

A SECOND-ORDER PERSPECTIVE ON MODEL COMPOSITIONALITY AND INCREMENTAL LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

The fine-tuning of deep pre-trained models has revealed compositional properties, with multiple specialized modules that can be arbitrarily composed into a single, multi-task model. However, identifying the conditions that promote compositionality remains an open issue, with recent efforts concentrating mainly on linearized networks. We conduct a theoretical study that attempts to demystify compositionality in standard non-linear networks through the second-order Taylor approximation of the loss function. The proposed formulation highlights the importance of staying within the pre-training basin to achieve composable modules. Moreover, it provides the basis for two dual incremental training algorithms: the one from the perspective of multiple models trained individually, while the other aims to optimize the composed model as a whole. We probe their application in incremental classification tasks and highlight some valuable skills. In fact, the pool of incrementally learned modules not only supports the creation of an effective multi-task model but also enables unlearning and specialization in certain tasks.

1 INTRODUCTION

In the last two decades, AI technologies have predominantly been viewed as monoliths, centered around a single model trained on a single (albeit large) dataset. This paradigm has undeniably yielded impressive results and is likely to continue doing so. However, to meet the numerous challenges of business, more modular and flexible approaches, such as *model averaging* (McMahan et al., 2017) and *Mixture Of Experts* (Yadav et al., 2024a), have experienced a renewed interest. In fact, the life cycle of monolithic models demands incremental and intensive updates to accommodate the diversity of markets, platforms, and customer needs. Moreover, many companies lack the financial means to assemble a proficient AI team or access GPU clusters to support the rapid evolution of these factors. In an effort to democratize AI, we advocate for modern paradigms (Li et al., 2023b; Pfeiffer et al., 2023) prioritizing fast adaptability, model customization, and data heterogeneity.

Recent frameworks (Liu & Soatto, 2023; Bowman et al., 2023) support these principles through *model compositionality*, allowing the creation of bespoke models through cheap editing operations. Surprisingly, fine-tuning pre-trained models not only enables efficient transfer with limited data (Zhuang et al., 2020) but also offers insights into compositionality (Ilharco et al., 2022). In fact, it has been shown (Wortsman et al., 2022) that a simple linear combination of individually fine-tuned weights yields meaningful outcomes. The resulting composed model exhibits rich and robust representations (Zhang & Bottou, 2023) without incurring additional inference or memory costs. Such a simple but powerful schema has been primarily employed in *model soups* (Wortsman et al., 2022), where individual models are trained on the same task but with different hyper-parameters. In contrast, *Task arithmetic* (Ilharco et al., 2022; Liu & Soatto, 2023; Ortiz-Jimenez et al., 2024) addresses a cross-dataset setting with each individual model focusing on a distinct task. Such a framework has also been recently explored for language modeling through *Cold fusion* (Don-Yehiya et al., 2023).

In this context, we elaborate on two points. Firstly, considering multiple models trained individually on different tasks, we aim to understand the conditions that allow the successful combination of their weights. This finding has been predominantly explored empirically (Ilharco et al., 2022), with a few notable exceptions (Ortiz-Jimenez et al., 2024) that aim to provide a more theoretical understanding. However, these approaches only focus on linearized networks and tangent fine-tuning (Liu & Soatto, 2023). We instead aim for a formulation that generalizes to standard non-linear deep networks

and diverse fine-tuning strategies (e.g., LoRA (Hu et al., 2021)). To do so, we propose leveraging the second-order Taylor approximation of the loss function around the pre-training weights. This enables us to derive a relationship between the capabilities of individual models and those of their composition, offering a key takeaway. Specifically, for the composed model to be accurate, each component should attain decent performance on examples outside its individual training distribution.

The second point of our study is a derivation of the former and revolves around learning composable models in an incremental fashion, dedicating an individual module to each incoming task. In this respect, previous works (Bowman et al., 2023; Liu & Soatto, 2023) build on the intuitive idea that modularity provides a natural foundation for incremental training. Through model editing and composition, in fact, one can simply introduce novel concepts without having to re-train the model. In this work, we turn the perspective and argue that compositional capabilities require incremental learning skills. Indeed, as mentioned above, preserving out-of-distribution performance is essential for compositionality. Nevertheless, this can be framed as a continual learning problem (Kirkpatrick et al., 2017), where the objective of each individual model is to maintain the pre-training general capabilities on examples that fall outside its specific training distribution.

On this basis, we propose two algorithms that tackle the incremental fine-tuning of composable models. The two algorithms build upon the same second-order formulation but approach the problem from opposed perspectives: optimizing the loss of each model individually *vs.* optimizing the loss of the composed model. We assess them on several *class-incremental* settings (Van de Ven & Tolias, 2019), differing in terms of alignment with pre-training knowledge. Notably, our approaches lead to an accurate multi-task model with further editing capabilities. These allow the efficient *specialization* of the model on certain tasks, as well as the ability of removing some others (*unlearning*).

2 FRAMEWORK

We consider $f(\cdot; \theta) : \mathcal{X} \rightarrow \mathcal{Y}$ as a twice-differentiable deep network with weights $\theta \in \Theta \subseteq \mathbb{R}^m$. It takes inputs $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^d$ and yields a conditional distribution $p_\theta(\mathbf{y}|\mathbf{x})$ over the targets $\mathbf{y} \in \mathcal{Y} \subseteq \mathbb{R}^c$. In this paper, we focus on **incremental training**, which progresses sequentially through a series of T classification tasks $\mathcal{T} = \{1, 2, \dots, T\}$. Each task $t \sim \mathcal{T}$ is characterized by a dataset \mathcal{D}_t with n_t training samples $\mathbf{x}, \mathbf{y} \sim p_t(\mathbf{x}, \mathbf{y})$ drawn from a distribution varying across tasks. We assume that tasks share the same loss function $\ell(\theta|\mathbf{x}, \mathbf{y})$, *i.e.*, the negative log-likelihood $-\log p_\theta(\mathbf{y}|\mathbf{x})$.

In this setting, we build upon a pool of fine-tuned composable networks $\mathcal{P} = \{f(\cdot; \theta_t) \mid \theta_t \triangleq \theta_0 + \tau_t\}_{t \in \mathcal{T}}$. In detail, each model is fine-tuned from a common set of pre-training weights θ_0 . The τ_t indicates the displacement in parameter space w.r.t. θ_0 after training on task t , often referred as *task vector* (Ilharco et al., 2022; Ortiz-Jimenez et al., 2024). We obtain the weights of the composed model $f_{\mathcal{P}}$ by averaging the weights within the pool, as follows:

$$f_{\mathcal{P}} \triangleq f(\cdot; \theta_{\mathcal{P}}) \quad \text{s.t.} \quad \theta_{\mathcal{P}} = \theta_0 + \sum_{t=1}^T w_t \tau_t, \quad \sum_{t=1}^T w_t = 1 \quad (1)$$

where w_t balances the contribution of the t -th learner. While some works (Asadi et al., 2024; Huang et al., 2023) optimize these coefficients, we devise uniform weights $w_t = 1/T$ in our algorithms.

Scope. How can we learn multiple disjoint tasks through a pool \mathcal{P} of models, so that the composed model performs well on their union? To answer this question, we introduce the concept of **empirical risk**, *i.e.*, the average loss $\hat{\ell}(\theta|\mathcal{D})$, computed over the union $\mathcal{D} = \bigcup_{t=1}^T \mathcal{D}_t$ of all training tasks:

$$\hat{\ell}(\theta|\mathcal{D}) = \frac{1}{\sum_{t=1}^T n_t} \sum_{\mathbf{x}, \mathbf{y} \in \mathcal{D}} \ell(\theta|\mathbf{x}, \mathbf{y}) \approx \mathbb{E}_{\substack{t \sim \mathcal{T} \\ \mathbf{x}, \mathbf{y} \sim p_t(\mathbf{x}, \mathbf{y})}} [\ell(\theta|\mathbf{x}, \mathbf{y})] \quad (2)$$

To simplify notation, we will henceforth omit the explicit dependence of the loss on the data, denoting the individual loss $\ell(\theta|\mathbf{x}, \mathbf{y})$ simply as $\ell(\theta)$, and the empirical risk $\hat{\ell}(\theta|\mathcal{D})$ as $\hat{\ell}(\theta)$. On this basis, the rest of this section will delve into the following research questions.

Question i) Given the empirical risk of each individual model, what can we say about the composed model $f(\cdot; \theta_{\mathcal{P}})$? *Question ii)* How can we train individual learners $f(\cdot; \theta_t)$ on distinct tasks (**individual training**) to still achieve a reliable composition $f(\cdot; \theta_{\mathcal{P}})$? *Question iii)* Instead of optimizing each model on its individual loss, could we optimize each model based on the loss of the whole composed model $f(\cdot; \theta_{\mathcal{P}})$ (**ensemble training**)?

2.1 INDIVIDUAL LEARNERS vs. THE COMPOSED MODEL: A PRE-TRAINING PERSPECTIVE

We now relate the composed model $f_{\mathcal{P}}$ to the individual components $f(\cdot; \theta_t)$ of the pool. To do so, we introduce the **second-order** Taylor approximation $\ell_{\text{cur}}(\theta)$ of the loss around the pre-trained weights θ_0 :

$$\ell(\theta) = \ell_{\text{cur}}(\theta) + \mathcal{O}(\|\theta - \theta_0\|^3) \quad (3)$$

$$\text{where } \ell_{\text{cur}}(\theta) = \ell(\theta_0) + (\theta - \theta_0)^{\text{T}} \nabla \ell(\theta_0) + \frac{1}{2} (\theta - \theta_0)^{\text{T}} \mathbf{H}_{\ell}(\theta_0) (\theta - \theta_0). \quad (4)$$

$\nabla \ell(\theta_0) \triangleq \nabla_{\theta} \ell(\theta_0 | \mathbf{x}, \mathbf{y})$ and $\mathbf{H}_{\ell}(\theta_0) \triangleq \nabla_{\theta}^2 \ell(\theta_0 | \mathbf{x}, \mathbf{y})$ are the gradient and the Hessian around θ_0 . Similarly, we can define the second-order approximation $\hat{\ell}_{\text{cur}}(\theta) \approx \hat{\ell}(\theta)$ of the empirical risk, which corresponds to averaging the approximated loss across examples from all tasks.

Assumption. We now assume that $\theta = \theta_0$ is a point of **local minimum** of the empirical risk $\hat{\ell}(\theta)$ ¹ across all tasks (Eq. 2). Under this hypothesis, the Hessian of the empirical risk $\mathbf{H}_{\hat{\ell}}(\theta_0) \succeq 0$ is positive semidefinite. In light of this and the quadratic nature of $\hat{\ell}_{\text{cur}}(\theta)$, we can state that the second-order approximation of the empirical risk $\hat{\ell}_{\text{cur}}(\theta)$ is **convex**. Therefore, we apply the **Jensen’s inequality** to derive a relation between the composed model and the individual components:

$$\hat{\ell}_{\text{cur}}(\theta_{\mathcal{P}} = \theta_0 + \sum_{t=1}^T w_t \tau_t) \leq \sum_{t=1}^T w_t \hat{\ell}_{\text{cur}}(\theta_t = \theta_0 + \tau_t). \quad (5)$$

The relation states that, under the second-order approximation, the empirical risk $\hat{\ell}_{\text{cur}}(\theta_{\mathcal{P}})$ of the composed model is upper-bounded by the convex combination of its individuals. In other words, if each individual model is trained to the optimum with near-zero loss value, there are some guarantees on the loss function attained by the composed model. Notably, this relation could help reduce the computational footprint during inference, as it enables the reduction of forward passes, from multiple (one for each individual) to a singular pass (performed on the composed model).

At a first glance, the result of Eq. 5 appears similar to the statement of Eq. 2 in (Liu & Soatto, 2023):

$$\hat{\ell}(\theta_{\mathcal{P}} = \theta_0 + \sum_{t=1}^T w_t \tau_t) \leq \sum_{t=1}^T w_t \hat{\ell}(\theta_t = \theta_0 + \tau_t) \quad (6)$$

$$\text{given that } f_t(\cdot; \theta_t) \triangleq f_{\text{lin}}(\cdot; \theta_t) = f(\cdot; \theta_0) + (\theta_t - \theta_0)^{\text{T}} \nabla f(\cdot; \theta_0) \quad (\text{tangentiality}) \quad (7)$$

However, some notable distinctions remain. Their inequality applies to the exact risk $\hat{\ell}$ but is valid only for linearized models (*i.e.*, fine-tuned in the tangent space of pre-training weights). In contrast, our result pertains to the *second-order* approximation $\hat{\ell}_{\text{cur}}$ of the risk and applies to *any* fine-tuning strategy (*e.g.*, LoRA, adapters, etc.). Intuitively, our inequality provides more flexibility to the training of individual learners, as long as: *i*) the learners remain in the pre-training basin, such that $\mathcal{O}(\|\theta - \theta_0\|^3) \rightarrow 0$ and ℓ_{cur} can be considered a good proxy of ℓ ; *ii*) θ_0 is a local minimum of ℓ .

2.2 ENABLING INDIVIDUAL TRAINING IN INCREMENTAL SCENARIOS

As mentioned above, a possible application of Eq. 5 is to devote each learner to a distinct task and optimize them in isolation (*i.e.*, **individual training**). Indeed, the upper bound in Eq. 5 describes a sort of worst-case scenario for the risk of the composed model: at worst, it collapses to that given by the upper bound. Nonetheless, if every individual model were *accurate* on all tasks, the right-side term of Eq. 5 would yield an appealing upper-bound to the risk of the composed model.

Limits of Eq. 5. Under the constraints of individual training, each individual learner $f(\cdot; \theta_t)$ can be optimized only on the current distribution $p_t(\mathbf{x}, \mathbf{y})$. Therefore, when considering examples from other data distributions $p_{t' \neq t}(\mathbf{x}, \mathbf{y})$, the loss $\ell_{\text{cur}}(\theta_t | \mathbf{x}, \mathbf{y})$ of the t -th individual learner is likely to be much higher for examples outside its training distribution $p_t(\mathbf{x}, \mathbf{y})$. As a consequence, the upper bound delivered by the right-side of Eq. 5 is likely to increase one task after the other, and we cannot rely on it to recover a reliable composed model $f(\cdot; \theta_{\mathcal{P}})$.

In this respect, our proposal is to tighten the upper bound of Eq. 5 through explicit regularization, which we devise during the optimization of each individual learner. In practice, each model is

¹As shown in Ostapenko et al. (2022), techniques like linear probing, latent replay, and incremental linear discriminant analysis can be employed to enforce the optimality of the base model θ_0 , along with Instruction Tuning for language modeling (Yadav et al., 2024c). See Sec. 3 for additional implementation details.

provided with a learning objective that extends beyond minimizing the loss on the assigned data distribution $p_t(\mathbf{x}, \mathbf{y})$. Specifically, to ensure decent performance on external distributions $p_{t' \neq t}(\mathbf{x}, \mathbf{y})$, we anchor the model to the pre-training knowledge for out-of-distribution examples. We do so by encouraging the predictions of $f(\cdot; \boldsymbol{\theta}_t)$ to be close to those generated by the pre-trained model $f(\cdot; \boldsymbol{\theta}_0)$, given examples $\mathbf{x}, \mathbf{y} \sim p_{t' \neq t}(\mathbf{x}, \mathbf{y})$. Denoting the pre-training posterior distribution as $p_{\boldsymbol{\theta}_0}(\mathbf{y}|\mathbf{x})$, we have:

$$\underset{\boldsymbol{\theta}_t}{\text{minimize}} \quad \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim p_t(\mathbf{x}, \mathbf{y})} [\ell_{\text{cur}}(\boldsymbol{\theta}_t | \mathbf{x}, \mathbf{y})] + \mathcal{D}_{\text{KL}}(p_{\boldsymbol{\theta}_0}(\mathbf{y}|\mathbf{x}) || p_{\boldsymbol{\theta}_t}(\mathbf{y}|\mathbf{x})). \quad (8)$$

It is noted that the computation of the KL term $\mathcal{D}_{\text{KL}}(\cdot)$ requires sampling from the external distributions $p_{t' \neq t}(\mathbf{x}, \mathbf{y})$. Theoretically, this clashes with the constraints of individual training. Thankfully, if $\tau_t \rightarrow 0$, the KL term can be approximated (Chaudhry et al., 2018) with the *distance* between $\boldsymbol{\theta}_t$ and pre-training weights $\boldsymbol{\theta}_0$:

$$\mathcal{D}_{\text{KL}}(p_{\boldsymbol{\theta}_0}(\mathbf{y}|\mathbf{x}) || p_{\boldsymbol{\theta}_t}(\mathbf{y}|\mathbf{x})) \approx \frac{1}{2}(\boldsymbol{\theta}_t - \boldsymbol{\theta}_0)^\top \mathbf{F}_{\boldsymbol{\theta}_0}(\boldsymbol{\theta}_t - \boldsymbol{\theta}_0) = \frac{1}{2} \|\boldsymbol{\tau}_t\|_{\mathbf{F}_{\boldsymbol{\theta}_0}}^2. \quad (9)$$

The distance is computed in the Riemannian manifold (Lee, 2006) induced by the Fisher Information Matrix (FIM) (Huszár, 2018); namely, a $|\theta| \times |\theta|$ positive semi-definite matrix given by:

$$\mathbf{F}_{\boldsymbol{\theta}_0} = \mathbb{E}_{\substack{t \sim \mathcal{T} \\ \mathbf{x} \sim p_t(\mathbf{x}, \mathbf{y})}} \left[\mathbb{E}_{\mathbf{y} \sim p_{\boldsymbol{\theta}_0}(\mathbf{y}|\mathbf{x})} \left[\nabla \log p_{\boldsymbol{\theta}_0}(\mathbf{y}|\mathbf{x}) \nabla \log p_{\boldsymbol{\theta}_0}(\mathbf{y}|\mathbf{x})^\top \right] \right]. \quad (10)$$

Under the local minimum hypothesis on pre-training weights, the FIM at $\boldsymbol{\theta}_0$ equals the expected² Hessian of the negative log-likelihood (Martens, 2020): $\mathbf{F}_{\boldsymbol{\theta}_0} = \mathbb{E}_{\mathbf{x}} [\mathbf{H}_{\ell}(\boldsymbol{\theta}_0)]$. Due to this connection, the FIM yields insights on the sensitivity of each parameter to changes in the data distribution.

