forgetting, and heavily suffer from noisy predictions on easy-to-forget old classes when distilling knowledge via an individual sample. Moreover, different forgetting speeds of old classes may result in heterogeneous class-relation consistency between Θ^{t-1} and Θ^t . It enforces gradient optimization of Θ^t to bias towards some specific relation distillations related to new classes, causing forgetting heterogeneity on old categories.

To overcome the above issues, we design a gradient-balanced relation distillation (GRD) loss \mathcal{L}_{RD} to distill relations among old and new classes from Θ^{t-1} to Θ^t . As introduced in Fig. 2, it can address heterogeneous forgetting from gradient aspect via rectifying imbalanced gradient propagation caused by heterogeneous class-relation consistency. Specifically, we construct class-wise relation prototype rather than a prediction of single sample to perform relation distillation, which can alleviate negative effect of noisy relation consistency. Then the task-adaptive gradient means $\{\Gamma_\eta^s\}_{\eta=1}^t$ in Eq. (8) are employed to reweight heterogeneous class-relation distillation.

A mini-batch $\{\mathbf{x}_i^t, \mathbf{y}_i^t\}_{i=1}^b \in \mathcal{T}^t \cup \mathcal{M}$ is fed into Θ^{t-1} and Θ^t to obtain softmax probabilities $\mathbf{P}^{t-1}(\mathbf{x}_i^t, \Theta^{t-1}) \in \mathbb{R}^{K^o}$ of old classes and $\mathbf{P}^t(\mathbf{x}_i^t, \Theta^t) \in \mathbb{R}^{K^o+K^t}$ of old and new classes, as shown in Fig. 2. The first K^o dimensions of $\mathbf{y}_i^t \in \mathbb{R}^{K^o+K^t}$ are replaced with $\mathbf{P}^{t-1}(\mathbf{x}_i^t, \Theta^{t-1})$ to get relation groundtruth $\mathbf{Y}^t(\mathbf{x}_i^t, \Theta^{t-1}) \in \mathbb{R}^{K^o+K^t}$ that implies underlying relations among old and new classes. To tackle noisy class relations, we construct class-wise relation prototype $\tilde{\mathbf{P}}_k^t$ and its relation groundtruth $\tilde{\mathbf{Y}}_k^t$ for the k -th class:

$$\tilde{\mathbf{P}}_k^t = \frac{1}{\Delta_k} \sum_{i=1}^b \mathbf{P}^t(\mathbf{x}_i^t, \Theta^t) \cdot \mathbb{I}_{\arg \max \mathbf{y}_i^t = k}, \quad (10)$$

$$\tilde{\mathbf{Y}}_k^t = \frac{1}{\Delta_k} \sum_{i=1}^b \mathbf{Y}^t(\mathbf{x}_i^t, \Theta^{t-1}) \cdot \mathbb{I}_{\arg \max \mathbf{y}_i^t = k}. \quad (11)$$

where $\Delta_k = \sum_i \mathbb{I}_{\arg \max \mathbf{y}_i^t = k}$. Then we formulate category-wise gradient mean Γ_k^s for the k -th category as follows:

$$\Gamma_k^s = \frac{\sum_{i=1}^b \log(|\Gamma_i^t|^{\frac{K^o}{K^o+K^t}} + 1) \cdot \mathbb{I}_{\arg \max \mathbf{y}_i^t = k}}{\sum_{i=1}^b \mathbb{I}_{\arg \max \mathbf{y}_i^t = k}}. \quad (12)$$

Consequently, we use Γ_k^s to reweight heterogeneous relation distillation, and express the proposed \mathcal{L}_{RD} as follows:

$$\mathcal{L}_{\text{RD}} = \frac{1}{K^o+K^t} \sum_{k=1}^{K^o+K^t} \frac{\Gamma_k^s \cdot \mathcal{D}_{\text{KL}}(\tilde{\mathbf{P}}_k^t, \tilde{\mathbf{Y}}_k^t)}{\sum_{\eta=1}^t \Gamma_\eta^s \cdot \mathbb{I}_{k \in \mathcal{Y}^\eta}}, \quad (13)$$

where $\mathcal{D}_{\text{KL}}(\cdot||\cdot)$ indicates the Kullback-Leibler divergence.

Table 1. Results on CIFAR-100 [33]. When $\mathcal{M} = 2000$, we set $T = \{5, 10, 20, 25, 50\}$ for $\mathcal{B} = 0\%$, and $T = \{5, 10, 25, 50\}$ for $\mathcal{B} = 50\%$.

