

Tangent Model Composition for Ensembling and Continual Fine-tuning

Tian Yu Liu
 University of California, Los Angeles
 tianyu@cs.ucla.edu

Stefano Soatto
 University of California, Los Angeles
 soatto@ucla.edu

Abstract

Tangent Model Composition (TMC) is a method to combine component models independently fine-tuned around a pre-trained point. Component models are tangent vectors to the pre-trained model that can be added, scaled, or subtracted to support incremental learning, ensembling, or unlearning. Component models are composed at inference time via scalar combination, reducing the cost of ensembling to that of a single model. TMC improves accuracy by 4.2% compared to ensembling non-linearly fine-tuned models at a $2.5\times$ to $10\times$ reduction of inference cost, growing linearly with the number of component models. Each component model can be forgotten at zero cost, with no residual effect on the resulting inference. When used for continual fine-tuning, TMC is not constrained by sequential bias and can be executed in parallel on federated data. TMC outperforms recently published continual fine-tuning methods almost uniformly on each setting – task-incremental, class-incremental, and data-incremental – on a total of 13 experiments across 3 benchmark datasets, despite not using any replay buffer. TMC is designed for composing models that are local to a pre-trained embedding, but could be extended to more general settings. The code is available at: <https://github.com/tianyu139/tangent-model-composition>

1. Introduction

Compositionality has long been considered central to cognition [5], possibly a reflection of compositionality in neural activity [50]. After decades of unsuccessful attempts to design compositional representations, deep neural networks trained on language data are beginning to exhibit emergent compositionality at scale, by capturing the compositional structure of the data. These models are now being used for visual inference, which has rekindled interest in the study of compositionality of visual representations. But beyond the data, any compositional structure that may be latent in the model is not directly accessible: One cannot simply “compose” weights or inner activations and expect

meaningful outcomes. The computational architecture of Transformers [52] has been leveraged extensively to co-opt the compositional structure of data through prompts or tokens [29, 55], but still the activations of trained models do not appear to be meaningfully composable. Compositionality of neural activity would allow one to combine activations from different models to capture novel concepts, or incorporate knowledge from different data without having to re-train or fine-tune the core models. This would enable open-universe classification and, more generally, combinatorial expansion of the hypothesis space. Continual learning could be performed simply by composing models trained on different data.

In this paper, we explore the simplest form of compositionality, that is *linear combination*. We leverage recent results on the linearization of deep neural networks around a pre-trained point, that can be trained by solving a convex optimization problem and yet perform on-par with non-linear fine-tuning [1]. This suggests that the tangent space at pre-trained models may be used to linearly compose neural activations, or equivalently compose different models trained or fine-tuned on different datasets and/or for different tasks.

We show that different models obtained from the linearization around a pre-trained point can be composed, combined, rescaled, and forgotten (“unlearning”), simply by scalar combinations. This fact can be used to perform ensembling at the inference cost of a single model (Sec. 3). It can also be used for continual fine-tuning, with each component model trained independently and in parallel, if necessary on federated data that can therefore be easily forgotten if needed.

Linear combination is not viable for general concept composition nor arbitrary multi-task learning. In Sec. 5, we discuss limitations of our approach, which can only meaningfully compose models that are “local” to a pre-trained embedding. If there are models trained on tasks that are *far* in representation space, or even *antagonistic*, they are likely to live on different tangent planes, making linear combination inappropriate. One such example, described in the appendix, concerns models trained on real images (Ima-




Figure 1. **Composition of Fine-tuned Models:** Purple indicates final models used for inference, yellow indicates specialist models fine-tuned on individual tasks. The paragon (a) is a model trained jointly on all tasks. However, in continual learning, different tasks are instantiated at different times, an re-training on their union is impractical. Ensembling (b) combines the output of different models trained separately on each task. Weight averaging (c) yields a “model soup” of fine-tuned non-linear models, which improves inference time; Tangent Model Composition (d) linearly composes models fine-tuned on the tangent space of a pre-trained model.