Application to incremental learning. Given that optimizing the regularization above does not necessitate data from external tasks, it can be readily adapted to incremental learning. Following Chaudhry et al. (2018); Schwarz et al. (2018), we introduce a few additional approximations. Firstly, we limit to estimate the diagonal $\hat{\mathbf{F}}_{\boldsymbol{\theta}_0}$ of the FIM, thus avoiding the prohibitive footprint required to treat a $|\theta| \times |\theta|$ matrix. Basically, the diagonal $\hat{\mathbf{F}}_{\boldsymbol{\theta}_0}$ consists of a Monte Carlo estimate of the (squared) gradients of the log-likelihood. Secondly, we note that the expectation $\mathbb{E}_{\mathbf{x}} [\mathbf{H}_{\ell}(\boldsymbol{\theta}_0)]$ in the FIM cannot be directly estimated, as examples from all tasks are not available simultaneously but rather sequentially. We hence compute the expectation incrementally (Chaudhry et al., 2018; Schwarz et al., 2018) one task at a time. As each new task becomes available, we calculate a *local* Fisher matrix on the data of the new task and then accumulate it into a *global* running Fisher estimate. As outlined by Alg. 1, the accumulation can be thought as a simple summation, net of re-normalizing operations reflecting the number of samples of each task (see App. E).

Given Eq. 9 and the diagonal FIM, the augmented optimization problem for the t -th learner becomes:

$$\underset{\boldsymbol{\theta}_t}{\text{minimize}} \quad \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim p_t(\mathbf{x}, \mathbf{y})} [\ell_{\text{cur}}(\boldsymbol{\theta}_t | \mathbf{x}, \mathbf{y})] + \frac{\alpha}{2} \text{EWC}_{\boldsymbol{\theta}_0}(\boldsymbol{\theta}_t) \quad (11)$$

$$\text{where} \quad \text{EWC}_{\boldsymbol{\theta}_0}(\boldsymbol{\theta}_t) = \sum_i^{|\theta|} \hat{\mathbf{F}}_{\boldsymbol{\theta}_0}^{(i)} (\boldsymbol{\theta}_t^{(i)} - \boldsymbol{\theta}_0^{(i)})^2. \quad (12)$$

where $\alpha \geq 0$ is an hyper-parameter and $\text{EWC}_{\boldsymbol{\theta}_0}(\cdot)$ indicates the Riemannian distance from the pre-training weights $\boldsymbol{\theta}_0$. The acronym $\text{EWC}_{\boldsymbol{\theta}_0}$ highlights the strong analogy with Elastic Weight Consolidation (Kirkpatrick et al., 2017), a well-established approach against *catastrophic forgetting* (McCloskey & Cohen, 1989). In a sense, our term prevents forgetting pre-training knowledge; however, while our anchor is fixed at $\boldsymbol{\theta}_0$, the anchor of EWC instead shifts and focuses on the weights learned during the preceding task.

2.3 JOINT TRAINING OF THE COMPOSED MODEL IN INCREMENTAL SCENARIOS

Individual training profitably aligns with *decentralized learning* (Bowman et al., 2023), emphasizing minimal interactions between learners and privacy preservation. Nevertheless, it might be inefficient when these constraints are not of interest and the goal is simply to create an accurate model. In fact, individual training prevents a learner from leveraging the knowledge of other ensemble members, eliminating the potential for beneficial mutual transfer. For these reasons, we adopt the dual perspective, in which each model is directly optimized using the loss of the composed model $f(\cdot; \boldsymbol{\theta}_{\mathcal{P}})$ (**ensemble training**). Since the inequality in Eq. 5 does not provide much help for the explicit optimization of $\ell_{\text{cur}}(\boldsymbol{\theta}_{\mathcal{P}})$, we quantify the exact gap between the two sides of the Jensen’s inequality:

²Precisely, we take the expectation w.r.t. data from other tasks $t' \neq t$ to reflect the regularization in Eq. 8.

Theorem 1. Let us assume a pool \mathcal{P} with $T \geq 2$ models, with the t -th model parameterized by $\theta_t = \theta_0 + \tau_t$. If we compose them through coefficients w_1, \dots, w_T s.t. $w_t \in [0, 1]$ and $\sum_{t=1}^T w_t = 1$, the 2nd order approximation $\ell_{\text{cur}}(\theta_{\mathcal{P}})$ evaluated on composed weights $\theta_{\mathcal{P}} = \theta_0 + \sum_{t=1}^T w_t \tau_t$ is:

$$\ell_{\text{cur}}(\theta_{\mathcal{P}}) + \Omega(\theta_1, \dots, \theta_T) = \sum_{t=1}^T w_t \ell_{\text{cur}}(\theta_t) \quad (13)$$

$$\text{where } \Omega(\theta_1, \dots, \theta_T) = \frac{1}{2} \sum_{t=1}^T \sum_{t' < t} w_t w_{t'} (\tau_t - \tau_{t'})^T \mathbf{H}_{\ell}(\theta_0) (\tau_t - \tau_{t'}). \quad (14)$$

Proof in App. B. The equality introduces a term $\Omega(\cdot)$ that is non-negative (due to $\mathbf{H}_{\ell}(\theta_0) \succeq 0$) and depends on weights $\theta_1, \dots, \theta_T$. Therefore, $\Omega(\cdot)$ is proportional to the cumulative *distance* between every pair of learners, within the Riemannian manifold induced by the Hessian $\mathbf{H}_{\ell}(\theta_0)$ (Lee, 2006). Notably, Eq. 13 permits to draw the following insights: *i*) optimizing both the loss of the composed model $\ell_{\text{cur}}(\theta_{\mathcal{P}})$ and the term $\Omega(\cdot)$ is equivalent to individual training; *ii*) during *inference*, if $\Omega(\cdot)$ tends towards zero, performing a prediction with the composed model is akin to conducting multiple forward passes, each corresponding to an individual learner. Based on that interpretation, we now transition to a setting that considers the explicit auxiliary minimization of $\Omega(\cdot)$, as follows:

$$\underset{\theta_1, \dots, \theta_T}{\text{minimize}} \quad \mathbb{E}_{\substack{t \sim \mathcal{T} \\ \mathbf{x}, \mathbf{y} \sim p_t(\mathbf{x}, \mathbf{y})}} [\ell_{\text{cur}}(\theta_{\mathcal{P}} | \mathbf{x}, \mathbf{y}) + \beta \Omega(\theta_1, \dots, \theta_T)] \quad (15)$$

where $\beta \geq 0$ is a hyper-parameter. It is noted that minimizing $\Omega(\cdot)$ encourages alignment among task vectors, especially for those weights that are sensitive/important in the pre-training loss landscape.

Similarly to Jeffares et al. (2024), the objective of Eq. 15 can be interpreted as a smooth transition between individual and ensemble training. Indeed, given that $\Omega(\cdot) = \sum_t w_t \ell_{\text{cur}}(\theta_t) - \ell_{\text{cur}}(\theta_{\mathcal{P}})$, Eq. 15 can be restated:

$$\ell_{\text{cur}}(\theta_{\mathcal{P}}) + \beta \Omega(\theta_1, \dots, \theta_T) \stackrel{\text{Eq. 13}}{=} (1 - \beta) \ell_{\text{cur}}(\theta_{\mathcal{P}}) + \beta \sum_{t=1}^T w_t \ell_{\text{cur}}(\theta_t) \quad (16)$$

This suggests that by minimizing both the loss $\ell_{\text{cur}}(\theta_{\mathcal{P}})$ of the joint composed model and $\Omega(\cdot)$, we also implicitly optimize the individual models. Notably, this result not only paves the way for a well-performing ensemble but also for components that are reliable when considered individually.

Incremental ensembles. We now refine the problem in Eq. 15 to account for incremental settings. We firstly bring the expectation inside $\Omega(\cdot)$ (Eq. 14) and replace the Hessian $\mathbf{H}_{\ell}(\theta_0)$ with its expectation, taken across data points $\mathbf{x}, \mathbf{y} \sim p_t(\mathbf{x}, \mathbf{y})$ from all tasks up to the current one. Afterwards, we capitalize on a property mentioned in Sec. 2.2, which states that, for a point of maximum likelihood like $\theta = \theta_0$, the expected Hessian of the negative log-likelihood coincides with the Fisher F_{θ_0} . As a result, we can approximate $\mathbb{E}_{\mathbf{x}}[\mathbf{H}_{\ell}(\theta_0)]$ with the tractable diagonal Fisher \hat{F}_{θ_0} . Based on $\mathbf{H}_{\ell}(\theta_0) \approx \hat{F}_{\theta_0}$ and further steps (see App. B), we can rearrange $\Omega(\cdot)$ as:

$$\Omega(\cdot) \approx \Omega_{\hat{F}}(\theta_1, \dots, \theta_T) = \frac{1}{2} \sum_{t=1}^T w_t (1 - w_t) \underbrace{\text{EWC}_{\theta_0}(\theta_t)}_{\text{see Eq. 12}} - \sum_{t=1}^T \sum_{t' < t} w_t w_{t'} \tau_t^T \hat{F}_{\theta_0} \tau_{t'}. \quad (17)$$

Intuitively, $\Omega_{\hat{F}}(\cdot)$ aims to preserve pre-training knowledge embodied by θ_0 through the first term; through the second one, instead, it encourages pairwise alignment between task vectors. Also, we highlight that the summations depend on the number of models T . Therefore, the *more* components we manage within the ensemble, the *more* crucial it becomes to minimize $\Omega_{\hat{F}}(\cdot)$ within Eq. 15.

Finally, we plug the approximated barrier term $\Omega_{\hat{F}}(\cdot)$ into the optimization problem outlined in Eq. 15. As learning proceeds in subsequent tasks, we can optimize the composed model only one task at a time. To prevent biasing the previously learned components of the pool toward current data, at each round t we optimize only the weights $\theta_t \triangleq \theta_0 + \tau_t$ of the corresponding t -th learner, while freezing the others components of the pool. This yields the following form for the t -th learning task:

$$\underset{\theta_t}{\text{minimize}} \quad \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim p_t(\mathbf{x}, \mathbf{y})} [\ell_{\text{cur}}(\theta_{\mathcal{P}} | \mathbf{x}, \mathbf{y})] + \beta \Omega_{\hat{F}}(\theta_1, \dots, \theta_t). \quad (18)$$

Finally, if we optimize only one θ_t at a time and $w_t = \frac{1}{t}$, then several terms in Eq. 17 can be ignored, such as $\text{EWC}_{\theta_0}(\theta_{t'})$ for $t' < t$. Hence, minimizing Eq. 18 is equivalent to minimize:

$$\mathbb{E}_{\mathbf{x}, \mathbf{y} \sim p_t(\mathbf{x}, \mathbf{y})} [\ell_{\text{cur}}(\theta_{\mathcal{P}} | \mathbf{x}, \mathbf{y})] + \frac{\beta}{t} \left[(1 - \frac{1}{t}) \frac{1}{2} \text{EWC}_{\theta_0}(\theta_t) - \frac{1}{t} \sum_{t' < t} \tau_t^T \hat{F}_{\theta_0} \tau_{t'} \right]. \quad (19)$$

As shown in App. C for both full and LoRA fine-tuning, the gradients of the regularization term in Eq. 19 can be computed in **closed form**. Importantly, this derivation enables a lightweight optimization, as the analytical gradients maintain **constant complexity** w.r.t. the number of tasks.

Algorithm 1 Incremental Task Arithmetic (ITA) vs. Incremental Ensemble Learning (IEL)

Input: T disjoint classification tasks $D_t = \{(\mathbf{x}, \mathbf{y})\}_{n_t}$, a pre-trained DNN $f(\cdot; \boldsymbol{\theta}_0)$, learning rate lr , hyper-parameters α and β , initialized pool $\mathcal{P} = \emptyset$ and the diagonal Fisher $\hat{\mathbf{F}}_{\boldsymbol{\theta}_0} = \mathbf{0}$.

for each task $t \in \{1, 2, \dots, T\}$ **do**

$\mathbf{h}_0^{(t)} \leftarrow$ Linear Probing on D_t with the pre-trained $f(\cdot; \boldsymbol{\theta}_0)$ Task pre-consolidation

$\boldsymbol{\theta}_0 \leftarrow \boldsymbol{\theta}_0 \cup \mathbf{h}_0^{(t)}$ \triangleright add the pre-training classification head

$\hat{\mathbf{F}}_{\boldsymbol{\theta}_0} \leftarrow \hat{\mathbf{F}}_{\boldsymbol{\theta}_0} + \hat{\mathbf{F}}_{\boldsymbol{\theta}_0}^{(t)}$ \triangleright update the global Fisher with the local Fisher $\hat{\mathbf{F}}_{\boldsymbol{\theta}_0}^{(t)}$ estimated on D_t

$\boldsymbol{\tau}_t \leftarrow \mathcal{N}(\mu_{\text{init}}, \sigma_{\text{init}})$ Fine-tuning

$\mathcal{P} \leftarrow \mathcal{P} \cup \{f(\cdot; \boldsymbol{\theta}_t = \boldsymbol{\theta}_0 + \boldsymbol{\tau}_t)\}$ \triangleright extend \mathcal{P} with the weights of the t -th learner

for each example (\mathbf{x}, \mathbf{y}) in D_t **do**

$p_{\boldsymbol{\theta}}(\mathbf{y}|\mathbf{x}) \leftarrow f(\mathbf{x}; \boldsymbol{\theta}_t)$ \triangleright predict with the t -th learner $\boldsymbol{\theta}_t \leftarrow \boldsymbol{\theta}_0 + \boldsymbol{\tau}_t$

$p_{\boldsymbol{\theta}}(\mathbf{y}|\mathbf{x}) \leftarrow f(\mathbf{x}; \boldsymbol{\theta}_{\mathcal{P}})$ \triangleright predict with the composed model $\boldsymbol{\theta}_{\mathcal{P}} \leftarrow \boldsymbol{\theta}_0 + \frac{1}{t} \sum_{t'=1}^t \boldsymbol{\tau}_{t'}$

$\ell \leftarrow -\log p_{\boldsymbol{\theta}}(\mathbf{y}|\mathbf{x})$

$\boldsymbol{\tau}_t \leftarrow \boldsymbol{\tau}_t - \text{lr} \cdot \nabla_{\boldsymbol{\tau}_t} [\ell + \frac{\alpha}{2} \text{EWC}_{\boldsymbol{\theta}_0}(\boldsymbol{\theta}_t)]$ \triangleright arithmetic-oriented regularization (Eq. 11)

$\boldsymbol{\tau}_t \leftarrow \boldsymbol{\tau}_t - \text{lr} \cdot \nabla_{\boldsymbol{\tau}_t} [\ell + \beta \Omega_{\hat{\mathbf{F}}}(\mathcal{P})]$ \triangleright ensemble-oriented regularization (Eqs. 18 and 19)

3 ALGORITHM(S)

We present **Incremental Task Arithmetic (ITA)** and **Incremental Ensemble Learning (IEL)**, two **distinct** algorithms for *individual* and *ensemble* incremental learning respectively. As shown in Alg. 1, they divide each task into the *pre-consolidation* and *fine-tuning* stages, with differences occurring in the latter. We direct the reader to App. E for comprehensive implementation guidelines.

Task pre-consolidation. Due to the pivotal role of the local optimality of $\boldsymbol{\theta}_0$ (Sec. 2.1), we now follow up on this aspect and consider *closed vocabulary* models. Unlike models like CLIP (Radford et al., 2021), closed vocabulary models require the addition of a tailored classification head to handle novel classes. In line with Ortiz-Jimenez et al. (2024), we denote the process of fine-tuning only the added classification head as **linear probing (LP)**: specifically, LP keeps the rest of the layers fixed to the pre-training $\boldsymbol{\theta}_0$. In our work, we basically exploit linear probing to enforce the optimality of pre-training weights $\boldsymbol{\theta}_0$. Namely, during the pre-consolidation phase, we conduct a few preliminary training epochs on the t -th incoming task, with the sole purpose of fine-tuning the new classification head. From that point onward, the fine-tuned head $\mathbf{h}_0^{(t)}$ is regarded as a part of pre-training weights $\boldsymbol{\theta}_0$. Finally, the pre-consolidation stage concludes with the update of the diagonal Fisher matrix $\hat{\mathbf{F}}_{\boldsymbol{\theta}_0}$.

Fine-tuning. While the pre-consolidation step is identical for both ITA and IEL, they differ during the fine-tuning phase. The shared goal is to learn a task vector $\boldsymbol{\tau}_t$ s.t. $\boldsymbol{\theta}_t \equiv \boldsymbol{\theta}_0 + \boldsymbol{\tau}_t$ for the current task. However, ITA treats $\boldsymbol{\tau}_t$ as the weights of an individual model, whereas IEL interprets it as a new learnable component of the composition. In other words, ITA computes predictions through the t -th learner $f(\mathbf{x}; \boldsymbol{\theta}_t)$, while IEL leverages the composed function $f(\mathbf{x}; \boldsymbol{\theta}_{\mathcal{P}})$. Moreover, their regularizing objectives differ: ITA builds upon Eq. 11 (*i.e.*, the additional EWC-like term computed w.r.t. $\boldsymbol{\theta}_0$), while IEL exploit Eqs. 18 and 19 to train the composed model. Notably, both approaches can be applied to any fine-tuning strategy in the form of $\boldsymbol{\theta}_0 + \Delta\boldsymbol{\theta}$. We conduct experiments on Full Fine-Tuning (*i.e.*, $\boldsymbol{\tau}_t \in \mathbb{R}^{F_{OUT} \times F_{IN}}$), Low-Rank Adaptation (LoRA) (Hu et al., 2021) (*i.e.*, $\boldsymbol{\tau}_t = B_t A_t$), and on (IA)³ (Liu et al., 2022). About the latter, let h represent the hidden dimension and $l \in \mathbb{R}^h$ denote the (IA)³ task-specific vector: it can be shown that (IA)³ is equivalent to $\boldsymbol{\theta}_t \equiv \boldsymbol{\theta}_0 + \boldsymbol{\tau}_t$ with $\boldsymbol{\tau}_t = \boldsymbol{\theta}_0 \odot ((l - \mathbf{1}_h) \otimes \mathbf{1}_h)$, where $\mathbf{1}_h$ is the all-ones vector. For each fine-tuning strategy, we initialize their parameters so that the resulting task vector $\boldsymbol{\tau}_t$ starts as a null vector at the beginning.

Computational analysis. As outlined in App. D, by treating the composed model as a cumulative average of individual models, both training/inference stages of ITA/IEL maintain constant complexity $\mathcal{O}(1)$ with respect to the number of tasks T . Indeed, a single forward pass is sufficient to compute the output of the composed model (**constant time**). Moreover, we do not need to store a separate set of weights for each task (**constant memory**), provided we are not interested in more complex forms of composition than the simplest uniform average (as required for specialization and unlearning).

4 RELATION WITH EXISTING WORKS

Task arithmetic. Standard and Parameter-Efficient (PEFT) fine-tuning have been shown to support addition/subtraction of task vectors. However, while the evidence in (Zhang et al., 2024; Ilharco et al., 2022) is primarily empirical, we derive theoretical insights about the pre-conditions for task arithmetic, emphasizing the importance of staying close to the pre-training basin. In this respect, our derivations ground previous findings regarding the efficacy of low learning rates (Ilharco et al., 2022; Ortiz-Jimenez et al., 2024). Remarkably, staying within the pre-training basin has also been proved beneficial in (Sadrtadinov et al., 2024) for ensemble learning. The conditions for compositionality are also studied by Ortiz-Jimenez et al. (2024) on *linearized* networks (Eq. 7). Albeit considering this work inspirational, we see the pros of task arithmetic in the non-linear regime. Firstly, non-linear models surpass their linearized counterparts in single-task accuracy, making them more attractive. To reduce the gap, linearization-aware fine-tuning has to be used (Liu & Soatto, 2023), contrarily to our approach that is compatible with the prevalent fine-tuning techniques. Secondly, linearized inference requires the demanding Jacobian-vector product (three times slower than a forward pass).