Comparison Methods	Backbone	#Params	$\mathcal{B} = 0\%$					$\mathcal{B} = 50\%$							
			5	10	20	25	50	Avg.	Imp.	5	10	25	50	Avg.	Imp.
iCaRL [40] (CVPR'2017)	ViT-Base	85.10M	83.5	80.6	78.5	77.6	76.4	79.3	↑5.5	80.6	78.6	75.4	71.0	76.4	↑2.9
BiC [50] (CVPR'2019)	ViT-Base	85.10M	83.5	82.5	78.7	80.3	78.4	80.7	↑4.2	78.2	75.7	74.2	71.9	75.0	↑4.3
PODNet [17] (ECCV'2020)	ViT-Base	85.10M	75.3	71.8	66.5	64.9	63.0	68.3	↑16.6	72.8	73.0	74.2	74.9	73.7	↑5.6
SS-IL [2] (ICCV'2021)	ViT-Base	85.10M	83.6	82.0	79.1	79.1	77.3	80.2	↑4.6	75.8	74.8	72.9	72.0	73.9	↑5.4
PODNet [17] + CSCCT [4] (ECCV'2022)	ViT-Base	85.10M	75.2	72.0	66.1	65.2	63.8	68.5	↑16.4	72.7	72.8	74.0	74.7	73.5	↑5.8
FOSTER [47] (ECCV'2022)	ViT-Base	85.10M	84.7	83.5	79.6	78.5	77.9	80.8	↑4.1	78.4	77.3	76.2	74.0	76.5	↑2.8
AFC [32] (CVPR'2022)	ViT-Base	85.10M	74.6	75.5	70.3	67.4	64.7	70.5	↑14.4	70.8	69.4	68.9	69.9	69.8	↑9.5
DyTox [18] (CVPR'2022)	ViT-Base	85.10M	85.5	85.8	81.3	80.8	77.4	82.2	↑2.7	82.5	81.0	77.3	74.5	78.8	↑0.5
HFC (Ours)	ViT-Base	85.10M	86.4	86.3	85.5	85.0	81.1	84.9	–	82.9	81.0	77.8	75.5	79.3	–
Upper Bound	ViT-Base	85.10M	94.2	94.2	94.2	94.2	94.2	94.2	–	94.2	94.2	94.2	94.2	94.2	–

Table 2. Results on ImageNet-100 [10]. When $\mathcal{M} = 2000$, we set $T = \{5, 10, 20, 25\}$ for $\mathcal{B} = 0\%$, and $T = \{5, 10, 25, 50\}$ for $\mathcal{B} = 50\%$.

Comparison Methods	Backbone	#Params	$\mathcal{B} = 0\%$					$\mathcal{B} = 50\%$						
			5	10	20	25	Avg.	Imp.	5	10	25	50	Avg.	Imp.
iCaRL [40] (CVPR'2017)	ViT-Base	85.10M	84.1	82.4	80.2	78.6	81.3	↑3.2	79.5	78.8	76.6	73.9	77.2	↑3.4
BiC [50] (CVPR'2019)	ViT-Base	85.10M	81.4	80.2	79.2	78.5	79.8	↑4.7	73.9	74.4	73.4	72.2	73.5	↑7.1
PODNet [17] (ECCV'2020)	ViT-Base	85.10M	82.4	80.1	79.5	78.3	80.1	↑4.5	68.4	71.3	75.4	77.6	73.2	↑7.4
SS-IL [2] (ICCV'2021)	ViT-Base	85.10M	83.0	81.7	80.9	79.4	81.2	↑3.3	75.8	74.8	72.9	72.0	73.9	↑6.7
PODNet [17] + CSCCT [4] (ECCV'2022)	ViT-Base	85.10M	82.3	81.0	79.3	79.0	80.4	↑4.1	68.3	71.2	75.0	76.9	72.8	↑7.8
FOSTER [47] (ECCV'2022)	ViT-Base	85.10M	85.6	84.0	83.8	83.1	84.1	↑0.4	81.0	80.5	80.0	79.8	80.3	↑0.3
AFC [32] (CVPR'2022)	ViT-Base	85.10M	83.2	81.2	80.5	79.4	81.1	↑3.5	77.1	78.4	78.9	78.1	78.1	↑2.5
DyTox [18] (CVPR'2022)	ViT-Base	85.10M	85.1	83.4	80.0	80.3	82.2	↑2.3	77.4	78.4	79.7	79.4	78.7	↑1.9
HFC (Ours)	ViT-Base	85.10M	86.1	85.4	84.0	82.6	84.5	–	81.5	81.1	80.3	79.5	80.6	–
Upper Bound	ViT-Base	85.10M	86.3	86.3	86.3	86.3	86.3	–	86.3	86.3	86.3	86.3	86.3	–

Table 3. Results on ImageNet-1000 [10]. When $\mathcal{M} = 20000$, we set $T = \{5, 10\}$ for $\mathcal{B} = 0\%$.

Comparison Methods	Backbone	#Params	5	10	Avg.	Imp.
iCaRL [40] (CVPR'2017)	ViT-Base	85.10M	75.4	72.4	73.9	↑3.5
BiC [50] (CVPR'2019)	ViT-Base	85.10M	68.9	66.0	67.5	↑10.0
PODNet [17] (ECCV'2020)	ViT-Base	85.10M	68.4	71.3	69.8	↑7.6
SS-IL [2] (ICCV'2021)	ViT-Base	85.10M	75.8	74.8	75.3	↑2.2
PODNet [17] + CSCCT [4] (ECCV'2022)	ViT-Base	85.10M	68.4	63.2	65.8	↑11.6
FOSTER [47] (ECCV'2022)	ViT-Base	85.10M	76.9	74.3	75.6	↑1.9
AFC [32] (CVPR'2022)	ViT-Base	85.10M	74.2	71.4	72.8	↑4.6
DyTox [18] (CVPR'2022)	ViT-Base	85.10M	78.1	74.7	76.4	↑1.0
HFC (Ours)	ViT-Base	85.10M	78.5	76.4	77.5	–
Upper Bound	ViT-Base	85.10M	86.3	86.3	86.3	–