geNet) and clip art or sketches (DomainNet). Nonetheless, in more common settings, TMC is competitive with general forms of ensembling such as averaging activations or logits of non-linearly fine-tuned models, and with more general forms of continual learning, including methods that use a replay buffer. For broader task coverage, one could use TMC to construct a *tangent bundle of models* around different pre-trained points, akin to a “tangent model zoo.” This concept is not too dissimilar from the architecture of some Foundation Models with a shared backbone and multiple distinct “heavy heads” [54], but far more compact, modular, efficient, interpretable, and easy to work with.

In the next section we place our contributions in the context of prior art, and in the following one we describe our method in more detail. Sec. 4 summarizes empirical evidence in support of our approach, and finally Sec. 5 discusses its main limitations.

2. Related work

Compositionality: Compositionality has been studied for decades as a means to expand the representative power of trained models, but has received increasing impetus in recent years thanks to large Transformer-based models that can be used with adaptable prompts. Prompt-tuning [29] has become commonplace to adapt large pre-trained models for specific downstream tasks. Composition of prompts has also been explored in [56, 55] for continual learning. More

general composition of parameters of trained deep networks has been explored for various purposes such as improving optimization and generalization, as we describe next.

Deep network linearization: Deep networks are linearized using the first-order Taylor expansion around a pre-trained weight [1, 18, 35, 49, 32]. [1] finetunes linearized networks modified to use Leaky-ReLU with gradient pre-conditioning to achieve similar performances to non-linear fine-tuning. We note that TMC is not just a matter of applying linearized networks to different datasets and averaging, since different scales would lead to imbalance as we describe in Sec. 3.3. Furthermore, unlike [1, 49], we do not require Leaky-ReLU nor gradient pre-conditioning.

Averaging model weights: Weight averaging has been used in deep learning to improve generalization [25, 17], improve pre-training [12, 34], perform distributed fine-tuning [57], and increase robustness against distribution shift [59]. [25] averages points along the trajectory of SGD to find flatter minima, [53] averages “late-phase weights” obtained from the later stages of SGD. [24] uses weight interpolation for model patching, and [58] averages weights of models fine-tuned using different hyperparameters from pre-trained models to improve generalization.

While weight averaging techniques for deep networks are related to ensembles [14, 19], in our case the two coincide since we operate in representation spaces that are linear in the parameters. We leverage this property to train on

disjoint subsets of data, and show that they can outperform existing weight averaging and ensembling techniques.

Continual learning/fine-tuning: Continual learning [37] aims to adapt models to new tasks and data without degrading performance in previously trained tasks. Continual fine-tuning [49, 6] additionally requires minimal forgetting of a pre-training objective to ensure all tasks can benefit from it. Existing approaches can be broadly classified into parameter isolation, regularization, architecture, and replay-based.

Parameter isolation methods [32, 3, 47] operate in the network architecture space by allocating each task a different set of parameters, using techniques such as network pruning [33] to reduce additional memory requirements.

Regularization-based methods [30, 2, 26, 36, 28] incorporate additional terms into the training objective to prevent catastrophic forgetting. [26] penalizes the distance between weights of the previous and current task, and [45] penalizes changes in weights using Hessian approximations.

Experience replay methods [44, 4, 32, 11, 43, 49, 7, 8, 26, 56, 6, 10, 9] assume the availability of a memory buffer to store samples from previous tasks. [49, 6] applies experience replay to continual fine-tuning. [49] shows that linearized models combined with replay and regularization alleviates catastrophic forgetting when fine-tuned, and [6] combines dark experience replay and an attention-based regularization loss to prevent catastrophic forgetting of the pre-training features, by encouraging similarity of features at each layer to the original pre-trained network. While replay methods have achieved state-of-the-art performance compared to non-replay ones, they require a sufficiently large memory buffer to store previously seen exemplars.