Ensemble learning. While original model soups (Wortsman et al., 2022) combine multiple weights fine-tuned on the same dataset, we herein managed to unlock *running model soups* in cross-dataset incremental settings, with possible returns in terms of forward transfer (Lopez-Paz & Ranzato, 2017). The optimization of the whole deep ensemble is also discussed in Jeffares et al. (2024) for standard ensembles, *i.e.*, averaging the outputs of different models. Their derivation is similar to ours and regards the decomposition of the ensemble loss into the strength of the individual learners and their diversity. However, Jeffares et al. (2024) use this result to elucidate the shortcomings of jointly trained deep ensembles, whereas we leverage it to provide effective regularization for model soups in incremental scenarios. Similar to our IEL, several works (Li et al., 2022; 2023a; Schmidt et al., 2023) build an ensemble through a cumulative mean of intermediate checkpoints sampled along the training trajectory, with benefits in terms of generalization and preservation of zero-shot pre-training capabilities (Zheng et al., 2023). Differently, Jolicoeur-Martineau et al. (2023) maintain a population of models trained with varying configurations (*e.g.*, data augmentations). They also gradually push each weight toward the population average, thus encouraging alignment across individual models. Notably, such a behaviour is also positively rewarded by our second-order formulation (see Eq. 17).

Incremental learning. Adding new knowledge in deep networks often degrades the ability on earlier tasks. Approaches against forgetting can be divided into three groups: those involving *regularization* (Kirkpatrick et al., 2017; Zenke et al., 2017), those retaining old data for *rehearsal* (Lopez-Paz & Ranzato, 2017; Aljundi et al., 2019), and those allocating new modules (Mallya & Lazebnik, 2018; Abati et al., 2020). Among the latter, SEED (Rypešć et al., 2024) manages an ensemble of expert networks learned incrementally, sharing similarities with our IEL. However, SEED stores separate models and combines their outputs, whereas our approach maintains a cumulative weight average of past experts, minimizing memory and computational overhead. This running average can also be achieved by InfLoRA (Liang & Li, 2024), which addresses interference among LoRA modules through tailored initialization. In contrast, our method addresses this issue through explicit regularization. In addition, a recent trend capitalizes on prompt-tuning (Wang et al., 2022a;b; Smith et al., 2023), devising a pool of learnable prompts. Notably, the extent to which these models support compositionality is investigated in Perera et al. (2023). Bowman et al. (2023) propose À-la-carte Prompt Tuning (APT), an attention mechanism that enables the creation of bespoke models by composing arbitrary prompts. Their setting called *à-la-carte* aligns with the scope of our ITA; however, our work extends to a broader range of PEFT modules, unlike APT that is limited to soft prompts.

NTK-based Incremental learning. The authors of Tangent Model Composition (TMC) (Liu & Soatto, 2023) build on the foundational work of Ortiz-Jimenez et al. (2024) to address incremental learning. They enforce task arithmetic across subsequent tasks through linearization-aware fine-tuning, which entails a first-order Taylor approximation of the output function around θ_0 . In this context, each task of TMC is effectively equivalent to training a kernel predictor using the **Neural Tangent Kernel (NTK)** (Jacot et al., 2018), defined as $k_{\text{NTK}}(x, x') = \nabla_{\theta} f(x; \theta_0)^{\top} \nabla_{\theta} f(x'; \theta_0)$. In contrast, our approach employs a higher-order approximation of the empirical risk, thereby focusing on the geometry of the loss landscape (Chaudhry et al., 2018). Notably, other recent works (Liu et al., 2024) have adopted the NTK framework to tackle incremental learning: *e.g.*, TKIL (Xiang & Shlizerman, 2023) exploit the NTK formulation to align representations for current and past tasks.

Table 1: Comparison with SOTA (Final Accuracy [\uparrow]). Best results in **bold**, second-best underlined. EWC, LwF-MC, DER++ (buffer size of 1, 000 examples), SEED, and TMC rely on full fine-tuning; L2P, CODA, and APT utilize prompt-based learning. Finally, InfLoRA adopts LoRA fine-tuning.

Model	IN-R	C-100	CUB	Caltech	MIT	RESISC	CropDis.
Joint	86.08	91.74	87.12	93.21	87.19	96.79	99.71
Finetune	22.31	21.57	29.35	44.03	18.85	12.90	20.54
EWC	58.64	73.49	40.33	73.23	64.44	58.80	70.33
LWF-MC	50.93	72.16	21.88	79.73	67.04	68.09	78.28
DER++	60.53	83.02	76.10	86.11	68.88	67.23	98.77
L2P	67.17	87.32	78.95	91.22	83.17	63.47	75.18
CODA	74.12	86.48	78.54	90.57	77.73	69.50	74.65
SEED	55.87	83.39	85.35	90.04	<u>86.34</u>	74.81	92.77
APT	65.32	86.19	69.51	87.71	75.83	49.99	59.37
InfLoRA	76.97	87.17	79.14	90.53	79.14	79.92	89.05
TMC	60.01	78.42	71.72	82.30	68.66	60.66	66.56
ITA-FFT	76.43	89.38	84.80	92.32	85.35	80.50	91.81
ITA-LoRA	77.79	<u>89.96</u>	<u>85.55</u>	92.65	86.60	82.00	95.85
ITA-(IA) ³	77.04	90.66	85.67	<u>92.67</u>	84.74	83.73	95.41
IEL-FFT	80.09	89.38	84.89	92.23	82.79	81.42	95.83
IEL-LoRA	<u>79.93</u>	89.53	84.95	92.19	84.49	<u>82.53</u>	<u>95.88</u>
IEL-(IA) ³	77.86	89.72	84.57	92.70	85.54	81.50	95.68

5 EXPERIMENTS

Datasets. Following affine works (Wang et al., 2022b; Bowman et al., 2023; Liu & Soatto, 2023), we evaluate on these **class-incremental** benchmarks: Split ImageNet-R (Hendrycks et al., 2021) (10 tasks \times 20 classes each), Split CIFAR-100 (Krizhevsky et al., 2009) (10 tasks \times 100 classes), Split CUB-200 (Wah et al., 2011) (10 tasks \times 200 classes), Split Caltech-256 (Griffin et al., 2007) (10 tasks, as in (Liu & Soatto, 2023)), and Split MIT-67 (Quattoni & Torralba, 2009) (10 tasks, as in (Liu & Soatto, 2023)). We conduct further tests on the **aerial** and **medical** domains using Split RESISC45 (Cheng et al., 2017) (9 tasks \times 5 classes) and Split CropDiseases (Hughes et al., 2015) (7 tasks \times 5 classes). They provide a challenging benchmark for our proposals due to their **low domain similarity** with the ImageNet pre-training (Oh et al., 2022). Further details can be found in App. G.

Benchmarking. We compare against recognized incremental methods as EWC (Kirkpatrick et al., 2017), LwF-MC (Rebuffi et al., 2017), DER++ (Buzzega et al., 2020b), L2P (Wang et al., 2022b), and CODA (Smith et al., 2023). We also assess SEED (Rypešć et al., 2024), InfLoRA (Liang & Li, 2024), APT (Bowman et al., 2023), and TMC (Liu & Soatto, 2023), four approaches featuring compositionality and ensemble learning. Their description and the main differences from our approaches are provided in App. F. All methods, including ours, utilize the **same backbone** – a ViT-B/16 (Dosovitskiy et al., 2020) with supervised pre-training on ImageNet21K (Ridnik et al., 2021) – and the same batch size (128). We compute the accuracy on all classes at the end of the final task (Final Accuracy, FA). Following Buzzega et al. (2020b), the **hyperparameters** are chosen through a grid search on a validation set (*i.e.*, 10% of the training set). This procedure was repeated for each method, ensuring careful tuning of the search space. See the supplementary for 1) the **standard deviation** of the FA (the results are averaged over three runs, see App. H.2); 2) the results expressed as Final Forgetting (Chaudhry et al., 2018) (App. H.1); 3) the selected hyper-parameters (App. I).

Comparison with SOTA. As reported in Tab. 1, ITA and IEL outperform existing approaches on all datasets except MIT-67 and CropDisease. At times, they match SEED; however, its reliance on Mixture of Experts makes its inference computationally demanding. Our results far exceed those of TMC, showcasing the potential of the non-linear regime over linearization. Considering the good results on RESISC and CropDis., ITA and IEL do not seem affected by large domain shifts, indicating that our formulation remains effective even when the pre-training optimality is challenged. Finally, given the comparable results of ITA and IEL, we recommend ITA as the preferred starting point for future research, in light of the greater flexibility offered by the individual training paradigm.

Table 2: For ITA, analysis of the impact of the proposed regularization loss (FA \uparrow).

Model	IN-R	C-100	CUB	Caltech	MIT	RESISC	CropDis.
ITA-FFT (reg)	76.43	89.38	84.80	92.32	85.35	80.50	91.81
without Eq. 12 reg.	8.61	17.59	10.47	12.76	12.01	17.17	20.64
Eq. 12 only on CLS	76.00	87.60	83.54	91.04	81.45	75.26	77.00
ITA-LoRA (reg)	77.79	89.96	85.55	92.65	86.60	82.00	95.85
without Eq. 12 reg.	50.17	66.58	60.58	74.87	52.74	37.59	55.86
Eq. 12 only on CLS	77.33	90.03	85.55	92.59	84.86	80.64	96.22

On the impact of regularization. Tab. 2 reports the results of ITA when the regularization term in Eq. 12 is removed (the same analysis for ITA- (IA)³ and IEL is available in App. H). To evaluate the effect on distinct layers, we additionally assess the model with only the last classification layer regularized (see Eq. 12 only on CLS in Tab. 2). As observed: *i*) applying Eq. 12 is beneficial for all examined fine-tuning strategies; *ii*) although regularizing all layers is the most consistent approach, applying Eq. 12 only on the classification head already yields good accuracy. This suggests that compositionality in closed-set models is largely dependent on the learning dynamics of the last layer. Finally, full fine-tuning (FFT) struggles the most when no regularization or partial regularization is applied, in contrast to PEFT modules like LoRA and (IA)³ that still manage to achieve decent results. We ascribe this evidence to the tendency of PEFT modules to forget less of the pre-trained knowledge (Biderman et al., 2024), a feature we identify as beneficial for model compositionality.

To gain insights into our regularization, we revisit the foundations of ITA: specifically the inequality in Eq. 5, which states that the risk of the composed model $\hat{\ell}_{\text{cur}}(\theta_{\mathcal{P}})$ is upper bounded by the weighted risk of the individual models $\sum w_t \hat{\ell}_{\text{cur}}(\theta_t)$. In practice, there is no guarantee that this upper bound is tight: as these individual models are trained on disjoint tasks, their risk is likely high for examples outside their training distribution. We hence expect the upper bound to be loose: this is shown in Fig. 1, where we measure the effect of our regularization on the empirical risk of the composed model $\hat{\ell}(\theta_{\mathcal{P}})$ and on the upper bound $\sum w_t \hat{\ell}(\theta_t)$.

From this lens, ITA significantly tightens the upper bound, with a remarkable reduction of the risk $\hat{\ell}(\theta_{\mathcal{P}})$ associated with the composed model $\hat{\ell}_{\text{cur}}(\theta_{\mathcal{P}})$. Notably, the latter consistently outperforms the pre-trained model $\hat{\ell}_{\text{cur}}(\theta_0)$ (obtained with linear probing). This indicates that there is a good margin to find a point in parameter space that improves pre-training. Vice versa, the upper bound $\sum w_t \hat{\ell}_{\text{cur}}(\theta_t)$ does not always exceed pre-training, even with regularization applied. This raises the theoretical possibility of *regret* between the composed model and pre-training. This work provides empirical evidence to address it but a theoretical analysis is needed, which we leave for future works.

On the compositional skills of ITA and IEL. We herein investigate *specialization* and *unlearning* capabilities of our approaches. Specifically, given the task vectors trained on the sequence of tasks, we freeze and use them to edit the composed model. As these adaptations involve simple additions and subtractions of task vectors, we can perform *specialization* and *unlearning* in a zero-shot fashion requiring no examples. In particular, through the former test on specialization, we examine whether composing only a subset of the task vectors can boost the composed model on the respective selected tasks. In doing so, we focus on three tasks of the incremental sequence *i.e.*, the first, the central, and the last one. The performance is then measured considering the average FA across the tasks we wish to specialize in (FA_{TGT}) and the average FA across the others (FA_{CTRL}). In the second test, instead, our goal is to eliminate knowledge of a specific task without degrading performance on other tasks (*unlearning*). Specifically, to unlearn a given task, we subtract its associated task vector from the composition, which still includes the weights of the tasks that need to be preserved. Afterwards: *i*) we measure the accuracy for the unlearned task (FA_{TGT}) and the others (FA_{CTRL}); *ii*) the procedure is then repeated for all tasks, with the performance averaged across them.

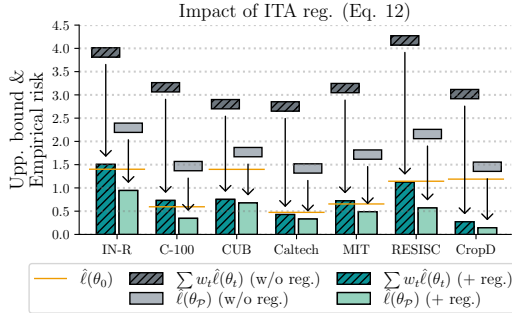


Figure 1: Effect of ITA. Best viewed in color.

Table 3: Analysis of compositional capabilities – see App. H for results on other datasets. In parentheses, we report the gain (or loss) in accuracy on the target task.

Dataset	Model	<i>zero-shot specialization</i>		<i>zero-shot unlearning</i>	
		FA _{TGT} [↑]	FA _{CTRL}	FA _{TGT} [↓]	FA _{CTRL}
IN-R	ITA-LoRA	80.83 (+11.40)	50.52	22.77 (−55.02)	52.72 (−25.07)
	IEL-LoRA	73.46 (−06.68)	38.46	18.55 (−61.38)	41.97 (−37.96)
	TMC	69.93 (+08.36)	34.08	45.77 (−14.24)	54.37 (−05.64)
C-100	ITA-LoRA	92.80 (+01.63)	60.06	28.67 (−61.29)	71.96 (−17.99)
	IEL-LoRA	77.77 (−13.22)	37.90	19.48 (−70.05)	56.52 (−33.01)
	TMC	87.53 (+06.49)	45.75	55.63 (−22.79)	71.83 (−06.59)

The results are in Tab. 3 and Tab. 7 of App. H, where we compare with TMC by Liu & Soatto (2023), *i.e.*, model linearization and tangent fine-tuning. Regarding **specialization**, both ITA and TMC improve the accuracy on the target tasks. IEL, instead, leads to a severe drop during specialization. This yields an interesting finding: while Tab. 1 highlights its performance as satisfying, we observe that the ensemble struggles when any of its members are removed. In an era where ensemble models are experiencing renewed attention (Liang & Li, 2024; Marouf et al., 2023), this evidence should be sobering. On the other hand, we underline the specialization capabilities of ITA, with the best absolute performance on the target tasks. We ascribe it to the devised regularization objective and the natural advantages of the non-linear regime. When focusing on **unlearning**, ITA shows mixed results. Notably, both ITA and TMC consistently reduce the performance of the task to be forgotten. While this also affects the tasks that should be preserved, ITA exhibits a greater distinction between the target and control tasks, indicating a more effective disentanglement of knowledge across tasks.

6 DISCUSSION OF LIMITATIONS AND FUTURE DIRECTIONS

Second-order Taylor approximations are valid tools to make the theory more tractable. For instance, a similar approach was used by Mirzadeh et al. (2020) to support the importance of wider local minima against forgetting. Nonetheless, these approximations often rely on certain concessions, the first regarding their validity during training. Since our approximation is performed at the pre-training weights, it may become inaccurate as parameters drift. While importance-based terms like ours may partially mitigate the drift, other techniques could be used (*e.g.*, a L2-norm penalty on task vectors or low learning rate). Even without explicit countermeasures, existing studies suggest that deep networks often fall into a regime where the loss tends to be almost convex in a neighborhood around their local minima (Goodfellow et al., 2015; Lucas et al., 2021; Yunis et al., 2022). Also, under some conditions, their training dynamics can enter the lazy regime (Chizat et al., 2019; Jacot et al., 2018), where these models rapidly achieve near-zero loss with minimal changes to their weights.

Moreover, two other concessions warrant a discussion: *approx. i*) we implicitly model the weight distribution with a Gaussian (Chaudhry et al., 2018) (Eq. 9); *approx. ii*) the Fisher matrix is approximated by its diagonal (Eq. 11). Regarding *approx. i*), while non-Gaussian posteriors are often cumbersome to manage, Farquhar et al. (2020) challenge the common belief that mean-field approximations are overly restrictive for deep networks. Moreover, although *approx. ii*) may seem crude, the diagonal approximation is efficient, with a low memory footprint. Still, our approach could profit from more accurate estimations of the Hessian, like the Kronecker factored approximation (Martens & Grosse, 2015), which considers the interactions between weights of the same layer.

Future work. Through theoretical and empirical analyses, we support the importance of remaining within the pre-training basin to achieve composable deep networks. However, there is still much to explore along this path. We mainly focus on closed-set classification models but it would be noteworthy to extend our analysis to self-supervised pre-training and open-vocabulary models like CLIP. Indeed, recent works (Zheng et al., 2023) have shown their tendency to forget zero-shot pre-training capabilities while fine-tuning. In this respect, our second-order regularization could significantly aid compositionality in CLIP-based models. Finally, we believe that staying within the pre-train basin is only one aspect to consider; future research should emphasize the **exploration** of the pre-train basin, to achieve composable modules with a higher degree of specialization on their respective tasks.