3.5. Optimization Pipeline

Overall, in the t -th ($t \geq 2$) learning task, the objective formulation to optimize Θ^t is expressed as follows:

$$\mathcal{L}_{\text{obj}} = \alpha_1 \mathcal{L}_{\text{FC}} + \alpha_2 \mathcal{L}_{\text{RD}}, \quad (14)$$

where α_1, α_2 are hyper-parameters. Θ^t along with learnable \mathbf{E}_0^t are optimized via \mathcal{L}_{CE} in Eq. (5) for the first task, and trained via \mathcal{L}_{obj} in Eq. (14) when $t \geq 2$. After optimizing new model Θ^t , we store Θ^t as the frozen old model Θ^{t-1} to perform the GRD loss \mathcal{L}_{RD} for the next incremental task.

4. Experiments

4.1. Implementation Details

For fair comparisons, we follow the same protocols utilized in baseline CIL methods [40, 50, 17, 17, 2, 18, 32] to

set incremental tasks (class order), and do experiments on CIFAR-100 [33], ImageNet-100 [10] and ImageNet-1000 [10]. Specifically, we consider two task settings in this paper: 1) We divide all categories of each dataset into T tasks equally [40, 18] (*i.e.*, the base class set $\mathcal{B} = 0\%$); 2) We start by training our model on half of the categories, while divide the rest of categories into T tasks equally [17, 32] (*i.e.*, the base class set $\mathcal{B} = 50\%$). When $\mathcal{B} = 0\%$, we set $T = \{5, 10, 20, 25, 50\}$ for CIFAR-100 [33], $T = \{5, 10, 20, 25\}$ for ImageNet-100 [10] and $T = \{5, 10\}$ for ImageNet-1000 [10]. When $\mathcal{B} = 50\%$, we set $T = \{5, 10, 25, 50\}$ for CIFAR-100 [33] and ImageNet-100 [10].

As introduced in [40, 50], the size of memory \mathcal{M} for all comparison methods is fixed as 2,000 for CIFAR-100 [33] and ImageNet-100 [10], and 20,000 for ImageNet-1000 [10]. We follow iCaRL [40] to update memory \mathcal{M} and use global memory for [18]. Besides, we employ the same data augmentation proposed in DyTox [18] for all comparison methods, but don't use mixup technology. Following [17, 32], we also consider another memory setting (*i.e.*, storing 20 exemplars per class), and set $T = \{5, 10\}$, $\mathcal{B} = 50\%$ for comparisons on CIFAR-100 [33] (see Tab. 6). For network architecture, we use ViT-Base [15] as feature extractor containing $l_m = 11$ MSA blocks and $l_t = 1$ TSA block, where the parameters are initialized via [22]. The task classifier includes only one fully-connected layer. The SGD optimizer is employed to optimize our model, where the learning rate is initialized as 6.25×10^{-5} . We apply two plug-and-play losses (*i.e.*, \mathcal{L}_{FC}

Table 4. Performance on CIFAR-100 [33] ($T = \{5, 10, 20, 25, 50\}$) and ImageNet-100 [10] ($T = \{5, 10, 20, 25\}$) when we apply **Ours**[‡] into existing distillation-based CIL methods and set $\mathcal{M} = 2000$, $\mathcal{B} = 0\%$. **Ours**[‡] denotes the proposed plug-and-play losses \mathcal{L}_{FC} and \mathcal{L}_{RD} .