A further line of work tackles exemplar-free class incremental learning (EFCIL) using regularization [60], distillation [22], and class-prototypes [61, 62, 40] to incrementally integrate new classes into a fixed-sized network without the memory costs of a replay buffer.

We show that composition of tangent models trained on individual tasks is a viable method for replay-free continual fine-tuning, and can even outperform recently published replay-based methods. As a side benefit, our method also yields fully parallel training across tasks, and enables zero-cost unlearning.

3. Method

In continual learning, models are trained on a sequence of tasks indexed by $t = 1, \dots, T$, with the goal of performing well on all, under three main scenarios [51]: (1) *Task-incremental*, where the task identity t (Task-ID) is known during inference, (2) *Data/Domain-incremental*, where Task-ID is not provided during inference, and (3) the most challenging *Class-incremental*, where Task-ID is unknown during inference and in addition the classes repre-

Table 1. **A Comparison of Composition Methods:** On average across 25 experiments, Tangent Model Ensembling (TME) improves accuracy by 5.1% over standard ensembling of non-linear fine-tuned models at a $2\times$ increase of inference cost. Tangent Model Composition (TMC) also improved accuracy, by 4.2%, while *reducing* inference cost by $2.5\times$ to $10\times$ in our experiments, and growing linearly with the number of models in the ensemble. Best results for single and multi model inference are indicated in bold and italic respectively. Joint training paragon accuracy (Original Model/Tangent Model): Caltech-256 - 89.17%/86.77%, MIT-67 - 77.29%/77.74%, OxfordPets - 92.68%/93.68%.

Dataset	Tasks	Soup	Ens-L	Ens-SM	TMC	TME
Inference/Memory Cost: $\mathcal{O}(1)$ $\mathcal{O}(T)$ $\mathcal{O}(T)$ $\mathcal{O}(1)$ $\mathcal{O}(T)$						
C-Caltech-256 (Class-IL)	5	84.00	82.80	83.09	84.82	85.53
	10	80.10	76.85	79.76	82.37	82.58
	20	74.86	66.76	75.47	78.66	79.30
C-MIT-67 (Class-IL)	5	59.48	56.79	58.68	69.43	71.72
	10	49.30	42.36	50.52	64.53	65.72
	20	28.51	23.61	35.62	48.98	54.93
C-OxfordPets (Class-IL)	5	80.71	81.14	81.86	84.42	85.04
	10	71.35	69.71	77.31	77.75	79.97
	5	85.96	86.27	86.59	84.86	85.13
D-Caltech-256 (Data-IL)	10	83.48	84.05	84.80	83.64	84.16
	20	79.64	80.72	83.17	81.99	82.94
	5	66.00	66.54	67.21	74.88	75.45
D-MIT-67 (Data-IL)	10	58.88	59.20	60.47	72.41	73.41
	20	50.10	51.99	55.20	70.22	72.11
	5	90.22	90.83	91.02	93.17	93.21
D-OxfordPets (Data-IL)	10	87.66	88.42	88.88	92.48	92.72
	20	85.57	85.97	86.58	91.28	91.77
	5	94.01	94.72	94.84	94.33	94.76
C-Caltech-256 (Task-IL)	10	95.00	95.87	96.04	96.03	96.23
	20	95.56	96.53	96.84	97.22	97.38
	5	83.36	86.24	86.37	91.27	91.79
C-MIT-67 (Task-IL)	10	87.19	90.42	90.97	94.45	95.03
	20	88.41	93.73	93.88	97.71	98.09
	5	97.27	97.27	97.35	98.57	98.74
C-OxfordPets (Task-IL)	10	97.59	97.92	97.98	99.05	99.13

sented in each task are disjoint.