REFERENCES

- 540
541
542 Davide Abati, Jakub Tomczak, Tijmen Blankevoort, Simone Calderara, Rita Cucchiara, and
543 Babak Ehteshami Bejnordi. Conditional channel gated networks for task-aware continual learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2020.
544
- 545 Rahaf Aljundi, Eugene Belilovsky, Tinne Tuytelaars, Laurent Charlin, Massimo Caccia, Min Lin,
546 and Lucas Page-Caccia. Online continual learning with maximal interfered retrieval. In *Advances
547 in Neural Information Processing Systems*, 2019.
548
- 549 Nader Asadi, Mahdi Beitollahi, Yasser Khalil, Yinchuan Li, Guojun Zhang, and Xi Chen. Does
550 combining parameter-efficient modules improve few-shot transfer accuracy? *arXiv preprint
551 arXiv:2402.15414*, 2024.
- 552 Dan Biderman, Jose Gonzalez Ortiz, Jacob Portes, Mansheej Paul, Philip Greengard, Connor Jen-
553 nings, Daniel King, Sam Havens, Vitaliy Chiley, Jonathan Frankle, et al. Lora learns less and
554 forgets less. *arXiv preprint arXiv:2405.09673*, 2024.
555
- 556 Matteo Boschini, Lorenzo Bonicelli, Pietro Buzzega, Angelo Porrello, and Simone Calderara. Class-
557 incremental continual learning into the extended der-verse. *IEEE Transactions on Pattern Analy-
558 sis and Machine Intelligence*, 2022.
- 559 Benjamin Bowman, Alessandro Achille, Luca Zancato, Matthew Trager, Pramuditha Perera, Gio-
560 vanni Paolini, and Stefano Soatto. A-la-carte prompt tuning (apt): Combining distinct data via
561 composable prompting. In *Proceedings of the IEEE conference on Computer Vision and Pattern
562 Recognition*, 2023.
563
- 564 Pietro Buzzega, Matteo Boschini, Lorenzo Bonicelli, and Angelo Porrello. Mammoth - an ex-
565 tendible (general) continual learning framework for pytorch, 2020a. URL <https://github.com/aimagelab/mammoth>.
566
- 567 Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark expe-
568 rience for general continual learning: a strong, simple baseline. *Advances in Neural Information
569 Processing Systems*, 2020b.
570
- 571 Lucas Caccia, Rahaf Aljundi, Nader Asadi, Tinne Tuytelaars, Joelle Pineau, and Eugene Belilovsky.
572 New insights on reducing abrupt representation change in online continual learning. In *Internat-
573 ional Conference on Learning Representations*, 2022.
- 574 Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr. Riemannian
575 walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the
576 European conference on computer vision (ECCV)*, pp. 532–547, 2018.
577
- 578 Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Bench-
579 mark and state of the art. *Proceedings of the IEEE*, 105(10):1865–1883, 2017.
- 580 Lenaic Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming.
581 *Advances in Neural Information Processing Systems*, 2019.
582
- 583 Yin Cui, Yang Song, Chen Sun, Andrew Howard, and Serge Belongie. Large scale fine-grained
584 categorization and domain-specific transfer learning. In *Proceedings of the IEEE conference on
585 Computer Vision and Pattern Recognition*, pp. 4109–4118, 2018.
586
- 587 Shachar Don-Yehiya, Elad Venezian, Colin Raffel, Noam Slonim, and Leshem Choshen. CoLD
588 fusion: Collaborative descent for distributed multitask finetuning. In *Proceedings of the 61st
589 Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp.
590 788–806, Toronto, Canada, 2023. Association for Computational Linguistics.
- 591 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
592 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An
593 image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint
arXiv:2010.11929*, 2020.

- 594 Sebastian Farquhar, Lewis Smith, and Yarin Gal. Liberty or depth: Deep bayesian neural nets do
595 not need complex weight posterior approximations. *Advances in Neural Information Processing*
596 *Systems*, 2020.
- 597 Thomas George. NNGeometry: Easy and Fast Fisher Information Matrices and Neural Tangent
598 Kernels in PyTorch, 2021. URL <https://doi.org/10.5281/zenodo.4532597>.
- 600 Ian J Goodfellow, Oriol Vinyals, and Andrew M Saxe. Qualitatively characterizing neural network
601 optimization problems. In *International Conference on Learning Representations*, 2015.
- 602 Gregory Griffin, Alex Holub, and Pietro Perona. Caltech-256 object category dataset. *CalTech*
603 *Report*, 2007.
- 605 Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul
606 Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer.
607 The many faces of robustness: A critical analysis of out-of-distribution generalization. *IEEE*
608 *International Conference on Computer Vision*, 2021.
- 610 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
611 and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint*
612 *arXiv:2106.09685*, 2021.
- 613 Chengsong Huang, Qian Liu, Bill Yuchen Lin, Tianyu Pang, Chao Du, and Min Lin. Lorahub: Effi-
614 cient cross-task generalization via dynamic lora composition. *arXiv preprint arXiv:2307.13269*,
615 2023.
- 616 David Hughes, Marcel Salathé, et al. An open access repository of images on plant health to enable
617 the development of mobile disease diagnostics. *arXiv preprint arXiv:1511.08060*, 2015.
- 619 Ferenc Huszár. Note on the quadratic penalties in elastic weight consolidation. *Proceedings of the*
620 *National Academy of Sciences*, 115(11):E2496–E2497, 2018.
- 622 Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt,
623 Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. *arXiv preprint*
624 *arXiv:2212.04089*, 2022.
- 625 Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and gen-
626 eralization in neural networks. *Advances in Neural Information Processing Systems*, 2018.
- 627 Alan Jeffares, Tennison Liu, Jonathan Crabbé, and Mihaela van der Schaar. Joint training of deep
628 ensembles fails due to learner collusion. *Advances in Neural Information Processing Systems*,
629 2024.
- 631 Xisen Jin, Xiang Ren, Daniel Preotiuc-Pietro, and Pengxiang Cheng. Dataless knowledge fusion by
632 merging weights of language models. *arXiv preprint arXiv:2212.09849*, 2022.
- 633 Alexia Jolicoeur-Martineau, Emy Gervais, Kilian Fatras, Yan Zhang, and Simon Lacoste-Julien.
634 Population parameter averaging (papa). *arXiv preprint arXiv:2304.03094*, 2023.
- 636 James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A
637 Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcom-
638 ing catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*,
639 114(13):3521–3526, 2017.
- 640 Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images.
641 *Master’s thesis, Department of Computer Science, University of Toronto*, 2009.
- 642 Frederik Kunstner, Philipp Hennig, and Lukas Balles. Limitations of the empirical fisher approx-
643 imation for natural gradient descent. *Advances in neural information processing systems*, 32,
644 2019.
- 645 John M Lee. *Riemannian manifolds: an introduction to curvature*, volume 176. Springer Science
646 & Business Media, 2006.

- 648 Tao Li, Zhehao Huang, Qinghua Tao, Yingwen Wu, and Xiaolin Huang. Trainable weight averaging:
649 Efficient training by optimizing historical solutions. In *The Eleventh International Conference on*
650 *Learning Representations*, 2022.
- 651 Tao Li, Zhehao Huang, Qinghua Tao, Yingwen Wu, and Xiaolin Huang. Trainable weight averaging:
652 A general approach for subspace training. *arXiv preprint arXiv:2205.13104*, 2023a.
- 653
654 Weishi Li, Yong Peng, Miao Zhang, Liang Ding, Han Hu, and Li Shen. Deep model fusion: A
655 survey. *arXiv preprint arXiv:2309.15698*, 2023b.
- 656
657 Yan-Shuo Liang and Wu-Jun Li. Inflora: Interference-free low-rank adaptation for continual learn-
658 ing. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2024.
- 659 Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohta, Tenghao Huang, Mohit Bansal, and
660 Colin A Raffel. Few-shot parameter-efficient fine-tuning is better and cheaper than in-context
661 learning. *Advances in Neural Information Processing Systems*, 35:1950–1965, 2022.
- 662
663 Jingren Liu, Zhong Ji, YunLong Yu, Jiale Cao, Yanwei Pang, Jungong Han, and Xuelong Li.
664 Parameter-efficient fine-tuning for continual learning: A neural tangent kernel perspective. *arXiv*
665 *preprint arXiv:2407.17120*, 2024.
- 666
667 Tian Yu Liu and Stefano Soatto. Tangent model composition for ensembling and continual fine-
668 tuning. In *IEEE International Conference on Computer Vision*, 2023.
- 669 David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning.
670 *Advances in neural information processing systems*, 30, 2017.
- 671 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint*
672 *arXiv:1711.05101*, 2017.
- 673
674 James Lucas, Juhan Bae, Michael R Zhang, Stanislav Fort, Richard Zemel, and Roger Grosse.
675 Analyzing monotonic linear interpolation in neural network loss landscapes. In *International*
676 *Conference on Machine Learning*, 2021.
- 677 Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative
678 pruning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*,
679 2018.
- 680 Imad Eddine Marouf, Subhankar Roy, Enzo Tartaglione, and Stéphane Lathuilière. Weighted en-
681 semble models are strong continual learners. *arXiv preprint arXiv:2312.08977*, 2023.
- 682
683 James Martens. New insights and perspectives on the natural gradient method. *Journal of Machine*
684 *Learning Research*, 21(146):1–76, 2020.
- 685
686 James Martens and Roger Grosse. Optimizing neural networks with kronecker-factored approximate
687 curvature. In *International Conference on Machine Learning*, 2015.
- 688 Michael S Matena and Colin A Raffel. Merging models with fisher-weighted averaging. *Advances*
689 *in Neural Information Processing Systems*, 2022.
- 690
691 Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The
692 sequential learning problem. In *Psychology of learning and motivation*. Elsevier, 1989.
- 693
694 Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agueria y Arcas.
695 Communication-efficient learning of deep networks from decentralized data. In *International*
696 *Conference on Artificial Intelligence and Statistics*, pp. 1273–1282. PMLR, 2017.
- 697
698 Seyed Iman Mirzadeh, Mehrdad Farajtabar, Razvan Pascanu, and Hassan Ghasemzadeh. Under-
699 standing the role of training regimes in continual learning. *Advances in Neural Information Pro-*
700 *cessing Systems*, 2020.
- 701
702 Jaehoon Oh, Sungnyun Kim, Namgyu Ho, Jin-Hwa Kim, Hwanjun Song, and Se-Young Yun. Un-
703 derstanding cross-domain few-shot learning based on domain similarity and few-shot difficulty.
704 *Advances in Neural Information Processing Systems*, 35:2622–2636, 2022.

- 702 Guillermo Ortiz-Jimenez, Alessandro Favero, and Pascal Frossard. Task arithmetic in the tangent
703 space: Improved editing of pre-trained models. *Advances in Neural Information Processing Sys-*
704 *tems*, 2024.
- 705
706 Oleksiy Ostapenko, Timothee Lesort, Pau Rodriguez, Md Rifat Arefin, Arthur Douillard, Irina Rish,
707 and Laurent Charlin. Continual learning with foundation models: An empirical study of latent
708 replay. In *Conference on lifelong learning agents*, pp. 60–91. PMLR, 2022.
- 709 Pramuditha Perera, Matthew Trager, Luca Zancato, Alessandro Achille, and Stefano Soatto. Prompt
710 algebra for task composition. *arXiv preprint arXiv:2306.00310*, 2023.
- 711
712 Jonas Pfeiffer, Sebastian Ruder, Ivan Vulić, and Edoardo Maria Ponti. Modular deep learning. *arXiv*
713 *preprint arXiv:2302.11529*, 2023.
- 714 Ariadna Quattoni and Antonio Torralba. Recognizing indoor scenes. In *Proceedings of the IEEE*
715 *conference on Computer Vision and Pattern Recognition*. IEEE, 2009.
- 716
717 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
718 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
719 models from natural language supervision. In *International conference on machine learning*, pp.
720 8748–8763. PMLR, 2021.
- 721 Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl:
722 Incremental classifier and representation learning. In *Proceedings of the IEEE conference on*
723 *Computer Vision and Pattern Recognition*, 2017.
- 724
725 Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik-Manor. Imagenet-21k pretraining for
726 the masses. *arXiv preprint arXiv:2104.10972*, 2021.
- 727 Grzegorz Rypeś, Sebastian Cygert, Valeriya Khan, Tomasz Trzcinski, Bartosz Michał Zieliński,
728 and Bartłomiej Twardowski. Divide and not forget: Ensemble of selectively trained experts in
729 continual learning. In *The Twelfth International Conference on Learning Representations*, 2024.
- 730
731 Ildus Sadrtidinov, Dmitrii Pozdeev, Dmitry P Vetrov, and Ekaterina Lobacheva. To stay or not to stay
732 in the pre-train basin: Insights on ensembling in transfer learning. *Advances in Neural Information*
733 *Processing Systems*, 2024.
- 734 Fabian David Schmidt, Ivan Vulić, and Goran Glavaš. Free lunch: Robust cross-lingual transfer
735 via model checkpoint averaging. In *Proceedings of the 61st Annual Meeting of the Association*
736 *for Computational Linguistics (Volume 1: Long Papers)*, pp. 5712–5730, Toronto, Canada, 2023.
737 Association for Computational Linguistics.
- 738
739 Jonathan Schwarz, Wojciech Czarnecki, Jelena Luketina, Agnieszka Grabska-Barwinska, Yee Whye
740 Teh, Razvan Pascanu, and Raia Hadsell. Progress & compress: A scalable framework for contin-
741 ual learning. In *International conference on machine learning*, pp. 4528–4537. PMLR, 2018.
- 742 James Seale Smith, Leonid Karlinsky, Vyshnavi Gutta, Paola Cascante-Bonilla, Donghyun Kim,
743 Assaf Arbelle, Rameswar Panda, Rogerio Feris, and Zsolt Kira. Coda-prompt: Continual de-
744 composed attention-based prompting for rehearsal-free continual learning. In *Proceedings of the*
745 *IEEE conference on Computer Vision and Pattern Recognition*, 2023.
- 746
747 George Stoica, Daniel Bolya, Jakob Bjorner, Taylor Hearn, and Judy Hoffman. Zipit! merging
748 models from different tasks without training. *arXiv preprint arXiv:2305.03053*, 2023.
- 749 Gido M Van de Ven and Andreas S Tolias. Three scenarios for continual learning. *arXiv preprint*
750 *arXiv:1904.07734*, 2019.
- 751
752 C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The caltech-ucsd birds-200-2011
753 dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.
- 754
755 Zifeng Wang, Zizhao Zhang, Sayna Ebrahimi, Ruoxi Sun, Han Zhang, Chen-Yu Lee, Xiaoqi Ren,
Guolong Su, Vincent Perot, Jennifer Dy, et al. Dualprompt: Complementary prompting for
rehearsal-free continual learning. In *European Conference on Computer Vision*, 2022a.

- 756 Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vin-
757 cent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. In *Pro-*
758 *ceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 139–149,
759 2022b.
- 760 Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes,
761 Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. Model
762 soups: averaging weights of multiple fine-tuned models improves accuracy without increasing
763 inference time. In *International Conference on Machine Learning*. PMLR, 2022.
- 764 Jinlin Xiang and Eli Shlizerman. Tkil: Tangent kernel optimization for class balanced incremental
765 learning. In *IEEE International Conference on Computer Vision*, pp. 3529–3539, 2023.
- 766 Prateek Yadav, Colin Raffel, Mohammed Muqeeth, Lucas Caccia, Haokun Liu, Tianlong Chen,
767 Mohit Bansal, Leshem Choshen, and Alessandro Sordoni. A survey on model moerging:
768 Recycling and routing among specialized experts for collaborative learning. *arXiv preprint*
769 *arXiv:2408.07057*, 2024a.
- 770 Prateek Yadav, Derek Tam, Leshem Choshen, Colin A Raffel, and Mohit Bansal. Ties-merging: Re-
771 solving interference when merging models. *Advances in Neural Information Processing Systems*,
772 2024b.
- 773 Prateek Yadav, Tu Vu, Jonathan Lai, Alexandra Chronopoulou, Manaal Faruqui, Mohit Bansal,
774 and Tsendsuren Munkhdalai. What matters for model merging at scale? *arXiv preprint*
775 *arXiv:2410.03617*, 2024c.
- 776 David Yunis, Kumar Kshitij Patel, Pedro Henrique Pamplona Savarese, Gal Vardi, Jonathan Frankle,
777 Matthew Walter, Karen Livescu, and Michael Maire. On convexity and linear mode connectivity
778 in neural networks. In *OPT: Optimization for Machine Learning (NeurIPS 2022 Workshop)*, 2022.
- 779 Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence.
780 In *International Conference on Machine Learning*, 2017.
- 781 Jianyu Zhang and Léon Bottou. Learning useful representations for shifting tasks and distributions.
782 In *International Conference on Machine Learning*. PMLR, 2023.
- 783 Jinghan Zhang, Junteng Liu, Junxian He, et al. Composing parameter-efficient modules with arith-
784 metic operation. *Advances in Neural Information Processing Systems*, 2024.
- 785 Zangwei Zheng, Mingyuan Ma, Kai Wang, Ziheng Qin, Xiangyu Yue, and Yang You. Preventing
786 zero-shot transfer degradation in continual learning of vision-language models. In *IEEE Interna-*
787 *tional Conference on Computer Vision*, 2023.
- 788 Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong,
789 and Qing He. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1):
790 43–76, 2020.
- 791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

A APPENDIX / SUPPLEMENTAL MATERIAL

The appendix is organized via the following contributions:

- App. B provides the proofs for Theorem 1 and Eq. 17.
- App. C illustrates that gradients of Eq. 19 can be derived in closed-form.
- App. D reports the computational analysis for our approaches Incremental Task Arithmetic (ITA) and Incremental Ensemble Learning (IEL).
- App. E describes the essential details required to reproduce the methods described in the main paper.
- Apps. F and G supplements the main paper with additional information regarding the relation with existing works and the adopted benchmarks.
- App. H provides additional experimental results: the forgetting metric, the standard deviation measured in each experiment, an additional ablation study regarding IEL, an extensive evaluation of model compositionality across multiple datasets, an analysis of the similarities between task vectors learned by ITA and IEL, and a timing analysis.
- App. I outlines the hyperparameter configurations adopted in our experiments.