Comparison Methods	CIFAR-100 Dataset [33]							ImageNet-100 Dataset [10]							
	Backbone	#Params	5	10	20	25	50	Avg.	Backbone	#Params	5	10	20	25	Avg.
iCaRL [40] (CVPR'2017)	ResNet-32	0.46M	66.4	63.9	53.1	50.2	39.1	54.5	ResNet-18	11.22M	76.4	69.5	60.6	57.9	66.1
iCaRL [40] + Ours [‡]	ResNet-32	0.46M	68.0	65.7	54.2	51.4	40.8	56.0	ResNet-18	11.22M	76.8	70.3	61.7	59.2	67.0
BiC [50] (CVPR'2019)	ResNet-32	0.46M	55.3	50.8	48.3	46.4	39.3	48.0	ResNet-18	11.22M	61.5	54.3	46.7	44.5	51.8
BiC [50] + Ours [‡]	ResNet-32	0.46M	63.2	53.8	52.4	49.3	41.5	52.0	ResNet-18	11.22M	63.2	57.0	50.3	48.4	54.7
PODNet [17] (ECCV'2020)	ResNet-32	0.46M	63.8	54.5	48.1	44.4	39.7	50.1	ResNet-18	11.22M	74.9	66.5	55.8	51.7	62.2
PODNet [17] + Ours [‡]	ResNet-32	0.46M	65.8	60.0	51.6	48.3	44.7	54.1	ResNet-18	11.22M	76.5	69.3	59.4	55.8	65.2
SS-IL [2] (ICCV'2021)	ResNet-32	0.46M	65.1	59.0	49.0	48.6	40.7	52.5	ResNet-18	11.22M	60.9	53.5	39.3	35.8	47.4
SS-IL [2] + Ours [‡]	ResNet-32	0.46M	67.7	64.9	51.6	51.2	42.5	55.6	ResNet-18	11.22M	75.2	67.5	57.9	56.0	64.2
PODNet [17] + CSCCT [4] (ECCV'2022)	ResNet-32	0.46M	63.7	54.3	47.8	44.0	39.0	49.8	ResNet-18	11.22M	74.8	66.3	55.4	51.3	62.0
PODNet [17] + CSCCT [4] + Ours [‡]	ResNet-32	0.46M	66.7	58.5	50.5	46.9	43.3	53.2	ResNet-18	11.22M	76.4	68.7	58.9	55.1	64.8
FOSTER [47] (ECCV'2022)	ResNet-32	0.46M	72.0	68.7	67.1	66.0	62.6	67.3	ResNet-18	11.22M	74.9	70.9	67.2	65.7	69.7
FOSTER [47] + Ours [‡]	ResNet-32	0.46M	72.2	69.9	67.4	66.7	62.8	67.8	ResNet-18	11.22M	75.0	71.3	67.5	66.2	70.0
AFC [32] (CVPR'2022)	ResNet-32	0.46M	65.6	57.1	50.9	48.1	42.0	52.7	ResNet-18	11.22M	76.3	69.7	61.1	58.1	66.3
AFC [32] + Ours [‡]	ResNet-32	0.46M	66.2	63.8	53.2	49.7	44.1	55.4	ResNet-18	11.22M	77.4	71.2	63.3	67.4	69.8
Upper Bound	ResNet-32	0.46M	76.6	76.6	76.6	76.6	76.6	76.6	ResNet-18	11.22M	73.2	73.2	73.2	73.2	73.2
DyTox [18] (CVPR'2022)	ViT-Tiny	10.71M	75.3	73.5	72.7	72.3	70.5	72.9	ViT-Tiny	10.71M	77.8	75.4	72.9	72.3	74.6
DyTox [18] + Ours [‡]	ViT-Tiny	10.71M	75.9	74.0	72.9	72.6	70.8	73.2	ViT-Tiny	10.71M	77.9	75.6	73.1	72.9	74.9
Upper Bound	ViT-Tiny	10.71M	76.1	76.1	76.1	76.1	76.1	76.1	ViT-Tiny	10.71M	79.1	79.1	79.1	79.1	79.1

Table 5. Ablation studies on CIFAR-100 [33] (the top block) and ImageNet-100 [10] (the bottom block), when we set $T = \{5, 10, 25\}$ under the settings of $\mathcal{M} = 2000$, $\mathcal{B} = 0\%$.

Ablation Variants	Backbone	#Params	TSA	GFC	GRD	5	10	25
Baseline	ViT-Base	85.10M	✗	✗	✗	81.6	80.2	78.1
Baseline + TSA	ViT-Base	85.10M	✓	✗	✗	83.1	82.4	81.0
Baseline + TSA + GFC	ViT-Base	85.10M	✓	✓	✗	84.0	83.7	82.1
HFC (Ours)	ViT-Base	85.10M	✓	✓	✓	86.4	86.3	85.0
Baseline	ViT-Base	85.10M	✗	✗	✗	82.0	80.5	78.8
Baseline + TSA	ViT-Base	85.10M	✓	✗	✗	83.7	82.8	81.2
Baseline + TSA + GFC	ViT-Base	85.10M	✓	✗	✗	84.3	83.4	82.0
HFC (Ours)	ViT-Base	85.10M	✓	✓	✓	86.1	85.4	84.0

and \mathcal{L}_{RD}) to existing CIL methods [40, 18] using ResNet-18 [23], ResNet-32 [23] or ViT-Tiny [15] as backbone. For evaluation, top-1 accuracy is used to compare performances of our model and baseline methods.

4.2. Comparison Experiments

Comparison Results: Tabs. 1–3 present comparison results of our HFC model and off-the-shelf CIL methods on CIFAR-100 [33], ImageNet-100 [10] and ImageNet-1000 [10]. Our model significantly outperforms existing CIL methods [40, 50, 17, 2, 18, 32, 47] by 0.7% ~ 16.6% accuracy under various incremental settings, when we use the same backbone (*i.e.*, ViT-Base [15]) as feature extractor for fair comparisons. Such large performance improvement illustrates the superiority of our model to address forgetting heterogeneity among easy-to-forget and hard-to-forget old categories from both representation and gradient aspects.

Improvement of Plug-and-Play Losses: Tab. 4 shows large performance improvement of existing knowledge-

based CIL methods [40, 50, 17, 4, 18, 32, 47], when we apply the proposed plug-and-play losses (*i.e.*, \mathcal{L}_{FC} and \mathcal{L}_{RD}) to them. The proposed losses \mathcal{L}_{FC} and \mathcal{L}_{RD} help existing CIL methods to address heterogeneous forgetting from gradient aspect, thus largely improving their performance. Besides, it validates generalization and robustness of our HFC model.