We assume the setting of continual fine-tuning of a pre-trained model. Naive fine-tuning causes catastrophic forgetting of prior tasks, including the one on which pre-training occurred, as the number of tasks increases [6]. Thus, only early tasks benefit from pre-training. Therefore, it is important for continual fine-tuning to ensure that all tasks can benefit from information acquired during pre-training.

We introduce a unified method to tackle the main scenarios (1)-(3) by composing separately-trained tangent component models. We show that each task takes full advantage of pre-training, and catastrophic forgetting of prior tasks is averted since information from each previous task is encoded separately and only fused for inference.

3.1. Model Linearization

The effectiveness of model aggregation in weight space hinges on whether the models lie in within the same small-loss training basin of convergence. It is known that weights of different models fine-tuned on the same dataset can be connected by linear paths with constant loss (“linear mode connectivity”) [16], but this does not generally occur for models fine-tuned on disjoint tasks/datasets, which is the case of interest to us. Therefore, we convexify the loss landscape locally by approximating the function represented by a deep (non-linear) network with its first-order Taylor approximation around the pre-trained weights. For any deep network f and initial weights $w \in \mathbb{R}^m$, we denote with $\mathcal{H}(w)$ the set of all such models tangent to f_w given by:

$$\mathcal{H}(w) = \{h_\delta(\cdot) \triangleq f_w(\cdot) + \nabla_w f_w(\cdot) \cdot \delta\}. \quad (1)$$

Since the tangent model h_δ is linear in the parameters δ , training it with standard losses such as mean-squared error (MSE) or empirical cross-entropy yields a convex loss landscape. In order to compose models, leveraging Jensen’s inequality, we are guaranteed that for any dataset \mathcal{D} :

$$L(\mathcal{D}, \sum_i^T h_{\theta_i \delta_i}) \leq \sum_i \theta_i L(\mathcal{D}, h_{\delta_i}), \quad \sum_i \theta_i = 1. \quad (2)$$

In other words, the loss resulting from a convex combination of tangent models in weight space will always be lower than or equal to the convex combination of their individual losses. This is particularly useful when combining models trained to some constant (possibly zero) loss, since it ensures that the composed model will have loss lower or equal to the loss of the component models.

While tangent models might seem computationally expensive to train, we note that the Jacobian vector product $\nabla_w f_w(\cdot) \cdot \delta$ can be computed in a single forward pass as shown by [39, 35, 1]. This reduces the computational cost to that of a forward pass through the original network. Also note that while the tangent model has twice the degrees of freedom of the original model, fixing the pre-trained model at w results in only a convex optimization in the remaining parameters, as many as the original model.

In the context of fine-tuning, w is initialized by pre-training on large datasets such as ImageNet [13]. In the rest of this paper, we use the short-form notation \mathcal{H} and h for $\mathcal{H}(w)$ and h_δ respectively.

3.2. Tangent Model Composition for Replay-free Continual Fine-tuning

Ensembling is an effective strategy for combining weak learners into a strong model [14]. For any finite $F \subseteq \mathbb{A}$, where \mathbb{A} is a set of functions, and $\lambda = \{\lambda_i\}_{i=1}^{|F|}$ a set of