B PROOFS

B.1 PROOF OF THEOREM 1

Proof. For the sake of notation, we will use the shortcuts $\ell_0 \triangleq \ell(\boldsymbol{\theta}_0)$, $\mathbf{J}_\ell \triangleq \nabla \ell(\boldsymbol{\theta}_0)$, and $\mathbf{H}_\ell \triangleq \mathbf{H}_\ell(\boldsymbol{\theta}_0)$. Given a setting with only two learners A and B , the second-order approximation of the loss function of the composed model \mathcal{P} is:

$$\begin{aligned}
\ell_{\text{cur}}(\boldsymbol{\theta}_{\mathcal{P}}) &= \ell_0 + (\boldsymbol{\theta}_{\mathcal{P}} - \boldsymbol{\theta}_0)^{\text{T}} \mathbf{J}_\ell + \frac{1}{2} (\boldsymbol{\theta}_{\mathcal{P}} - \boldsymbol{\theta}_0)^{\text{T}} \mathbf{H}_\ell (\boldsymbol{\theta}_{\mathcal{P}} - \boldsymbol{\theta}_0) \\
&= \underbrace{(w_a + w_b)}_{=1} \ell_0 + \underbrace{(\boldsymbol{\theta}_{\mathcal{P}} - \boldsymbol{\theta}_0)^{\text{T}} \mathbf{J}_\ell}_{w_a \boldsymbol{\tau}_a + w_b \boldsymbol{\tau}_b} + \frac{1}{2} (\boldsymbol{\theta}_{\mathcal{P}} - \boldsymbol{\theta}_0)^{\text{T}} \mathbf{H}_\ell (\boldsymbol{\theta}_{\mathcal{P}} - \boldsymbol{\theta}_0) \\
&= w_a \ell_0 + w_b \ell_0 + (w_a \boldsymbol{\tau}_a + w_b \boldsymbol{\tau}_b)^{\text{T}} \mathbf{J}_\ell + \frac{1}{2} (w_a \boldsymbol{\tau}_a + w_b \boldsymbol{\tau}_b)^{\text{T}} \mathbf{H}_\ell (w_a \boldsymbol{\tau}_a + w_b \boldsymbol{\tau}_b) \\
&= w_a \ell_0 + w_b \ell_0 + w_a \boldsymbol{\tau}_a^{\text{T}} \mathbf{J}_\ell + w_b \boldsymbol{\tau}_b^{\text{T}} \mathbf{J}_\ell + \frac{1}{2} (w_a \boldsymbol{\tau}_a^{\text{T}} + w_b \boldsymbol{\tau}_b^{\text{T}}) \mathbf{H}_\ell (w_a \boldsymbol{\tau}_a + w_b \boldsymbol{\tau}_b) \\
&= w_a [\ell_0 + \boldsymbol{\tau}_a^{\text{T}} \mathbf{J}_\ell] + w_b [\ell_0 + \boldsymbol{\tau}_b^{\text{T}} \mathbf{J}_\ell] + \frac{1}{2} (w_a \boldsymbol{\tau}_a^{\text{T}} \mathbf{H}_\ell + w_b \boldsymbol{\tau}_b^{\text{T}} \mathbf{H}_\ell) (w_a \boldsymbol{\tau}_a + w_b \boldsymbol{\tau}_b) \\
&= w_a [\ell_0 + \boldsymbol{\tau}_a^{\text{T}} \mathbf{J}_\ell] + w_b [\ell_0 + \boldsymbol{\tau}_b^{\text{T}} \mathbf{J}_\ell] + \frac{1}{2} (w_a^2 \boldsymbol{\tau}_a^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_a + w_b^2 \boldsymbol{\tau}_b^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_b + \\
&\quad + w_a w_b \boldsymbol{\tau}_a^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_b + w_a w_b \boldsymbol{\tau}_b^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_a)
\end{aligned}$$

Given that $\boldsymbol{\tau}_a^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_b$ is a scalar and $(ABC)^{\text{T}} = (C^{\text{T}} B^{\text{T}} A^{\text{T}})$

$$\begin{aligned}
&\rightarrow \boldsymbol{\tau}_a^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_b = (\boldsymbol{\tau}_a^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_b)^{\text{T}} = \underbrace{(\boldsymbol{\tau}_b^{\text{T}} \mathbf{H}_\ell)^{\text{T}}}_{\text{symm.}} \boldsymbol{\tau}_a = \boldsymbol{\tau}_b^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_a
\end{aligned}$$

$$\begin{aligned}
\ell_{\text{cur}}(\boldsymbol{\theta}_{\mathcal{P}}) &= w_a [\ell_0 + \boldsymbol{\tau}_a^{\text{T}} \mathbf{J}_\ell] + w_b [\ell_0 + \boldsymbol{\tau}_b^{\text{T}} \mathbf{J}_\ell] + \frac{1}{2} (w_a^2 \boldsymbol{\tau}_a^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_a + w_b^2 \boldsymbol{\tau}_b^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_b + 2w_a w_b \boldsymbol{\tau}_a^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_b) \\
&= w_a [\ell_0 + \boldsymbol{\tau}_a^{\text{T}} \mathbf{J}_\ell] + w_b [\ell_0 + \boldsymbol{\tau}_b^{\text{T}} \mathbf{J}_\ell] + \frac{1}{2} \left(\underbrace{w_a^2 \boldsymbol{\tau}_a^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_a}_{\substack{w_a^2 = w_a(1-w_b) \\ = w_a - w_a w_b}} + w_b^2 \boldsymbol{\tau}_b^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_b + 2w_a w_b \boldsymbol{\tau}_a^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_b \right) \\
&= w_a [\ell_0 + \boldsymbol{\tau}_a^{\text{T}} \mathbf{J}_\ell + \frac{1}{2} \boldsymbol{\tau}_a^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_a] + w_b [\ell_0 + \boldsymbol{\tau}_b^{\text{T}} \mathbf{J}_\ell] + \frac{1}{2} (w_a \boldsymbol{\tau}_a^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_a + w_b \boldsymbol{\tau}_b^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_b + \\
&\quad - w_a w_b \boldsymbol{\tau}_a^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_a - w_a w_b \boldsymbol{\tau}_b^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_b + 2w_a w_b \boldsymbol{\tau}_a^{\text{T}} \mathbf{H}_\ell \boldsymbol{\tau}_b)
\end{aligned}$$

$$\begin{aligned}
&= w_a[\ell_0 + \boldsymbol{\tau}_a^T \mathbf{J}_\ell + \frac{1}{2} \boldsymbol{\tau}_a^T \mathbf{H}_\ell \boldsymbol{\tau}_a] + w_b[\ell_0 + \boldsymbol{\tau}_b^T \mathbf{J}_\ell + \frac{1}{2} \boldsymbol{\tau}_b^T \mathbf{H}_\ell \boldsymbol{\tau}_b] + \\
&\quad - \frac{1}{2} w_a w_b (\boldsymbol{\tau}_a^T \mathbf{H}_\ell \boldsymbol{\tau}_a + \boldsymbol{\tau}_b^T \mathbf{H}_\ell \boldsymbol{\tau}_b - 2 \boldsymbol{\tau}_a^T \mathbf{H}_\ell \boldsymbol{\tau}_b) \\
&= w_a \ell_{\text{cur}}(\boldsymbol{\theta}_A) + w_b \ell_{\text{cur}}(\boldsymbol{\theta}_B) - \frac{1}{2} w_a w_b (\boldsymbol{\tau}_a^T \mathbf{H}_\ell \boldsymbol{\tau}_a + \boldsymbol{\tau}_b^T \mathbf{H}_\ell \boldsymbol{\tau}_b - 2 \boldsymbol{\tau}_a^T \mathbf{H}_\ell \boldsymbol{\tau}_b) \\
&= w_a \ell_{\text{cur}}(\boldsymbol{\theta}_A) + w_b \ell_{\text{cur}}(\boldsymbol{\theta}_B) - \frac{1}{2} w_a w_b (\boldsymbol{\tau}_a - \boldsymbol{\tau}_b)^T \mathbf{H}_\ell (\boldsymbol{\tau}_a - \boldsymbol{\tau}_b).
\end{aligned}$$

In the multiple learning setting, we have that $\boldsymbol{\theta}_{\mathcal{P}} = \boldsymbol{\theta}_0 + \boldsymbol{\tau}_{\mathcal{P}} = \boldsymbol{\theta}_0 + \sum_{t=1}^T w_t \boldsymbol{\tau}_t$ with $\sum_{t=1}^T w_t = 1$. Therefore:

$$\begin{aligned}
\ell_{\text{cur}}(\boldsymbol{\theta}_{\mathcal{P}}) &= \ell_0 + (\boldsymbol{\theta}_{\mathcal{P}} - \boldsymbol{\theta}_0)^T \mathbf{J}_\ell + \frac{1}{2} (\boldsymbol{\theta}_{\mathcal{P}} - \boldsymbol{\theta}_0)^T \mathbf{H}_\ell (\boldsymbol{\theta}_{\mathcal{P}} - \boldsymbol{\theta}_0) \\
&= \underbrace{\sum_{t=1}^T w_t \ell_0}_{=1} + \boldsymbol{\tau}_{\mathcal{P}}^T \mathbf{J}_\ell + \frac{1}{2} \boldsymbol{\tau}_{\mathcal{P}}^T \mathbf{H}_\ell \boldsymbol{\tau}_{\mathcal{P}} = \sum_{t=1}^T w_t \ell_0 + \left(\sum_{t=1}^T w_t \boldsymbol{\tau}_t \right)^T \mathbf{J}_\ell + \frac{1}{2} \boldsymbol{\tau}_{\mathcal{P}}^T \mathbf{H}_\ell \boldsymbol{\tau}_{\mathcal{P}} \\
&= \sum_{t=1}^T w_t [\ell_0 + \boldsymbol{\tau}_t^T \mathbf{J}_\ell] + \frac{1}{2} \boldsymbol{\tau}_{\mathcal{P}}^T \mathbf{H}_\ell \boldsymbol{\tau}_{\mathcal{P}}.
\end{aligned}$$

Let us now focus on the quadratic term:

$$\begin{aligned}
\boldsymbol{\tau}_{\mathcal{P}}^T \mathbf{H}_\ell \boldsymbol{\tau}_{\mathcal{P}} &= \left(\sum_{t=1}^T w_t \boldsymbol{\tau}_t \right)^T \mathbf{H}_\ell \left(\sum_{t=1}^T w_t \boldsymbol{\tau}_t \right) = \left(\sum_{t=1}^T w_t \boldsymbol{\tau}_t^T \mathbf{H}_\ell \right) \left(\sum_{t=1}^T w_t \boldsymbol{\tau}_t \right) \\
&= \left(\sum_{t=1}^T w_t \boldsymbol{\tau}_t^T \mathbf{H}_\ell \right) \left(w_t \boldsymbol{\tau}_t + \sum_{t' \neq t} w_{t'} \boldsymbol{\tau}_{t'} \right) = \sum_{t=1}^T w_t^2 \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t + w_t \boldsymbol{\tau}_t^T \mathbf{H}_\ell \sum_{t' \neq t} w_{t'} \boldsymbol{\tau}_{t'} \\
&= \sum_{t=1}^T w_t^2 \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t + \sum_{t=1}^T w_t \boldsymbol{\tau}_t^T \mathbf{H}_\ell \sum_{t' \neq t} w_{t'} \boldsymbol{\tau}_{t'} \\
&= \sum_{t=1}^T w_t^2 \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t + \sum_{t=1}^T \sum_{t' \neq t} w_t w_{t'} \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_{t'} \\
&\text{Since } \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_{t'} = \boldsymbol{\tau}_{t'}^T \mathbf{H}_\ell \boldsymbol{\tau}_t \text{ (symmetry)} \\
&= \sum_{t=1}^T w_t^2 \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t + 2 \sum_{t, t' < t} w_t w_{t'} \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_{t'}
\end{aligned}$$

Given that $w_t^2 = w_t \cdot w_t = w_t (1 - \sum_{t' \neq t} w_{t'}) = w_t - w_t \sum_{t' \neq t} w_{t'}$:

$$\begin{aligned}
&= \sum_{t=1}^T \left(w_t - w_t \sum_{t' \neq t} w_{t'} \right) \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t + 2 \sum_{t=1, t' < t} w_t w_{t'} \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_{t'} \\
&= \sum_{t=1}^T w_t \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t - \sum_{t=1}^T \left(w_t \sum_{t' \neq t} w_{t'} \right) \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t + 2 \sum_{t=1, t' < t} w_t w_{t'} \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_{t'} \\
&= \sum_{t=1}^T w_t \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t - \sum_{t=1}^T \sum_{t' \neq t} w_t w_{t'} \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t + 2 \sum_{t=1, t' < t} w_t w_{t'} \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_{t'} \\
&= \sum_{t=1}^T w_t \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t - \sum_{t=1, t' < t} w_t w_{t'} (\boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t + \boldsymbol{\tau}_{t'}^T \mathbf{H}_\ell \boldsymbol{\tau}_{t'}) + 2 \sum_{t=1, t' < t} w_t w_{t'} \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_{t'}
\end{aligned}$$

$$\begin{aligned}
&= \sum_{t=1}^T w_t \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t - \sum_{t=1, t' < t}^T w_t w_{t'} (\boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t + \boldsymbol{\tau}_{t'}^T \mathbf{H}_\ell \boldsymbol{\tau}_{t'}) - 2w_t w_{t'} \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_{t'} \\
&= \sum_{t=1}^T w_t \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t - \sum_{t=1, t' < t}^T w_t w_{t'} (\boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t + \boldsymbol{\tau}_{t'}^T \mathbf{H}_\ell \boldsymbol{\tau}_{t'} - 2\boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_{t'}) \\
&= \sum_{t=1}^T w_t \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t - \sum_{t=1, t' < t}^T w_t w_{t'} (\boldsymbol{\tau}_t - \boldsymbol{\tau}_{t'})^T \mathbf{H}_\ell (\boldsymbol{\tau}_t - \boldsymbol{\tau}_{t'})
\end{aligned}$$

Therefore:

$$\begin{aligned}
\ell_{\text{cur}}(\boldsymbol{\theta}_{\mathcal{P}}) &= \ell_0 + (\boldsymbol{\theta}_{\mathcal{P}} - \boldsymbol{\theta}_0)^T \mathbf{J}_\ell + \frac{1}{2} (\boldsymbol{\theta}_{\mathcal{P}} - \boldsymbol{\theta}_0)^T \mathbf{H}_\ell (\boldsymbol{\theta}_{\mathcal{P}} - \boldsymbol{\theta}_0) \\
&= \sum_{t=1}^T w_t [\ell_0 + \boldsymbol{\tau}_t^T \mathbf{J}_\ell] + \frac{1}{2} \boldsymbol{\tau}_{\mathcal{P}}^T \mathbf{H}_\ell \boldsymbol{\tau}_{\mathcal{P}}. \\
&= \sum_{t=1}^T w_t [\ell_0 + \boldsymbol{\tau}_t^T \mathbf{J}_\ell] + \frac{1}{2} \sum_{t=1}^T w_t \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t - \frac{1}{2} \sum_{t=1, t' < t}^T w_t w_{t'} (\boldsymbol{\tau}_t - \boldsymbol{\tau}_{t'})^T \mathbf{H}_\ell (\boldsymbol{\tau}_t - \boldsymbol{\tau}_{t'}) \\
&= \sum_{t=1}^T w_t \left[\ell_0 + \boldsymbol{\tau}_t^T \mathbf{J}_\ell + \frac{1}{2} w_t \boldsymbol{\tau}_t^T \mathbf{H}_\ell \boldsymbol{\tau}_t \right] - \frac{1}{2} \sum_{t=1, t' < t}^T w_t w_{t'} (\boldsymbol{\tau}_t - \boldsymbol{\tau}_{t'})^T \mathbf{H}_\ell (\boldsymbol{\tau}_t - \boldsymbol{\tau}_{t'}) \\
&= \sum_{t=1}^T w_t \ell_{\text{cur}}(\boldsymbol{\theta}_t) - \frac{1}{2} \sum_{t=1, t' < t}^T w_t w_{t'} (\boldsymbol{\tau}_t - \boldsymbol{\tau}_{t'})^T \mathbf{H}_\ell (\boldsymbol{\tau}_t - \boldsymbol{\tau}_{t'}),
\end{aligned}$$

which ends the proof of Theorem 1. \square

B.2 PROOF OF EQ. 17

Proof. In the dual-learner setting we have:

$$\Omega(A, B) = \frac{1}{2} w_A w_B (\boldsymbol{\tau}_A - \boldsymbol{\tau}_B)^T \mathbf{H}_\ell(\boldsymbol{\theta}_0) (\boldsymbol{\tau}_A - \boldsymbol{\tau}_B).$$

If we replace the Hessian matrix with the diagonal Fisher matrix $\hat{\mathbf{F}}_{\boldsymbol{\theta}_0}$, we get:

$$\begin{aligned}
\Omega(A, B) &= \frac{1}{2} w_A w_B (\boldsymbol{\tau}_A - \boldsymbol{\tau}_B)^T \hat{\mathbf{F}}_{\boldsymbol{\theta}_0} (\boldsymbol{\tau}_A - \boldsymbol{\tau}_B) = \\
&= \frac{1}{2} w_A w_B \left[\boldsymbol{\tau}_A^T \hat{\mathbf{F}}_{\boldsymbol{\theta}_0} \boldsymbol{\tau}_A - 2\boldsymbol{\tau}_A^T \hat{\mathbf{F}}_{\boldsymbol{\theta}_0} \boldsymbol{\tau}_B + \boldsymbol{\tau}_B^T \hat{\mathbf{F}}_{\boldsymbol{\theta}_0} \boldsymbol{\tau}_B \right].
\end{aligned}$$

If we focus on the first term inside the parenthesis, we obtain:

$$\begin{aligned}
\boldsymbol{\tau}_A^T \hat{\mathbf{F}}_{\boldsymbol{\theta}_0} \boldsymbol{\tau}_A &= \boldsymbol{\tau}_A^T \hat{\mathbf{F}}_{\boldsymbol{\theta}_0}^{1/2} \hat{\mathbf{F}}_{\boldsymbol{\theta}_0}^{1/2} \boldsymbol{\tau}_A = (\hat{\mathbf{F}}_{\boldsymbol{\theta}_0}^{1/2} \boldsymbol{\tau}_A)^T (\hat{\mathbf{F}}_{\boldsymbol{\theta}_0}^{1/2} \boldsymbol{\tau}_A) = \|\hat{\mathbf{F}}_{\boldsymbol{\theta}_0}^{1/2} \boldsymbol{\tau}_A\|_2^2 \\
&= \sum_{i=1}^{|\theta_A|} \hat{\mathbf{F}}_{\boldsymbol{\theta}_0}^{(i)} (\boldsymbol{\tau}_A^{(i)})^2 = \sum_{i=1}^{|\theta_A|} \hat{\mathbf{F}}_{\boldsymbol{\theta}_0}^{(i)} (\boldsymbol{\theta}_A^{(i)} - \boldsymbol{\theta}_0^{(i)})^2 = \text{EWC}_{\boldsymbol{\theta}_0}(\boldsymbol{\theta}_A).
\end{aligned}$$

Therefore, we can rewrite $\Omega(A, B)$ as:

$$\Omega(A, B) = \frac{1}{2} w_A w_B \left[\text{EWC}_{\boldsymbol{\theta}_0}(\boldsymbol{\theta}_A) + \text{EWC}_{\boldsymbol{\theta}_0}(\boldsymbol{\theta}_B) - 2\boldsymbol{\tau}_A^T \hat{\mathbf{F}}_{\boldsymbol{\theta}_0} \boldsymbol{\tau}_B \right].$$

We now generalize to $T \geq 2$. Starting from Eq. 14 and replacing the Hessian with the diagonal Fisher approximation $\hat{\mathbf{F}}_{\boldsymbol{\theta}_0}$, we obtain:

$$\frac{1}{2} \sum_{t=1, t' < t}^T w_t w_{t'} (\boldsymbol{\tau}_t - \boldsymbol{\tau}_{t'})^T \mathbf{H}_\ell (\boldsymbol{\tau}_t - \boldsymbol{\tau}_{t'}) \approx$$

$$\begin{aligned}
& \approx \frac{1}{2} \sum_{t=1, t' < t}^T w_t w_{t'} \left[\text{EWC}_{\theta_0}(\theta_t) + \text{EWC}_{\theta_0}(\theta_{t'}) - 2\tau_t^T \hat{\mathbf{F}}_{\theta_0} \tau_{t'} \right] \\
& = \frac{1}{2} \sum_{t=1, t' < t}^T w_t w_{t'} \left[\text{EWC}_{\theta_0}(\theta_t) + \text{EWC}_{\theta_0}(\theta_{t'}) \right] - \sum_{t=1, t' < t}^T w_t w_{t'} \tau_t^T \hat{\mathbf{F}}_{\theta_0} \tau_{t'} \\
& = \frac{1}{2} \sum_{t=1, t' \neq t}^T w_t w_{t'} \text{EWC}_{\theta_0}(\theta_t) - \sum_{t=1, t' < t}^T w_t w_{t'} \tau_t^T \hat{\mathbf{F}}_{\theta_0} \tau_{t'} \\
& = \frac{1}{2} \sum_{t=1}^T w_t \text{EWC}_{\theta_0}(\theta_t) \underbrace{\sum_{t' \neq t}^T w_{t'}}_{1-w_t} - \sum_{t=1, t' < t}^T w_t w_{t'} \tau_t^T \hat{\mathbf{F}}_{\theta_0} \tau_{t'} \\
& = \frac{1}{2} \sum_{t=1}^T w_t (1 - w_t) \text{EWC}_{\theta_0}(\theta_t) - \sum_{t=1, t' < t}^T w_t w_{t'} \tau_t^T \hat{\mathbf{F}}_{\theta_0} \tau_{t'},
\end{aligned}$$

which aligns with the result of Eq. 17. \square

C CLOSED FORM GRADIENTS FOR EQ. 19

We start by deriving the gradients of $\Omega_{\hat{\mathbf{F}}}(\cdot)$ (Eq. 14) w.r.t. the generic task vector τ_k . We initially consider the case of full-fine tuning, and then generalize the results to LoRA and (IA)³.