4.3. Ablation Studies

As shown in Tab. 5, we present ablation studies on CIFAR-100 [33] and ImageNet-100 [10] to investigate the effectiveness of each module. TSA, GFC and GRD are task-semantic aggregation blocks, the gradient-balanced forgetting compensation loss \mathcal{L}_{FC} and gradient-balanced relation distillation loss \mathcal{L}_{RD} . Baseline denotes the performance of our model using the traditional classification loss and knowledge distillation loss proposed in DyTox (CVPR'2022) [18], but without the TSA, GFC and GRD modules. Compared with Ours, the performance of all ablation variants degrades significantly. It validates superiority of each module to overcome forgetting heterogeneity of different old categories.

4.4. Analysis of Task-Wise Comparisons

Figs. 3–4 present task-wise performance comparisons between our model and other CIL methods [40, 50, 17, 18, 32] when we set backbone as ViT-Base and different tasks as $T = \{5, 10, 20, 25\}$. Our HFC model outperforms baseline methods for most task-wise comparisons, which illustrates the superior performance of our model to surmount heterogeneous forgetting from representation and gradient aspects. To identify new classes continually, the task-semantic aggregation (TSA) block explores task-shared global representa-

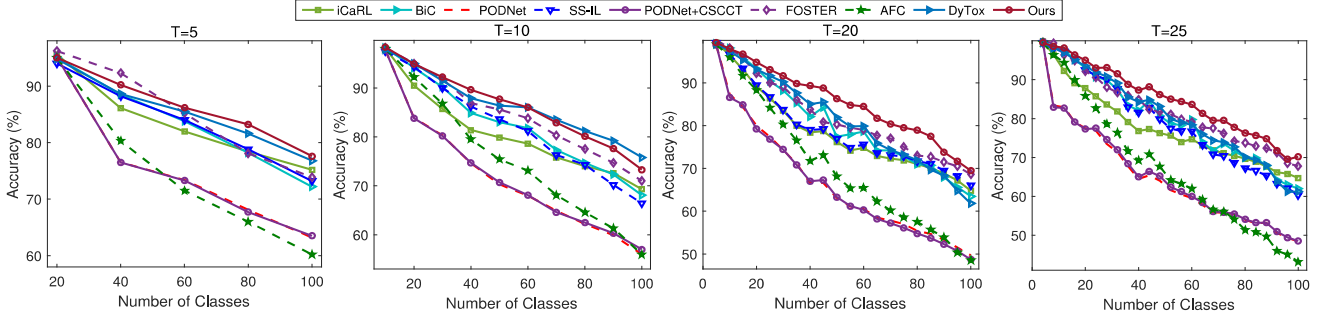


Figure 3. Analysis of task-wise performance comparisons on CIFAR-100 [33]. We set $\mathcal{M} = 2000$, $\mathcal{B} = 0\%$ when the backbone is ViT-Base.

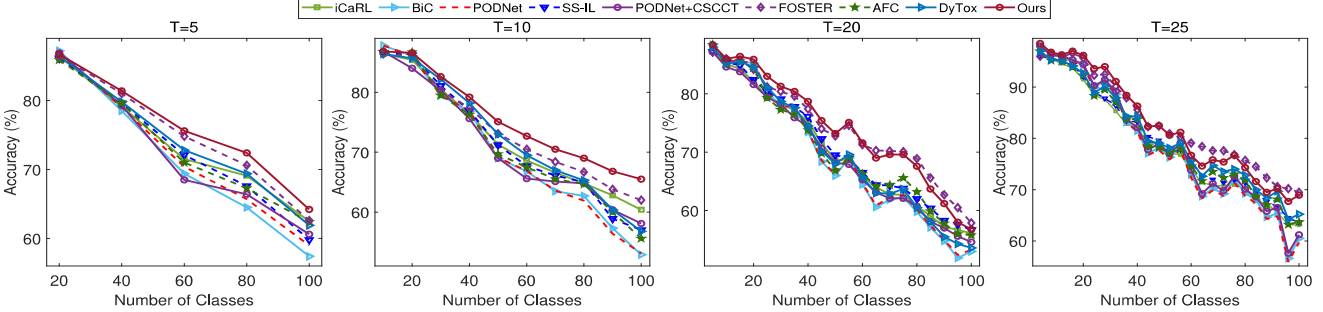


Figure 4. Analysis of task-wise performance on ImageNet-100 [10]. We set $\mathcal{M} = 2000$, $\mathcal{B} = 0\%$ when the backbone is ViT-Base.

Table 6. Results on CIFAR-100 [33] when storing 20 exemplars for each class. We set $T = \{5, 10\}$ for $\mathcal{B} = 50\%$, and **Ours[†]** denotes the proposed plug-and-play losses \mathcal{L}_{FC} and \mathcal{L}_{RD} .