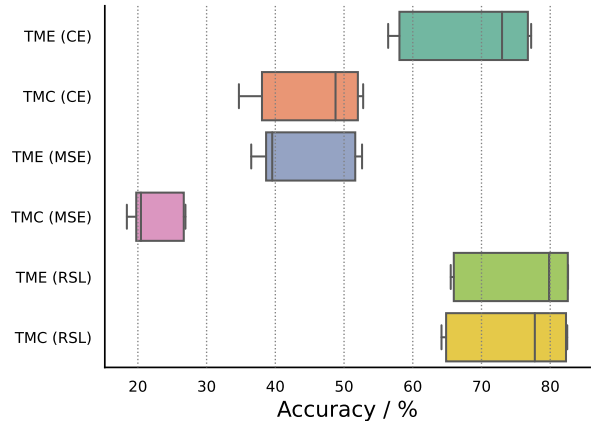


Figure 2. Accuracy plot over models trained on C-Caltech-256, C-MIT-67, and C-Oxford Pets, split into 10 tasks each. Soft-max ensembling (TME) consistently yields the best results compared to logit ensembling (TMC) for each loss function. Composition of component models trained with the Rescaled Square Loss (RSL) outperforms those trained with empirical cross-entropy (CE) or MSE loss, trading off only 1.2% accuracy against TME to reduce inference time and memory requirement from $\mathcal{O}(T)$ to $\mathcal{O}(1)$.

weights, the logit ensemble operation is given by

$$Ensemble(F, \lambda) = \sum_{i=1}^{|F|} \lambda_i f_i(\cdot) \quad (3)$$

In the context of continual learning, a simple yet effective approach is to simply train a new model on each task t , and ensemble the outputs at inference time. Unfortunately, this has the obvious shortcoming that both storage requirements and inference time grows along with the number of tasks.

However, consider the case where $\mathbb{A} = \mathcal{H}$. Then for any finite $H \subseteq \mathcal{H}$ and associated weights λ , we can write

$$Ensemble(H, \lambda) = \sum_{i=1}^{|H|} \lambda_i h_i(\cdot) \quad (4)$$

$$= f_w(\cdot) + \sum_{i=1}^{|H|} \nabla_w f_w(\cdot) \cdot \lambda_i \delta_i \quad (5)$$

$$= h_{\sum \lambda_i \delta_i}(\cdot) \quad (6)$$

In other words, since \mathcal{H} defines a vector space with addition and scalar multiplication based on δ , any linear combination (or ensemble) of models in \mathcal{H} is equivalent to a single model in \mathcal{H} derived simply by the linear combination over δ . This reduces the inference time for any ensemble of T tangent models from $\mathcal{O}(T)$ to $\mathcal{O}(1)$.

In continual learning, it is also important for storage space to not scale with T , since the number of tasks can be arbitrarily large. However, since tasks arrive sequentially in the continual learning framework, we can further reduce

Table 2. **Class-Incremental Learning:** Comparison against existing methods using an ImageNet pre-trained ResNet-50. TMC is almost uniformly better than all replay-based and replay-free methods tested, outperforming the best replay-free (RF) method by 14.48% and even the best replay-based method [6] by 1.53%.

Method / Buffer	C-Caltech-256			C-MIT-67			
	Tasks	5	10	20	5	10	20
ER [46] / 300		70.14 \pm 0.22	65.75 \pm 0.92	65.49 \pm 0.60	54.61 \pm 0.41	49.57 \pm 1.46	45.33 \pm 1.73
LWF [30] / None		34.55 \pm 0.18	22.76 \pm 0.17	14.37 \pm 0.14	25.14 \pm 0.57	13.30 \pm 0.49	6.41 \pm 0.17
oEWC [48] / None		28.71 \pm 1.31	16.32 \pm 4.37	15.59 \pm 3.83	14.18 \pm 0.77	7.12 \pm 0.49	4.33 \pm 0.16
DER++ [7] / 300		80.50 \pm 0.26	77.84 \pm 1.02	74.46 \pm 0.21	64.56 \pm 0.55	59.73 \pm 1.02	50.28 \pm 1.27
Co2L [10] / 300		20.71 \pm 2.89	17.13 \pm 2.53	14.62 \pm 0.61	19.64 \pm 12.93	21.21 \pm 0.50	17.78 \pm 1.33
ER-ACE [9] / 300		78.47 \pm 1.25	73.85 \pm 0.53	70.87 \pm 1.38	64.94 \pm 0.50	60.25 \pm 0.26	52.48 \pm 0.29
TWF (RF) [6] / None		77.95 \pm 0.52	70.67 \pm 0.72	62.79 \pm 0.82	56.58 \pm 0.62	45.03 \pm 0.78	28.91 \pm 0.92
TWF [6] / 300		81.34 \pm 0.91	77.03 \pm 0.53	72.44 \pm 0.42	68.33 \pm 0.46	64.21 \pm 0.65	56.29 \pm 0.68
TMC / None		84.82 \pm 0.03	82.37 \pm 0.08	78.66 \pm 0.13	69.43 \pm 0.21	64.53 \pm 0.28	48.98 \pm 0.13