To simplify the calculations, we first rearrange $\Omega_{\hat{\mathbf{F}}}(\cdot)$ as:

$$\begin{aligned}
\Omega_{\hat{\mathbf{F}}}(\theta_1, \dots, \theta_T) & = \frac{1}{2} \sum_{t=1}^T \sum_{t' < t} w_t w_{t'} (\tau_t - \tau_{t'})^T \hat{\mathbf{F}}_{\theta_0} (\tau_t - \tau_{t'}) \\
& \approx \frac{1}{2} \sum_{t=1}^T w_t (1 - w_t) \text{EWC}_{\theta_0}(\theta_t) - \sum_{t=1, t' < t}^T w_t w_{t'} \tau_t^T \hat{\mathbf{F}}_{\theta_0} \tau_{t'} \quad (\text{see Eq. 17}),
\end{aligned}$$

Therefore, we can write:

$$\begin{aligned}
\frac{\partial \Omega_{\hat{\mathbf{F}}}}{\partial \tau_k} & = \frac{\partial \left[\frac{1}{2} \sum_{t=1}^T w_t (1 - w_t) \text{EWC}_{\theta_0}(\theta_t) \right]}{\partial \tau_k} - \frac{\partial \left[\sum_{t=1, t' < t}^T w_t w_{t'} \tau_t^T \hat{\mathbf{F}}_{\theta_0} \tau_{t'} \right]}{\partial \tau_k} \\
& = \frac{\partial \left[\frac{1}{2} w_k (1 - w_k) \text{EWC}_{\theta_0}(\theta_k) \right]}{\partial \tau_k} - \\
& \quad - \frac{\partial \left[\sum_{t' < k} w_k w_{t'} \tau_k^T \hat{\mathbf{F}}_{\theta_0} \tau_{t'} + \sum_{t \neq k, t' < t} w_t w_{t'} \tau_t^T \hat{\mathbf{F}}_{\theta_0} \tau_{t'} \right]}{\partial \tau_k} \\
& = w_k (1 - w_k) \frac{\partial \left[\frac{1}{2} \text{EWC}_{\theta_0}(\theta_k) \right]}{\partial \tau_k} - \sum_{t' \neq k}^T w_k w_{t'} \frac{\partial \tau_k^T \hat{\mathbf{F}}_{\theta_0} \tau_{t'}}{\partial \tau_k} \\
& = w_k (1 - w_k) \hat{\mathbf{F}}_{\theta_0} \odot \tau_k - w_k \sum_{t' \neq k}^T w_{t'} \hat{\mathbf{F}}_{\theta_0} \odot \tau_{t'} \\
& = w_k \left[(1 - w_k) \hat{\mathbf{F}}_{\theta_0} \odot \tau_k - \sum_{t' \neq k}^T w_{t'} \hat{\mathbf{F}}_{\theta_0} \odot \tau_{t'} \right] \\
& = w_k \left[\hat{\mathbf{F}}_{\theta_0} \odot \left[(1 - w_k) \tau_k - \sum_{t' \neq k}^T w_{t'} \tau_{t'} \right] \right].
\end{aligned} \tag{20}$$

We now suppose that training is performed incrementally on a sequence of $1, 2, \dots, k, \dots, T$ tasks. During the k -th task, IEL optimizes only τ_k , namely the task vector instantiated at the beginning of the k -th task. Indeed, as discussed in Sec. 2.3 (Eq. 19), the task vectors $\tau_1, \tau_2, \dots, \tau_{k-1}$ introduced in preceding tasks are kept frozen. Moreover, at each round, we devise uniform weights, such that $w_t = \frac{1}{k}$, with $t < k$. On top of that, we can rewrite the gradients in Eq. 20 as:

$$\frac{\partial \Omega_{\hat{F}}}{\partial \tau_k} = \frac{1}{k} \left[\hat{F}_{\theta_0} \odot \left[\left(1 - \frac{1}{k}\right) \tau_k - \frac{1}{k} \sum_{t < k} \tau_t \right] \right] \quad (22)$$

Notably, the term $\frac{1}{k} \sum_{t < k} \tau_t$ is a running average of the preceding task vectors. As also discussed in App. D, this reduces the memory/time computational cost for computing the gradients of $\Omega_{\hat{F}}$ to $\mathcal{O}(1)$.

As a final step, we can exploit the result in Eq. 22 to compute the gradients of LoRA ($\tau_k = B_k A_k$) and (IA)³. For LoRA, we have that:

$$\begin{aligned} \frac{\partial \Omega}{\partial B_k} &= \frac{\partial \Omega}{\partial \tau_k} \frac{\partial \tau_k}{\partial B_k} = \frac{1}{k} \left[\hat{F}_{\theta_0} \odot \left[\left(1 - \frac{1}{k}\right) \tau_k - \frac{1}{k} \sum_{t < k} \tau_t \right] \right] A^T \\ \frac{\partial \Omega}{\partial A_k} &= \frac{\partial \Omega}{\partial \tau_k} \frac{\partial \tau_k}{\partial A_k} = \frac{1}{k} B^T \left[\hat{F}_{\theta_0} \odot \left[\left(1 - \frac{1}{k}\right) \tau_k - \frac{1}{k} \sum_{t < k} \tau_t \right] \right] \end{aligned}$$

In the case of (IA)³, where $\tau_k = \theta_0 \odot ((l_k - \mathbf{1}_d) \otimes \mathbf{1}_d)$:

$$\frac{\partial \Omega}{\partial l_k} = \frac{\partial \Omega}{\partial \tau_k} \frac{\partial \tau_k}{\partial l_k} = \left[\frac{1}{k} \left[\hat{F}_{\theta_0} \odot \left[\left(1 - \frac{1}{k}\right) \tau_k - \frac{1}{k} \sum_{t < k} \tau_t \right] \right] \odot \theta_0^T \right] \mathbf{1}_d$$

where $\mathbf{1}_d$ is a d -dimensional column vector of ones with shape $d \times 1$.

D COMPUTATIONAL ANALYSIS

Analysis of ITA. Regarding ITA (*i.e.*, individual training), the learning phase has constant $\mathcal{O}(1)$ time and memory complexity. As each model is optimized in isolation, the training cost is similar to that of EWC (Kirkpatrick et al., 2017) and stems from the computation and storage of the Fisher Information Matrix, along with the calculations of the penalty term. During the evaluation phase, the time complexity of ITA remains $\mathcal{O}(1)$, as it involves a single forward pass on the composed model $f(\cdot; \theta_{\mathcal{P}})$. Memory complexity, on the other hand, is $\mathcal{O}(T)$ if maintaining separate models with distinct expertise is desired (*e.g.*, for later re-composition). Otherwise, this can be avoided by considering the composed model as a cumulative average of individual models (see later), resulting in a memory cost of $\mathcal{O}(1)$.

Analysis of IEL. Referring to IEL (*i.e.*, ensemble training), it might initially appear expensive due to the joint training of the ensemble. However, the complexity of IEL remains constant to $\mathcal{O}(1)$ in terms of both memory and time. This reduction is intuitive when we view the ensemble as a cumulative average of individual weights. To demonstrate this, we begin by considering the following cascade of models to be learned:

$$\begin{aligned} \text{Task}_{\#1} &\rightarrow f(\cdot; \theta_{\mathcal{P}} = \theta_0 + \tau_1) \\ \text{Task}_{\#2} &\rightarrow f(\cdot; \theta_{\mathcal{P}} = \theta_0 + \frac{1}{2} \tau_1 + \frac{1}{2} \tau_2) \\ &\dots \\ \text{Task}_{\#t} &\rightarrow f(\cdot; \theta_{\mathcal{P}} = \theta_0 + \frac{1}{t} \tau_1 + \frac{1}{t} \tau_2 + \dots + \frac{1}{t} \tau_{t-1} + \frac{1}{t} \tau_t) \end{aligned}$$

At each task, only the latter component of the composed model is learnable, while the preceding components are frozen:

$$\text{Task}_{\#t} \rightarrow f(\cdot; \theta_{\mathcal{P}} = \underbrace{\theta_0 + \frac{1}{t} \tau_1 + \frac{1}{t} \tau_2 + \dots + \frac{1}{t} \tau_{t-1}}_{\text{frozen components}} + \underbrace{\frac{1}{t} \tau_t}_{\text{learnable}}).$$

Therefore, as the preceding $\#t - 1$ components are kept frozen, we can incorporate them into the initialization weights $\theta_0 \rightarrow \theta_0^{(t)}$ and optimize:

$$\text{Task}_{\#t} \rightarrow f(\cdot; \theta_{\mathcal{P}} = \theta_0^{(t)} + \frac{1}{t} \tau_t) \quad \text{where} \quad \theta_0^{(t)} = \theta_0 + \frac{1}{t} \sum_{t'=1}^{t-1} \tau_{t'}.$$

Under this perspective, the learning of the t -th task is comparable to standard fine-tuning with a re-scaling factor $= 1/t$. Therefore, provided that $\theta_0^{(t)}$ is computed once at the beginning of the t -th task, the learning of the t -th task has constant $\mathcal{O}(1)$ time complexity.

It is noted that optimizing Eq. 19 via gradient descent can be similarly simplified, resulting in a constant $\mathcal{O}(1)$ time complexity. In fact, the gradients derived in App. C involve averaging the weights learned during previous tasks, specifically $\frac{1}{t} \sum_{t'=1}^{t-1} \tau_{t'}$. Since the weights from previous tasks are frozen during the current task, we can compute this average once and cache it.

To reduce memory complexity to $\mathcal{O}(1)$, we must avoid storing $\tau_1, \dots, \tau_{t-1}$ separately. This can be accomplished by assuming an initial null displacement $\tau_0 = \mathbf{0}$ and redefining $\theta_0^{(t)}$ as:

$$\begin{aligned} \theta_0^{(t)} &= \theta_0 + \frac{1}{t} \sum_{t'=1}^{t-1} \tau_{t'} = \theta_0 + \tau_0 + \frac{1}{t} \sum_{t'=1}^{t-1} \tau_{t'} = \theta_0 + \frac{1}{t} \sum_{t'=0}^{t-1} \tau_{t'} = \\ &= \theta_0 + \tau_{\text{AVG}}^{(t)} \end{aligned}$$

The term $\tau_{\text{AVG}}^{(t)} = \frac{1}{t} \sum_{t'=0}^{t-1} \tau_{t'}$ is basically the cumulative average of the displacements up to the current task (*excluded*). The cumulative average is straightforward to compute and, as is well-known, eliminates the need to store all previous values appearing in the sum, with resulting memory complexity $\mathcal{O}(1)$.

E IMPLEMENTATION DETAILS OF ITA AND IEL

Task pre-consolidation – Linear Probing. At the beginning of each task, during the pre-consolidation phase, we train a new classification head to account for the classes introduced by the current task. While the rest of the backbone remains frozen, a new linear classification layer is then trained with standard Stochastic Gradient Descent (SGD) for a varying number of epochs, depending on the dataset (typically either 3 or 8 epochs, as detailed in App. D).

Task pre-consolidation – Update of the FIM. Afterward, the diagonal Fisher Information Matrix (FIM) has to be updated to incorporate new information from the current task. Similar to Schwarz et al. (2018), we consider the estimated FIM as an online cumulative average of the squared gradients of the negative log-likelihood. Unlike Schwarz et al. (2018), we do not introduce the hyperparameter γ to down-weight the importance of the previous estimate. Moreover, following Kunstner et al. (2019), we compute the **true** FIM as defined in Eq. 10: *i.e.*, taking the expected gradient on the prediction vector \hat{y} . This approach differs from the majority of existing methods (Schwarz et al., 2018; Kirkpatrick et al., 2017; Chaudhry et al., 2018), which apply a further approximation by relying on the empirical FIM: *i.e.*, only considering the gradient of the ground truth label. Consequently, our estimate is more accurate but requires multiple backward passes, but the computational impact of this operation can be significantly reduced through batch-wise parallelization as in George (2021).

Fine-tuning – Initialization. All methods utilize the same pre-training (supervised) on ImageNet21K (`'vit_base_patch16_224.augreg_in21k'` from the `'timm'` library). Following the original works, we initialize the learnable parameters such that task vectors start from the pre-train initialization:

- Full Fine-Tuning: we apply zero-initialization.
- LoRA: we use Gaussian initialization for matrix A and zero initialization for matrix B .
- (IA)³: we initialize the vectors l with ones.

Fine-tuning – Loss function. Importantly, while our derivations regard the second-order approximation ℓ_{cur} , the full loss ℓ is instead employed in our algorithms. Indeed, as is common in

frameworks similar to ours (Chaudhry et al., 2018; Mirzadeh et al., 2020), the Taylor approximation is employed to build a surrogate of the exact loss that is both accurate and mathematically tractable. However, when it comes to practice, this proxy is often relaxed, and the full target function is used instead for simplicity.

Following existing works (Smith et al., 2023; Wang et al., 2022b), we employ the **local cross-entropy loss** as the learning objective. Given an example from the current task, while the standard cross-entropy loss considers logits related to all classes, including those learned in previous tasks, the local cross-entropy focuses only on logits corresponding to the classes introduced in the current task. This approach prevents the logits of past classes from being overly penalized during the current task (Caccia et al., 2022; Boschini et al., 2022). To ensure a fair comparison in our experiments, we apply this modification to other competing methods (*e.g.*, EWC (Kirkpatrick et al., 2017)).

Fine-tuning – Optimization. In each experiment, we use the AdamW optimizer (Loshchilov & Hutter, 2017) with a learning rate of 3×10^{-4} for LoRA and (IA)³ fine-tuning, and 1×10^{-4} for full fine-tuning. Importantly, in both ITA and IEL, we employ a decoupled strategy (Loshchilov & Hutter, 2017) to incorporate the gradients of the regularization term. Specifically, we apply the gradients of the regularizing objectives directly to the parameters before the gradient update step, ensuring that the regularization term does not interfere with momentum in the optimization dynamics. By adopting this approach, we observe an empirical improvement in final accuracy, attributed to a more effective minimization of the Riemannian distance relative to the pre-training initialization (see Eq. 12). Finally, we apply this decoupled gradient update exclusively to LoRA and (IA)³. For full fine-tuning, we refrain from using it as we observed numerical instabilities (*i.e.*, exploding loss).

Furthermore, we decouple the regularization strength applied to the final classification layer from that applied to the rest of the learnable parameters. This introduces two additional hyper-parameters: α_{CLS} for ITA and β_{CLS} for IEL. Intuitively, by decoupling the regularization weights, we can increase the regularization strength of intermediate layers without causing numerical instabilities, which often stem from the final classification layer.

F EXTENDED DISCUSSION ON RELATED WORKS AND COMPETING METHODS

To deliver the most comprehensive and significant comparison with the state of the art in incremental learning, we chose a combination of well-established standard approaches (such as EWC, DER++, L2P, and CODA) and recent proposals emphasizing compositionality skills (*e.g.*, TMC and APT) and ensemble learning (*i.e.*, SEED). It is important to note that these methods have been heavily influenced by the research trends prevalent at the time they were originally proposed, and therefore, they tend to employ different strategies for fine-tuning the model. To sum up:

- EWC, LwF, DER++, SEED and TMC leverage on full-fine tuning.
- L2P, CODA, APT resort instead to prompting.

Therefore, achieving a direct apple-to-apple comparison is challenging, as the approaches present in the literature are themselves prone to this issue.

We herein summarize the main aspects of the more important recent methods we compared with.

Elastic Weight Consolidation (EWC) As discussed in the main paper, while our approach regularizes the distance in parameter space with respect to θ_0 , the regularization in EWC (Kirkpatrick et al., 2017) instead focuses on the weights learned during the preceding task. However, beyond the different regularizing strategies, there is another significant difference between ITA and EWC. While ITA fine-tunes each task starting from the same original pre-trained model θ_0 , EWC begins with the weights of the previous task.

It is noted that an EWC-like term that protects the last task weights could work as well in our framework based on task vectors. We chose to anchor the model to the pre-training weights to allow for more flexible decentralized learning, wherein multiple task vectors can be trained on different tasks *in parallel*, with minimal interactions. In contrast, an EWC-like term, which regularizes each successor based on its predecessor, would necessitate training the multiple task vectors *in sequence*.

Learning to Prompt (L2P) (Wang et al., 2022b) and **CODA-Prompt** (Smith et al., 2023) are two continual learning techniques based on prompting. They fine-tune a pre-trained model through a few learnable parameters stored in a prompt pool, which can be either shared or tied to different tasks. At inference time, the prompts are retrieved from the pool through a query-key search. In terms of memory complexity, our ITA is in line with L2P and CODA when coupled with a parameter-efficient fine-tuning technique such as IA3. Moreover, ITA does not require the additional forward pass on the frozen pre-trained model required to retrieve the prompts.

À-la-carte Prompt Tuning (APT) (Bowman et al., 2023) is a prompt-based strategy similar to L2P and CODA. The prompts are trained in isolation (similarly to ITA) and concatenated at inference time to create the composed predictor (no prior query-key search is required). To avoid destructive interference, a tailored masking mechanism is employed in self-attention layers. Differently, our approaches fuse parameters through linear combinations.