Comparison Methods	Backbone	#Params	5	10	Avg. Imp.
iCaRL [40] (CVPR'2017)	ViT-Base	85.10M	78.8	76.2	77.5 $\uparrow 1.9$
BiC [50] (CVPR'2019)	ViT-Base	85.10M	75.3	74.5	74.9 $\uparrow 4.5$
PODNet [17] (ECCV'2020)	ViT-Base	85.10M	71.6	71.6	71.6 $\uparrow 7.8$
SS-IL [2] (ICCV'2021)	ViT-Base	85.10M	73.0	71.8	72.4 $\uparrow 7.0$
PODNet [17] + CSCCT [4] (ECCV'2022)	ViT-Base	85.10M	71.5	71.7	71.6 $\uparrow 7.8$
FOSTER [47] (ECCV'2022)	ViT-Base	85.10M	78.8	78.7	78.8 $\uparrow 0.6$
AFC [32] (CVPR'2022)	ViT-Base	85.10M	68.2	65.3	66.8 $\uparrow 12.6$
DyTox [18] (CVPR'2022)	ViT-Base	85.10M	79.3	78.4	78.8 $\uparrow 0.6$
HFC (Ours)	ViT-Base	85.10M	79.8	78.9	79.4 —
Upper Bound	ViT-Base	85.10M	94.2	94.2	94.2 —
iCaRL [40] (CVPR'2017)	ResNet-32	0.46M	56.6	53.2	54.9 $\uparrow 0.7$
iCaRL [40] + Ours[†]	ResNet-32	0.46M	57.3	53.9	55.6 —
PODNet [17] + CSCCT [4] (ECCV'2022)	ResNet-32	0.46M	62.2	61.1	61.7 $\uparrow 0.6$
PODNet [17] + CSCCT [4] + Ours[†]	ResNet-32	0.46M	62.9	61.6	62.3 —
FOSTER [47] (ECCV'2022)	ResNet-32	0.46M	67.3	66.4	66.8 $\uparrow 0.5$
FOSTER [47] + Ours[†]	ResNet-32	0.46M	67.7	66.6	67.2 —
AFC [32] (CVPR'2022)	ResNet-32	0.46M	63.8	63.6	63.7 $\uparrow 0.7$
AFC [32] + Ours[†]	ResNet-32	0.46M	64.4	64.3	64.4 —
Upper Bound	ResNet-32	0.46M	76.6	76.6	76.6 —
DyTox [18] (CVPR'2022)	ViT-Tiny	10.71M	67.9	66.2	67.1 $\uparrow 0.5$
DyTox [18] + Ours[†]	ViT-Tiny	10.71M	68.5	66.7	67.6 —
Upper Bound	ViT-Tiny	10.71M	76.1	76.1	76.1 —

tions to alleviate forgetting from representation aspect, while \mathcal{L}_{FC} and \mathcal{L}_{RD} can achieve gradient-balanced compensation.

4.5. Analysis of Forgetting Heterogeneity

Tab. 7 presents forgetting heterogeneity (FH) of different old classes via measuring their variance of gradient updating. We formulate forgetting heterogeneity (FH) as $FH = \frac{1}{T} \sum_{t=1}^T (\frac{1}{S_t} \sum_{i=1}^{S_t} (|\Gamma_i^t| - \sum_{\eta=1}^t \Gamma_{\eta} \cdot \mathbb{I}_{y_i^t \in \mathcal{Y}_{\eta}})^2)$, where Γ_i^t and Γ_{η} are obtained from Eqs. (6)–(7), and S_t is

Table 7. Analysis of forgetting heterogeneity (FH) on CIFAR-100 [33] when $T = 5$, $\mathcal{B} = 0\%$. AC denotes the averaged accuracy.

$\mathcal{M} = 2000, \mathcal{B} = 0\%$	Backbone	TSA	GFC	GRD	FH	AC
iCaRL [40] (CVPR'2017)	ViT-Base	—	—	—	131.5	83.5
FOSTER [47] (ECCV'2022)	ViT-Base	—	—	—	119.4	85.1
AFC [32] (CVPR'2022)	ViT-Base	—	—	—	108.5	74.6
DyTox [18] (CVPR'2022)	ViT-Base	—	—	—	105.1	85.5
Baseline	ViT-Base	✗	✗	✗	128.9	81.6
Baseline + TSA	ViT-Base	✓	✗	✗	125.4	83.1
Baseline + TSA + GFC	ViT-Base	✓	✓	✗	120.5	84.0
HFC (Ours)	ViT-Base	✓	✓	✓	103.6	86.4

number of samples in test set. Tab. 7 verifies superiority of our model against other CIL methods to tackle heterogeneous forgetting. It also shows effectiveness of all modules to collaboratively minimize forgetting heterogeneity.

5. Conclusion

In this paper, we develop a novel Heterogeneous Forgetting Compensation (HFC) model to surmount heterogeneous forgetting from representation and gradient aspects. To be specific, a task-semantic aggregation block is designed to tackle heterogeneous forgetting from representation aspect via exploring task-shared global representations. Moreover, we propose a gradient-balanced forgetting compensation loss and a relation distillation loss to compensate forgetting heterogeneity from gradient aspect via performing balanced gradient propagation and distilling heterogeneous class relations. Experiments verify the superiority of our proposed HFC model against baseline methods. We will further consider addressing noisy forgetting heterogeneity in the future.