Tangent Model Composition (TMC) (Liu & Soatto, 2023) is a recent approach addressing continual learning through task arithmetic in the tangent space. TMC builds upon task vectors and is similar to ITA (full fine-tuning). They differ in two aspects: *i*) TMC applies a first-order approximation of the forward pass to support compositionality, making it two to three times slower than a non-linear forward pass; *ii*) TMC does not include auxiliary regularization during training.

Selection of Experts for Ensemble Diversification (SEED) (Rypeś et al., 2024) is a recent approach that trains an ensemble of models incrementally. For each incoming task, an expert model is chosen from the pool and trained. The major difference with respect to our IEL concerns the inference stage: while IEL performs an ensemble prediction in $\mathcal{O}(1)$ time (thanks to weight averaging), SEED makes inference on all models at test time and averages their predictions.

Model merging. While we address compositionality during training, other approaches focus on post-training techniques, as simple averaging leads to interference (Yadav et al., 2024b) when parameters are redundant or have conflicting signs. TIES (Yadav et al., 2024b) discards uninformative parameters and addresses conflicts via majority voting. Zipit! (Stoica et al., 2023) merges redundant parameters that produce similar features, while RegMean (Jin et al., 2022) exploits a closed-form solution for linear layers. Notably, (Matena & Raffel, 2022) weighs the contribution of each parameter through the Fisher matrix, computed by each individual model at its optimum. This differs from our approach that evaluates the FIM at the pre-training optimum.

G DATASETS

We conduct a comprehensive evaluation using a variety of benchmarks. Following the current literature on pre-trained CL models (Wang et al., 2022b;a; Smith et al., 2023), we include conventional image datasets such as Split CIFAR-100 and Split ImageNet-R. We also include Split CUB-200, Split Caltech-256 and Split MIT-67, recently used in the context of composable incremental methods (Bowman et al., 2023; Liu & Soatto, 2023). Finally, we assess the adaptability of these pre-trained methods in settings with decreasing domain similarity to the ImageNet pre-training utilized by our backbone model (Oh et al., 2022; Cui et al., 2018). Specifically, these settings include the satellite and medical domains, represented by Split RESISC45 and Split CropDiseases, respectively. In the following, we outline the due details:

- **Standard domains:** *Split CIFAR-100* (Krizhevsky et al., 2009) and *Split ImageNet-R* (Hendrycks et al., 2021), with respectively 100 and 200 classes split into 10 tasks. We train each task of Split ImageNet-R for 30 epochs and each task of Split CIFAR-100 for 20 epochs. In particular, IN-R is a variant of the ImageNet dataset that includes artistic renditions such as sketches, cartoons, and paintings. It is used to evaluate the robustness and generalization capabilities of models trained on the original ImageNet when tested on out-of-distribution data. Following (Liu & Soatto, 2023), we also employ Split Caltech-256 (Griffin et al., 2007) and Split MIT-67 (Quattoni & Torralba, 2009), dividing both into 10 tasks (5 epoch each).
- **Specialized domain:** We adopt *Split CUB-200* (Wah et al., 2011) to evaluate compositional capabilities in a more fine-grained classification scenario, namely recognizing 200 species of birds. The classes are split across 10 tasks, each lasting for 50 epochs.

Table 4: Comparison with the SOTA across several benchmarks (Final Forgetting [↓]).

Model	IN-R	C-100	CUB	Caltech	MIT	RESISC	CropDis.
Joint	\times	\times	\times	\times	\times	\times	\times
Finetune	75.5	85.7	65.79	58.76	87.89	97.6	92.04
EWC (ON)	32.43	25.37	53.04	23.52	34.65	43.12	31.58
LWF-MC	10.39	6.13	26.2	5.23	3.56	5.66	7.94
DER++	35.81	15.66	13.58	11.19	31.51	43.04	0.76
L2P	5.6	6.08	6.63	1.78	5.55	26.29	17.81
CODA	4.97	5.28	8.82	2.64	7.49	22	15.37
SEED	8.96	6.39	5.12	2.67	4.53	8.66	1.19
InfLoRA	5.67	4.05	4.58	1.89	6.85	10.77	4.79
APT	8.99	6.71	8.82	3.64	6.33	27.74	14.01
TMC	11.66	9.2	7.37	6.39	12.43	16.43	16.94
ITA-FFT	7.72	3.89	1.39	2.44	4.80	6.19	0.95
ITA-LoRA	8.00	3.22	2.08	2.38	4.68	8.30	0.68
ITA-(IA) ³	8.72	3.57	2.23	2.39	4.48	7.33	0.76
IEL-FFT	5.40	2.20	2.04	1.50	2.29	3.46	1.40
IEL-LoRA	7.63	5.14	3.46	2.46	8.34	8.30	2.14
IEL-(IA) ³	8.34	3.27	1.69	2.20	5.82	6.86	1.04

- **Aerial domain:** we use *Split RESISC45* (Cheng et al., 2017), which comprises 30000 RGB satellite images for land use and land cover classification. The dataset contains 45 classes (*e.g.*, airport, cloud, island, and so on) divided into 9 tasks, with each task lasting 30 epochs.
- **Medical domain:** we finally explore the medical setting (*i.e.*, plant diseases) and conduct experiments on *Split CropDiseases* (Hughes et al., 2015). It regards infected leaves with 7 tasks of 5 classes each (5 epochs).

We base our code on Mammoth (Buzzega et al., 2020b;a), a widely adopted framework in the class-incremental learning literature.

H ADDITIONAL RESULTS

H.1 FORGETTING

Tab. 4 reports the Final Forgetting metric (Chaudhry et al., 2018) for our experiments.

H.2 STANDARD DEVIATION OF FA

Tab. 5 reports the standard deviation of the Final Accuracy metrics reported in Tab. 1.

H.3 ABLATION STUDY ON IEL

We herein present an ablative analysis regarding the ensemble-oriented regularization applied to the proposed IEL approach. Specifically, we evaluate the impact of Eq. 19 on the final accuracy and report the results in Tab. 6. After examining them, we can draw conclusions similar to those made for ITA. In particular, the regularization driven by the second-order formulation proves beneficial for achieving effective composition (especially the full-fine tuning), with the final classification layer playing an important role in attaining good performance.

H.4 ADDITIONAL EXPERIMENTS ON INCREMENTAL MODEL COMPOSITIONALITY

Tab. 7 reports the results of ITA, IEL and TMC on additional datasets, with a focus on their compositionality capabilities (*i.e.*, specialization and unlearning).

Table 5: For each experiment of Tab. 1, the standard deviation of the Final Accuracy (FA).

Model	IN-R	C-100	CUB	Caltech	MIT	RESISC	CropDis.
Joint	0.87	0.16	0.41	0.43	0.31	0.22	0.21
Finetune	4.53	2.83	1.65	2.57	3.49	1.12	0.43
EWC	0.73	1.20	0.87	1.29	1.45	0.61	0.99
LWF-MC	0.89	1.12	3.23	3.94	2.98	2.85	2.85
DER++	1.12	1.14	0.82	1.27	0.95	1.79	1.29
L2P	0.56	1.23	2.12	0.63	0.81	3.42	0.45
CODA	0.63	0.95	2.87	0.82	0.69	4.61	0.98
SEED	0.18	0.47	0.65	0.55	0.41	0.86	0.23
APT	0.84	1.30	2.45	0.95	0.73	1.94	0.68
InfLoRA	0.10	0.21	0.89	0.68	2.03	1.16	0.42
TMC	0.45	0.85	1.67	0.71	0.92	0.82	0.81
ITA-FFT	0.27	0.27	0.39	0.18	0.24	0.56	2.10
ITA-LoRA	0.26	0.20	0.13	0.16	0.40	0.89	0.55
ITA-(IA) ³	0.28	0.04	0.18	0.18	0.69	0.68	0.27
IEL-FFT	0.11	0.17	0.21	0.09	1.12	0.77	0.90
IEL-LoRA	0.36	0.27	0.37	0.41	0.93	0.30	0.43
IEL-(IA) ³	0.23	0.20	0.32	0.21	0.55	0.48	0.41

Table 6: Ablation study for ITA-(IA)³ and IEL on several benchmarks (FA [↑]).

Model	IN-R	C-100	CUB	Caltech	MIT	RESISC	CropDis.
ITA-(IA)³ (reg)	77.04	90.66	85.67	92.67	84.74	83.73	95.41
<i>without Eq. 12 reg.</i>	71.82	88.43	77.61	90.66	69.14	69.01	63.72
<i>Eq. 12 only on CLS</i>	76.72	90.48	85.56	92.56	85.25	84.37	95.45
IEL-FFT (reg)	80.09	89.38	84.89	92.23	82.79	81.42	95.83
<i>without Eq. 19 reg.</i>	40.85	52.56	14.02	53.76	47.63	39.20	31.24
<i>Eq. 19 reg. only on CLS</i>	77.99	85.82	85.30	91.43	77.58	76.87	96.18
IEL-LoRA (reg)	79.93	89.53	84.95	92.19	84.49	82.53	95.88
<i>without Eq. 19 reg.</i>	51.15	66.01	60.39	70.71	55.38	42.72	45.25
<i>Eq. 19 reg. only on CLS</i>	76.14	86.11	84.43	91.77	82.50	70.05	95.54
IEL-(IA)³ (reg)	77.86	89.72	84.57	92.70	85.54	81.50	95.68
<i>without Eq. 19 reg.</i>	73.72	84.00	74.72	89.58	69.82	62.52	66.29
<i>Eq. 19 reg. only on CLS</i>	77.23	89.38	84.70	92.76	85.43	81.60	95.72

H.5 ALIGNMENT BETWEEN TASK VECTORS OF ITA AND IEL

To better understand the relationship between the task vectors learned by ITA and IEL, in Fig. 2 we conduct an additional evaluation by measuring the alignment – in terms of cosine similarity – between the parameters learned by the two proposed approaches. The analysis is twofold: first, the alignment is evaluated between individual task vectors τ_t^{ITA} and τ_t^{IEL} by averaging the similarity across tasks; second, it is assessed between composed models $\tau_{\mathcal{P}}^{\text{ITA}}$ and $\tau_{\mathcal{P}}^{\text{IEL}}$. To eliminate potential confounding effects due to PEFT strategies, we perform the experiment using full fine-tuning.

Our results reveal mostly positive alignment, even for out-of-distribution datasets such as RESISC45 and CropDisease, and low-resolution datasets like CIFAR-100. On the other hand, we observe smaller similarities between the task vectors of the individual models learned on more fine-grained and complex datasets such as CUB-200 and MIT67. Indeed, CUB-200 demands highly detailed distinctions between bird species, while MIT67 involves subtle contextual differences, such as distinguishing children’s rooms from living rooms or art studios from classrooms.

Table 7: Analysis of compositional capabilities. In parentheses, we report the gain (or loss) in accuracy on the target task.

Dataset	Model	<i>zero-shot specialization</i>		<i>zero-shot unlearning</i>	
		FA _{TGT} [↑]	FA _{CTRL}	FA _{TGT} [↓]	FA _{CTRL}
IN-R	ITA-LoRA	80.83 (+11.40)	50.52	22.77 (-55.02)	52.72 (-25.07)
	IEL-LoRA	73.46 (-06.68)	38.46	18.55 (-61.38)	41.97 (-37.96)
	TMC	69.93 (+08.36)	34.08	45.77 (-14.24)	54.37 (-05.64)
C-100	ITA-LoRA	92.80 (+01.63)	60.06	28.67 (-61.29)	71.96 (-17.99)
	IEL-LoRA	77.77 (-13.22)	37.90	19.48 (-70.05)	56.52 (-33.01)
	TMC	87.53 (+06.49)	45.75	55.63 (-22.79)	71.83 (-06.59)
CUB	ITA-LoRA	90.46 (+05.21)	57.87	68.19 (-17.36)	74.63 (-10.92)
	IEL-LoRA	74.03 (-10.57)	47.95	22.44 (-62.51)	30.99 (-53.96)
	TMC	71.06 (+09.92)	44.93	67.91 (-03.81)	53.22 (-18.50)
Caltech	ITA-LoRA	89.84 (-01.21)	65.47	80.02 (-12.63)	82.40 (-10.25)
	IEL-LoRA	79.70 (-12.99)	62.45	32.00 (-60.19)	42.03 (-50.16)
	TMC	89.23 (+10.63)	49.46	64.52 (-17.78)	73.33 (-08.97)
MIT	ITA-LoRA	89.66 (+08.17)	57.30	36.99 (-49.61)	74.75 (-11.85)
	IEL-LoRA	56.14 (-19.38)	37.54	06.39 (-78.10)	32.57 (-52.02)
	TMC	88.03 (+10.11)	30.56	29.77 (-38.89)	62.03 (-06.63)
RESISC	ITA-LoRA	89.47 (+06.70)	49.75	33.38 (-48.62)	64.06 (-17.94)
	IEL-LoRA	90.13 (+04.92)	48.06	32.77 (-49.76)	65.19 (-17.34)
	TMC	75.77 (+27.63)	17.00	06.64 (-54.02)	52.31 (-08.35)
CropDis.	ITA-LoRA	97.63 (+00.11)	53.05	65.87 (-29.98)	79.24 (-16.61)
	IEL-LoRA	90.12 (-04.68)	48.31	34.01 (-61.87)	54.43 (-41.45)
	TMC	73.60 (+09.23)	15.95	06.24 (-60.32)	53.21 (-13.35)

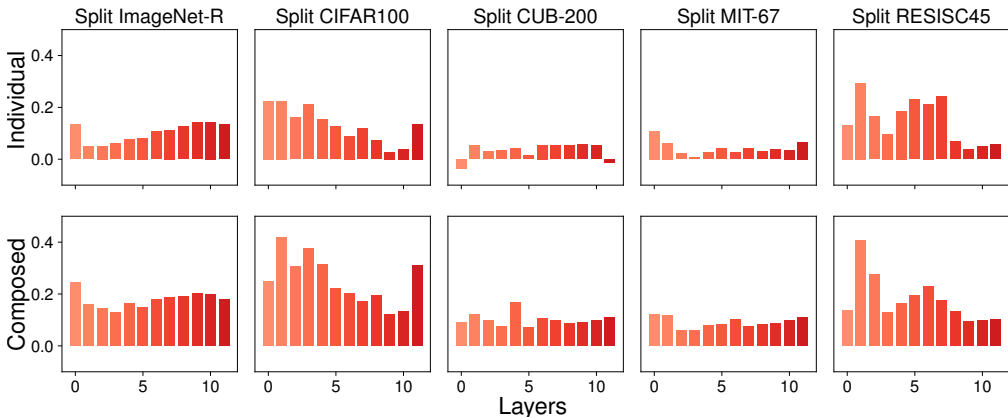


Figure 2: Alignment – *i.e.*, cosine similarity – between the task vectors produced by ITA and IEL for both the *composed* model $\theta_{\mathcal{P}}$ and *individual* learners θ_t (averaged across tasks t).

Despite the lower alignment for individual task vectors in certain datasets, we observe that the similarity is consistently higher when comparing the composed task vectors to the individual ones.

H.6 TIMING

We herein include a wall-clock time analysis of the training algorithms. Specifically, considering Split ImageNet-R, Split RESISC45, and Split CropDisease, we measured the per-task runtime of ITA and IEL (both trained with full fine-tuning), as well as that of three existing approaches to

1404
 1405
 1406
 1407
 1408
 1409
 1410
 1411
 1412
 1413
 1414
 1415
 1416
 1417
 1418
 1419
 1420
 1421
 1422
 1423
 1424
 1425
 1426
 1427
 1428
 1429
 1430
 1431
 1432
 1433
 1434
 1435
 1436
 1437
 1438
 1439
 1440
 1441
 1442
 1443
 1444
 1445
 1446
 1447
 1448
 1449
 1450
 1451
 1452
 1453
 1454
 1455
 1456
 1457




Figure 3: Comparative timing analysis (in minutes). The plot illustrates the per-task runtime of ITA and IEL, alongside baseline methods (DER++, TMC, and SEED). Runtimes include both the setup phase (e.g., steps required to compute FIM statistics) and the training phase.

incremental learning: DER++ (*rehearsal*), TMC (*model compositionality*), and SEED (*ensemble learning*).

Among the methods compared, SEED appears to be the most efficient. However, the runtimes of both ITA and IEL are comparable to, or even better than, those of baseline methods like DER++ and TMC. Specifically, on ImageNet-R and RESISC45, the average runtime of ITA/IEL is approximately $2/3$ that of DER++. We attribute the reasonable training times of IEL/ITA to the efficient procedure we employed for estimating the FIM, which leverages batch-wise parallelization (see App. E and George (2021)), combined with the use of closed-form gradient computation (refer to App. C).

I HYPERPARAMETERS

The hyperparameters employed for each experiment are reported in the following subsections (one for each dataset).