References

- [1] Davide Abati, Jakub Tomczak, Tijmen Blankevoort, Simone Calderara, Rita Cucchiara, and Babak Ehteshami Bejnordi. Conditional channel gated networks for task-aware continual learning. In *CVPR*, pages 3930–3939, 2020.
- [2] Hongjoon Ahn, Jihwan Kwak, Subin Lim, Hyeonsu Bang, Hyojun Kim, and Taesup Moon. Ss-il: Separated softmax for incremental learning. In *ICCV*, pages 844–853, October 2021.
- [3] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *ECCV*, pages 144–161, 2018.
- [4] Arjun Ashok, K. J. Joseph, and Vineeth N. Balasubramanian. Class-incremental learning with cross-space clustering and controlled transfer. In *ECCV*, pages 105–122, 2022.
- [5] Matteo Boschini, Lorenzo Bonicelli, Pietro Buzzega, Angelo Porrello, and Simone Calderara. Class-incremental continual learning into the extended der-verse. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–16, 2022.
- [6] Francisco M. Castro, Manuel J. Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *ECCV*, pages 241–257, Cham, 2018. Springer International Publishing.
- [7] Fei-Long Chen, Du-Zhen Zhang, Ming-Lun Han, Xiu-Yi Chen, Jing Shi, Shuang Xu, and Bo Xu. Vlp: A survey on vision-language pre-training. *Machine Intelligence Research*, 20(1):38–56, 2023.
- [8] Huitong Chen, Yu Wang, and Qinghua Hu. Multi-granularity regularized re-balancing for class incremental learning. *IEEE Transactions on Knowledge and Data Engineering*, pages 1–15, 2022.
- [9] Stéphane d’Ascoli, Hugo Touvron, Matthew Leavitt, Ari Morcos, Giulio Biroli, and Levent Sagun. Convit: Improving vision transformers with soft convolutional inductive biases. In *ICML*, 2021.
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009.
- [11] Jiahua Dong, Yang Cong, Gan Sun, Zhen Fang, and Zhengming Ding. Where and how to transfer: Knowledge aggregation-induced transferability perception for unsupervised domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 11–20, 2021.
- [12] Jiahua Dong, Lixu Wang, Zhen Fang, Gan Sun, Shichao Xu, Xiao Wang, and Qi Zhu. Federated class-incremental learning. In *CVPR*, pages 10164–10173, June 2022.
- [13] Jiahua Dong, Duzhen Zhang, Yang Cong, Wei Cong, Henghui Ding, and Dengxin Dai. Federated incremental semantic segmentation. In *CVPR*, pages 3934–3943, June 2023.
- [14] Xin Dong, Junfeng Guo, Ang Li, Wei-Te Ting, Cong Liu, and H.T. Kung. Neural mean discrepancy for efficient out-of-distribution detection. In *CVPR*, pages 19217–19227, June 2022.
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021.
- [16] Arthur Douillard, Yifu Chen, Arnaud Dapogny, and Matthieu Cord. Plop: Learning without forgetting for continual semantic segmentation. In *CVPR*, 2021.
- [17] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *ECCV*, pages 86–102, 2020.
- [18] Arthur Douillard, Alexandre Ramé, Guillaume Couairon, and Matthieu Cord. Dytox: Transformers for continual learning with dynamic token expansion. In *CVPR*, pages 9285–9295, June 2022.
- [19] Zhen Fang, Yixuan Li, Jie Lu, Jiahua Dong, Bo Han, and Feng Liu. Is out-of-distribution detection learnable? In *NeurIPS*, volume 35, pages 37199–37213, 2022.
- [20] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. Convolutional sequence to sequence learning. In *ICML*, page 1243–1252, 2017.
- [21] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, volume 27, 2014.
- [22] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, pages 16000–16009, June 2022.
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [24] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NeurIPS Workshop*, 2015.
- [25] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *CVPR*, June 2019.
- [26] Weiming Hu, Xi Li, Guodong Tian, Stephen Maybank, and Zhongfei Zhang. An incremental dpmm-based method for trajectory clustering, modeling, and retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(5):1051–1065, 2013.
- [27] Xinting Hu, Kaihua Tang, Chunyan Miao, Xian-Sheng Hua, and Hanwang Zhang. Distilling causal effect of data in class-incremental learning. In *CVPR*, 2021.
- [28] Ching-Yi Hung, Cheng-Hao Tu, Cheng-En Wu, Chien-Hung Chen, Yi-Ming Chan, and Chu-Song Chen. Compacting, picking and growing for unforgetting continual learning. In *Advances in Neural Information Processing Systems*, pages 13647–13657, 2019.
- [29] Ahmet Iscen, Jeffrey Zhang, Svetlana Lazebnik, and Cordelia Schmid. Memory-efficient incremental learning through feature adaptation. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *ECCV*, pages 699–715, 2020.

- [30] Andrew Jaegle, Sebastian Borgeaud, Jean-Baptiste Alayrac, Carl Doersch, Catalin Ionescu, and et al. Perceiver io: A general architecture for structured inputs & outputs. In *ICLR*, 2022.
- [31] KJ Joseph, Salman Khan, Fahad Shahbaz Khan, Rao Muhammad Anwar, and Vineeth Balasubramanian. Energy-based latent aligner for incremental learning. In *CVPR*, 2022.
- [32] Minsoo Kang, Jaeyoo Park, and Bohyung Han. Class-incremental learning by knowledge distillation with adaptive feature consolidation. In *CVPR*, pages 16071–16080, June 2022.
- [33] Alex Krizhevsky. Learning multiple layers of features from tiny images. *Technical report*, pages 32–33, 2009.
- [34] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12):2935–2947, 2018.
- [35] Chenxi Liu, Lixu Wang, Lingjuan Lyu, Chen Sun, Xiao Wang, and Qi Zhu. Deja vu: Continual model generalization for unseen domains. In *ICLR*, 2022.
- [36] Yaoyao Liu, Yingying Li, Bernt Schiele, and Qianru Sun. Online hyperparameter optimization for class-incremental learning. *AAAI*, 37(7):8906–8913, Jun. 2023.
- [37] Yaoyao Liu, Bernt Schiele, and Qianru Sun. RMM: reinforced memory management for class-incremental learning. In *NeurIPS*, pages 3478–3490, 2021.
- [38] John M. Pierre. Incremental lifelong deep learning for autonomous vehicles. In *ITSC*, pages 3949–3954, 2018.
- [39] Jathushan Rajasegaran, Munawar Hayat, Salman H Khan, Fahad Shahbaz Khan, and Ling Shao. Random path selection for continual learning. In *NeurIPS*, 2019.
- [40] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. icarl: Incremental classifier and representation learning. In *CVPR*, 2017.
- [41] Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *ICML*, volume 80, pages 4548–4557, 2018.
- [42] Yu-Ming Tang, Yi-Xing Peng, and Wei-Shi Zheng. Learning to imagine: Diversify memory for incremental learning using unlabeled data. In *CVPR*, 2022.
- [43] Zhenheng Tang, Yonggang Zhang, Shaohuai Shi, Xin He, Bo Han, and Xiaowen Chu. Virtual homogeneity learning: Defending against data heterogeneity in federated learning. In *ICML*, pages 21111–21132, 2022.
- [44] Xiaoyu Tao, Xinyuan Chang, Xiaopeng Hong, Xing Wei, and Yihong Gong. Topology-preserving class-incremental learning. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *ECCV*, pages 254–270, Cham, 2020. Springer International Publishing.
- [45] Rishabh Tiwari, Krishnateja Killamsetty, Rishabh Iyer, and Pradeep Shenoy. Gcr: Gradient coreset based replay buffer selection for continual learning. In *CVPR*, pages 99–108, June 2022.
- [46] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers. In *ICCV*, pages 32–42, 2021.
- [47] Fu-Yun Wang, Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Foster: Feature boosting and compression for class-incremental learning. In *ECCV*, pages 398–414, 2022.
- [48] Lixu Wang, Shichao Xu, Xiao Wang, and Qi Zhu. Addressing class imbalance in federated learning. In *AAAI*, volume 35, pages 10165–10173, 2021.
- [49] Liyuan Wang, Xingxing Zhang, Kuo Yang, Longhui Yu, Chongxuan Li, Lanqing Hong, Shifeng Zhang, Zhenguo Li, Yi Zhong, and Jun Zhu. Memory replay with data compression for continual learning. In *ICLR*, 2022.
- [50] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *CVPR*, pages 374–382, 2019.
- [51] Ziyang Wu, Christina Baek, Chong You, and Yi Ma. Incremental learning via rate reduction. In *CVPR*, pages 1125–1133, 2021.
- [52] Rongtao Xu, Changwei Wang, Shibiao Xu, Weiliang Meng, and Xiaopeng Zhang. Wave-like class activation map with representation fusion for weakly-supervised semantic segmentation. *IEEE Transactions on Multimedia*, 2023.
- [53] Rongtao Xu, Changwei Wang, Jiguang Zhang, Shibiao Xu, Weiliang Meng, and Xiaopeng Zhang. Rssformer: Foreground saliency enhancement for remote sensing land-cover segmentation. *IEEE Transactions on Image Processing*, 32:1052–1064, 2023.
- [54] Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *CVPR*, pages 3014–3023, 2021.
- [55] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. In *ICLR*, 2018.
- [56] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *ICML*, page 3987–3995, 2017.
- [57] Duzhen Zhang, Feilong Chen, and Xiuyi Chen. Dualgats: Dual graph attention networks for emotion recognition in conversations. In *ACL*, pages 7395–7408, 2023.
- [58] Yonggang Zhang, Mingming Gong, Tongliang Liu, Gang Niu, Xinmei Tian, Bo Han, Bernhard Schölkopf, and Kun Zhang. Causaladv: Adversarial robustness through the lens of causality. *ICLR*, 2022.
- [59] Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shu-Tao Xia. Maintaining discrimination and fairness in class incremental learning. In *CVPR*, pages 13205–13214, 2020.
- [60] Hanbin Zhao, Yongjian Fu, Mintong Kang, Qi Tian, Fei Wu, and Xi Li. Mgsvf: Multi-grained slow vs. fast framework for few-shot class-incremental learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [61] Hanbin Zhao, Fengyu Yang, Xinghe Fu, and Xi Li. Rbc: Rectifying the biased context in continual semantic segmentation. In *ECCV*, pages 55–72, 2022.