I.1 IMAGENET-R

CODA-Prompt: $\text{lr} = 1.0 \times 10^{-3}$;

DER++: $\alpha = 3.0 \times 10^{-1}$; $\beta = 8.0 \times 10^{-1}$; $\text{lr} = 1.0 \times 10^{-3}$;

EWC: $\epsilon = 1.0 \times 10^2$; $\gamma = 1.0$; $\text{lr} = 1.0 \times 10^{-2}$;

L2P: $\text{lr} = 2.5 \times 10^{-3}$;

LWF-MC: $\text{wd} = 0.0$; $\text{lr} = 1.0 \times 10^{-2}$;

SEED: $\text{lr} = 3.0 \times 10^{-4}$;

SGD: $\text{lr} = 3.0 \times 10^{-2}$;

TMC: $\text{lr} = 1.0 \times 10^{-4}$;

InfLoRA: $\text{lr} = 5.0 \times 10^{-4}$ $r = 10$; $\epsilon = 0.98$;

For both ITA and IEL: $\# \text{ epochs}_{\text{pre-tuning}} = 3$; $\text{lr}_{\text{pre-tuning}} = 1.0 \times 10^{-2}$;

ITA-FFT: $\alpha = 5.0 \times 10^1$; $\alpha_{\text{CLS}} = 1.0 \times 10^{-1}$; $\alpha_{\text{CLS-prior}} = 1.0 \times 10^{-2}$; $\text{lr} = 1.0 \times 10^{-4}$;

ITA-LoRA: $\alpha = 2.0 \times 10^{-2}$; $\alpha_{\text{CLS}} = 1.0 \times 10^{-1}$; $\alpha_{\text{CLS-prior}} = 1.0 \times 10^{-3}$; $\text{lr} = 3.0 \times 10^{-4}$;

ITA-(IA)³: $\alpha = 7.0 \times 10^{-1}$; $\alpha_{\text{CLS}} = 1.0 \times 10^{-1}$; $\alpha_{\text{CLS-prior}} = 1.0 \times 10^{-2}$; $\text{lr} = 3.0 \times 10^{-3}$;

IEL-FFT: $\beta = 5.0 \times 10^1$; $\beta_{\text{CLS}} = 1.0 \times 10^{-1}$; $\beta_{\text{CLS-prior}} = 3.0 \times 10^{-3}$; $\text{lr} = 1.0 \times 10^{-4}$;

IEL-LoRA: $\beta = 2.0 \times 10^1$; $\beta_{\text{CLS}} = 1.0 \times 10^{-1}$; $\beta_{\text{CLS-prior}} = 1.0 \times 10^{-3}$; $\text{lr} = 3.0 \times 10^{-4}$;

IEL-(IA)³: $\beta = 2.0 \times 10^1$; $\beta_{\text{CLS}} = 1.0 \times 10^{-1}$; $\beta_{\text{CLS-prior}} = 1.0 \times 10^{-2}$; $\text{lr} = 3.0 \times 10^{-3}$;

1458 I.2 CIFAR-100

1459 **CODA-Prompt**: $\text{lr} = 1.0 \times 10^{-3}$;
 1461 **DER++**: $\alpha = 3.0 \times 10^{-1}$; $\beta = 8.0 \times 10^{-1}$; $\text{lr} = 1.0 \times 10^{-4}$;
 1462 **EWC**: $\epsilon = 1.0 \times 10^2$; $\gamma = 1.0$; $\text{lr} = 1.0 \times 10^{-3}$;
 1464 **L2P**: $\text{lr} = 2.5 \times 10^{-3}$;
 1465 **LWF-MC**: $\text{wd} = 0.0$; $\text{lr} = 1.0 \times 10^{-2}$;
 1467 **SEED**: $\text{lr} = 3.0 \times 10^{-4}$;
 1468 **SGD**: $\text{lr} = 1.0 \times 10^{-2}$;
 1469 **TMC**: $\text{lr} = 1.0 \times 10^{-4}$;
 1471 **InfLoRA**: $\text{lr} = 5.0 \times 10^{-4}$ $r = 10$; $\epsilon = 0.95$;

1473 For both ITA and IEL: $\# \text{ epochs}_{\text{pre-tuning}} = 3$; $\text{lr}_{\text{pre-tuning}} = 1.0 \times 10^{-2}$;
 1474 **ITA-FFT**: $\alpha = 2.0 \times 10^3$; $\alpha_{\text{CLS}} = 1.0 \times 10^{-1}$; $\alpha_{\text{CLS-prior}} = 3.0 \times 10^{-3}$; $\text{lr} = 1.0 \times 10^{-4}$;
 1475 **ITA-LoRA**: $\alpha = 2.0 \times 10^{-2}$; $\alpha_{\text{CLS}} = 1.0 \times 10^{-1}$; $\alpha_{\text{CLS-prior}} = 1.0 \times 10^{-3}$; $\text{lr} = 3.0 \times 10^{-4}$;
 1477 **ITA-(IA)³**: $\alpha = 2.0 \times 10^{-2}$; $\alpha_{\text{CLS}} = 1.0 \times 10^{-1}$; $\alpha_{\text{CLS-prior}} = 3.0 \times 10^{-3}$; $\text{lr} = 3.0 \times 10^{-3}$;
 1478 **IEL-FFT**: $\beta = 2.0 \times 10^3$; $\beta_{\text{CLS}} = 2.5 \times 10^{-2}$; $\beta_{\text{CLS-prior}} = 1.0 \times 10^{-2}$; $\text{lr} = 1.0 \times 10^{-4}$;
 1480 **IEL-LoRA**: $\beta = 2.0 \times 10^1$; $\beta_{\text{CLS}} = 1.0 \times 10^{-1}$; $\beta_{\text{CLS-prior}} = 1.0 \times 10^{-3}$; $\text{lr} = 3.0 \times 10^{-4}$;
 1481 **IEL-(IA)³**: $\beta = 2.0 \times 10^1$; $\beta_{\text{CLS}} = 2.5 \times 10^{-2}$; $\beta_{\text{CLS-prior}} = 1.0 \times 10^{-3}$; $\text{lr} = 3.0 \times 10^{-4}$;

1484 I.3 CUB-200

1485 **CODA-Prompt**: $\text{lr} = 1.0 \times 10^{-3}$;
 1487 **DER++**: $\alpha = 3.0 \times 10^{-1}$; $\beta = 8.0 \times 10^{-1}$; $\text{lr} = 1.0 \times 10^{-3}$;
 1488 **EWC**: $\epsilon = 1.0 \times 10^1$; $\gamma = 9.0 \times 10^{-1}$; $\text{lr} = 1.0 \times 10^{-2}$;
 1489 **L2P**: $\text{lr} = 2.5 \times 10^{-3}$;
 1491 **LWF-MC**: $\text{wd} = 1.0 \times 10^{-4}$; $\text{lr} = 1.0 \times 10^{-2}$;
 1492 **SEED**: $\text{lr} = 3.0 \times 10^{-4}$;
 1494 **SGD**: $\text{lr} = 3.0 \times 10^{-2}$;
 1495 **TMC**: $\text{lr} = 1.0 \times 10^{-4}$;
 1497 **InfLoRA**: $\text{lr} = 5.0 \times 10^{-4}$ $r = 10$; $\epsilon = 0.98$;

1498 For both ITA and IEL: $\text{lr}_{\text{pre-tuning}} = 1.0 \times 10^{-2}$; $\# \text{ epochs}_{\text{pre-tuning}} = 8$.
 1499 **ITA-FFT**: $\alpha = 5.0 \times 10^2$; $\alpha_{\text{CLS}} = 1.0 \times 10^{-1}$; $\alpha_{\text{CLS-prior}} = 1.0 \times 10^{-2}$; $\text{lr} = 1.0 \times 10^{-4}$;
 1501 **ITA-LoRA**: $\alpha = 5.0$; $\alpha_{\text{CLS}} = 1.0 \times 10^{-1}$; $\alpha_{\text{CLS-prior}} = 1.0 \times 10^{-2}$; $\text{lr} = 3.0 \times 10^{-4}$;
 1502 **ITA-(IA)³**: $\alpha = 7.0 \times 10^{-1}$; $\alpha_{\text{CLS}} = 1.0 \times 10^{-1}$; $\alpha_{\text{CLS-prior}} = 1.0 \times 10^{-2}$; $\text{lr} = 3.0 \times 10^{-3}$;
 1503 **IEL-FFT**: $\beta = 5.0 \times 10^3$; $\beta_{\text{CLS}} = 1.0 \times 10^{-1}$; $\beta_{\text{CLS-prior}} = 1.0 \times 10^{-2}$; $\text{lr} = 1.0 \times 10^{-5}$;
 1504 **IEL-LoRA**: $\beta = 2.0 \times 10^1$; $\beta_{\text{CLS}} = 1.0 \times 10^{-1}$; $\beta_{\text{CLS-prior}} = 1.0 \times 10^{-2}$; $\text{lr} = 3.0 \times 10^{-4}$;
 1505 **IEL-(IA)³**: $\beta = 5.0 \times 10^2$; $\beta_{\text{CLS}} = 1.0 \times 10^{-1}$; $\beta_{\text{CLS-prior}} = 1.0 \times 10^{-2}$; $\text{lr} = 3.0 \times 10^{-4}$;

1510 I.4 CALTECH-256

1511 **CODA-Prompt**: $\text{lr} = 1.0 \times 10^{-3}$;

1512 **DER++:** $\alpha = 3.0 \times 10^{-1}$; $\beta = 8.0 \times 10^{-1}$; $\text{lr} = 1.0 \times 10^{-3}$;

1513 **EWC:** $\epsilon = 1.0 \times 10^2$; $\gamma = 1.0$; $\text{lr} = 1.0 \times 10^{-2}$;

1514 **L2P:** $\text{lr} = 2.5 \times 10^{-3}$;

1515 **LWF-MC:** $\text{wd} = 0.0$; $\text{lr} = 1.0 \times 10^{-2}$;

1516 **SEED:** $\text{lr} = 3.0 \times 10^{-4}$;

1517 **SGD:** $\text{lr} = 1.0 \times 10^{-2}$;

1518 **TMC:** $\text{lr} = 1.0 \times 10^{-4}$;

1519 **InfLoRA:** $\text{lr} = 5.0 \times 10^{-4}$ $r = 10$; $\epsilon = 0.99$;

1520

1521 For both ITA and IEL: $\text{lr}_{\text{pre-tuning}} = 1.0 \times 10^{-2}$.

1522 **ITA-FFT:** $\# \text{ epochs}_{\text{pre-tuning}} = 3$; $\alpha = 2.0 \times 10^3$; $\alpha_{\text{CLS}} = 1.0 \times 10^{-1}$; $\alpha_{\text{CLS-prior}} = 1.0 \times 10^{-2}$;

1523 **ITA-LoRA:** $\# \text{ epochs}_{\text{pre-tuning}} = 3$; $\alpha = 2.0 \times 10^{-2}$; $\alpha_{\text{CLS}} = 1.0 \times 10^{-1}$; $\alpha_{\text{CLS-prior}} = 3.0 \times 10^{-3}$;

1524 **ITA-(IA)³:** $\# \text{ epochs}_{\text{pre-tuning}} = 8$; $\alpha = 7.0 \times 10^{-1}$; $\alpha_{\text{CLS}} = 1.0 \times 10^{-1}$; $\alpha_{\text{CLS-prior}} = 1.0 \times 10^{-2}$;

1525 **IEL-FFT:** $\# \text{ epochs}_{\text{pre-tuning}} = 3$; $\beta = 5.0 \times 10^1$; $\beta_{\text{CLS}} = 1.0 \times 10^{-1}$; $\beta_{\text{CLS-prior}} = 1.0 \times 10^{-2}$;

1526 **IEL-LoRA:** $\# \text{ epochs}_{\text{pre-tuning}} = 3$; $\beta = 2.0 \times 10^1$; $\beta_{\text{CLS}} = 5.0$; $\beta_{\text{CLS-prior}} = 3.0 \times 10^{-3}$;

1527 **IEL-(IA)³:** $\# \text{ epochs}_{\text{pre-tuning}} = 3$; $\beta = 2.0 \times 10^1$; $\beta_{\text{CLS}} = 1.0 \times 10^{-1}$; $\beta_{\text{CLS-prior}} = 3.0 \times 10^{-3}$;

1528

1529 I.5 MIT-67

1530

1531 **CODA-Prompt:** $\text{lr} = 1.0 \times 10^{-3}$;

1532 **DER++:** $\alpha = 3.0 \times 10^{-1}$; $\beta = 8.0 \times 10^{-1}$; $\text{lr} = 1.0 \times 10^{-3}$;

1533 **EWC:** $\epsilon = 1.0 \times 10^2$; $\gamma = 1.0$; $\text{lr} = 1.0 \times 10^{-3}$;

1534 **L2P:** $\text{lr} = 2.5 \times 10^{-3}$;

1535 **LWF-MC:** $\text{wd} = 0.0$; $\text{lr} = 1.0 \times 10^{-2}$;

1536 **SEED:** $\text{lr} = 3.0 \times 10^{-4}$;

1537 **SGD:** $\text{lr} = 1.0 \times 10^{-2}$;

1538 **TMC:** $\text{lr} = 1.0 \times 10^{-4}$;

1539 **InfLoRA:** $\text{lr} = 1.0 \times 10^{-3}$ $r = 10$; $\epsilon = 0.95$;

1540

1541 For both ITA and IEL: $\text{lr}_{\text{pre-tuning}} = 1.0 \times 10^{-2}$.

1542 **ITA-FFT:** $\# \text{ epochs}_{\text{pre-tuning}} = 8$; $\alpha = 5.0 \times 10^3$; $\alpha_{\text{CLS}} = 1.0 \times 10^{-1}$; $\alpha_{\text{CLS-prior}} = 1.0 \times 10^{-1}$;

1543 **ITA-LoRA:** $\# \text{ epochs}_{\text{pre-tuning}} = 8$; $\alpha = 5.0$; $\alpha_{\text{CLS}} = 1.0 \times 10^{-1}$; $\alpha_{\text{CLS-prior}} = 1.0 \times 10^{-2}$;

1544 **ITA-(IA)³:** $\# \text{ epochs}_{\text{pre-tuning}} = 8$; $\alpha = 5.0$; $\alpha_{\text{CLS}} = 1.0 \times 10^{-1}$; $\alpha_{\text{CLS-prior}} = 3.0 \times 10^{-3}$;

1566 **IEL-FFT**: $\# \text{epochs}_{\text{pre-tuning}} = 8; \beta = 5.0 \times 10^3; \beta_{\text{CLS}} = 1.0 \times 10^{-1}; \beta_{\text{CLS-prior}} = 1.0 \times 10^{-2};$
 1567 $\text{lr} = 1.0 \times 10^{-4};$

1568 **IEL-LoRA**: $\# \text{epochs}_{\text{pre-tuning}} = 3; \beta = 2.0 \times 10^1; \beta_{\text{CLS}} = 1.0 \times 10^{-1}; \beta_{\text{CLS-prior}} = 1.0 \times 10^{-2};$
 1569 $\text{lr} = 3.0 \times 10^{-4};$

1571 **IEL-(IA)³**: $\# \text{epochs}_{\text{pre-tuning}} = 3; \beta = 2.0 \times 10^1; \beta_{\text{CLS}} = 1.0 \times 10^{-1}; \beta_{\text{CLS-prior}} = 1.0 \times 10^{-2};$
 1572 $\text{lr} = 3.0 \times 10^{-4};$

1574 I.6 RESISC

1576 **CODA-Prompt**: $\text{lr} = 1.0 \times 10^{-3};$

1577 **DER++**: $\alpha = 3.0 \times 10^{-1}; \beta = 8.0 \times 10^{-1}; \text{lr} = 1.0 \times 10^{-3};$

1578 **EWC**: $\epsilon = 1.0 \times 10^2; \gamma = 9.0 \times 10^{-1}; \text{lr} = 1.0 \times 10^{-2};$

1580 **L2P**: $\text{lr} = 2.5 \times 10^{-3};$

1581 **LWF-MC**: $\text{wd} = 0.0; \text{lr} = 1.0 \times 10^{-2};$

1583 **SEED**: $\text{lr} = 3.0 \times 10^{-4};$

1584 **SGD**: $\text{lr} = 1.0 \times 10^{-2};$

1585 **TMC**: $\text{lr} = 3.0 \times 10^{-4};$

1587 **InfLoRA**: $\text{lr} = 5.0 \times 10^{-4} \ r = 10; \epsilon = 0.98;$

1589 For both ITA and IEL: $\# \text{epochs}_{\text{pre-tuning}} = 8; \text{lr}_{\text{pre-tuning}} = 1.0 \times 10^{-2}.$

1590 **ITA-FFT**: $\alpha = 1.0 \times 10^4; \alpha_{\text{CLS}} = 1.0 \times 10^{-1}; \alpha_{\text{CLS-prior}} = 1.0 \times 10^{-2}; \text{lr} = 1.0 \times 10^{-4};$

1591 **ITA-LoRA**: $\alpha = 2.0 \times 10^{-2}; \alpha_{\text{CLS}} = 1.0 \times 10^{-1}; \alpha_{\text{CLS-prior}} = 1.0 \times 10^{-2}; \text{lr} = 3.0 \times 10^{-4};$

1592 **ITA-(IA)³**: $\alpha = 5.0; \alpha_{\text{CLS}} = 1.0 \times 10^{-1}; \alpha_{\text{CLS-prior}} = 3.0 \times 10^{-3}; \text{lr} = 3.0 \times 10^{-3};$

1593 **IEL-FFT**: $\beta = 5.0 \times 10^3; \beta_{\text{CLS}} = 1.0 \times 10^{-1}; \beta_{\text{CLS-prior}} = 3.0 \times 10^{-3}; \text{lr} = 1.0 \times 10^{-4};$

1594 **IEL-LoRA**: $\beta = 2.0 \times 10^2; \beta_{\text{CLS}} = 1.0 \times 10^{-2}; \beta_{\text{CLS-prior}} = 1.0 \times 10^{-3}; \text{lr} = 3.0 \times 10^{-4};$

1595 **IEL-(IA)³**: $\beta = 2.0; \beta_{\text{CLS}} = 1.0 \times 10^{-2}; \beta_{\text{CLS-prior}} = 1.0 \times 10^{-3}; \text{lr} = 3.0 \times 10^{-4};$

1600 I.7 CROPDISEASE

1601 **CODA-Prompt**: $\text{lr} = 1.0 \times 10^{-3};$

1602 **DER++**: $\alpha = 3.0 \times 10^{-1}; \beta = 8.0 \times 10^{-1}; \text{lr} = 1.0 \times 10^{-3};$

1603 **EWC**: $\epsilon = 1.0; \gamma = 1.0; \text{lr} = 1.0 \times 10^{-2};$

1604 **L2P**: $\text{lr} = 2.5 \times 10^{-3};$

1605 **LWF-MC**: $\text{wd} = 1.0 \times 10^{-4}; \text{lr} = 1.0 \times 10^{-2};$

1606 **SEED**: $\text{lr} = 3.0 \times 10^{-4};$

1607 **SGD**: $\text{lr} = 1.0 \times 10^{-2};$

1608 **TMC**: $\text{lr} = 1.0 \times 10^{-4};$

1609 **InfLoRA**: $\text{lr} = 5.0 \times 10^{-4} \ r = 10; \epsilon = 0.98;$

1615 For both ITA and IEL: $\# \text{epochs}_{\text{pre-tuning}} = 8; \text{lr}_{\text{pre-tuning}} = 1.0 \times 10^{-2};$

1616 **ITA-FFT**: $\alpha = 5.0 \times 10^3; \alpha_{\text{CLS}} = 2.5 \times 10^{-2}; \alpha_{\text{CLS-prior}} = 1.0 \times 10^{-2}; \text{lr} = 1.0 \times 10^{-4};$

1617 **ITA-LoRA**: $\alpha = 5.0; \alpha_{\text{CLS}} = 1.0 \times 10^{-1}; \alpha_{\text{CLS-prior}} = 1.0 \times 10^{-2}; \text{lr} = 3.0 \times 10^{-4};$

1618 **ITA-(IA)³**: $\alpha = 2.0 \times 10^{-2}; \alpha_{\text{CLS}} = 1.0 \times 10^{-1}; \alpha_{\text{CLS-prior}} = 1.0 \times 10^{-2}; \text{lr} = 3.0 \times 10^{-4};$

1620 **IEL-FFT**: $\beta = 2.0 \times 10^2$; $\beta_{\text{CLS}} = 1.0 \times 10^{-1}$; $\beta_{\text{CLS-prior}} = 1.0 \times 10^{-2}$; lr = 1.0×10^{-5} ;
1621
1622 **IEL-LoRA**: $\beta = 5.0 \times 10^1$; $\beta_{\text{CLS}} = 1.0 \times 10^{-1}$; $\beta_{\text{CLS-prior}} = 1.0 \times 10^{-2}$; lr = 3.0×10^{-4} ;
1623 **IEL-(IA)³**: $\beta = 2.0$; $\beta_{\text{CLS}} = 2.5 \times 10^{-2}$; $\beta_{\text{CLS-prior}} = 1.0 \times 10^{-2}$; lr = 3.0×10^{-4} ;
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673