the memory capacity required for storage of models from $\mathcal{O}(T)$, to $\mathcal{O}(1)$ by using a simple autoregressive model: We train on the first task t_1 to produce model $\tilde{h}_1 = h_1$. Then, for any subsequent task $t > 1$, we train a new model \tilde{h}_t , and compose it with h_{t-1} to give $h_t = \lambda_{(t,1)}\tilde{h}_t + \lambda_{(t,2)}h_{t-1} := h_{\delta'}$ where $\delta' = \frac{\lambda_{(t,1)}\delta_t + \lambda_{(t,2)}\delta_{t-1}}{\lambda_{(t,1)} + \lambda_{(t,2)}}$. Since we only require model h_{t-1} to produce h_t , all prior models up to and including h_{t-1} can be discarded upon constructing h_t , hence reducing memory requirement from $\mathcal{O}(T)$ to $\mathcal{O}(1)$, equal to the size of a single model. In contrast to methods that expands the model at each task, or require a memory buffer of examples from previous tasks, this method does not incur additional storage requirements, nor assume availability of an extra replay buffer for storing past exemplars (*i.e.* our method is replay-free).

Remark 1 (Free Parallelization): Additionally, for any given convex loss function L , the optimization objective

$$\min_{\delta \in \mathbb{R}^m} \sum_{(x,y) \in \mathcal{D}} L(h_{\delta}(x), y) \quad (7)$$

is convex and thus converges to a global optimum. This implies that each model \tilde{h}_t can be trained from a standard initialization as opposed to from h_{t-1} . Hence, if multiple tasks arrive at the same time, such as when training in a federated manner, training can be parallelized across tasks and composed in a zero-cost, zero-shot manner to yield equivalent results (assuming equivalence of global optima) to sequential training and composition. To the best of our knowledge, apart from [41] which simply stores samples in a replay buffer and trains on them at test time, no other work in continual learning can be parallelized across tasks.

Remark 2 (Free Forgetting/Unlearning): If we relax memory constraints to store models from each task, unlearning any specific task t (*e.g.*, private data associated with specific samples) is straightforward and can be done in a zero-cost, zero-shot manner. Since subtraction is well-defined in the vector space \mathcal{H} , any task i can be unlearned

from the model h_T trained on T tasks by simply subtracting the weights associated to \tilde{h}_i :

$$h_{T \setminus i} = h_T - \lambda_{i,1} \prod_{j=i+1}^T \lambda_{(j,2)} \tilde{h}_i \quad (8)$$

Remark 3 (Unified Continual Fine-tuning): Our method does not assume any specific task splits, and is hence generalizable across all three forms of continual fine-tuning: task-incremental, class-incremental, and data-incremental.

3.3. Scale Standardization

Training with the cross-entropy loss does not ensure that the weights are comparable in scale. We can study the implications of this for composition on \mathcal{H} through the lens of logits ensembling. While logits ensembling has been shown to be effective [58], Fig. 2 shows that ensembling the normalized soft-max outputs (Tangent Model Ensembling, TME) is uniformly better than logit ensembling/weight averaging via TMC across various loss functions. However, TME cannot be represented by a linear combination of perturbations δ_i and hence incurs $\mathcal{O}(T)$ model storage costs and $\mathcal{O}(T)$ inference time.

Instead, we propose to induce a waterbed effect similar to soft-max ensembling using standardization. In particular, we train using the Rescaled Square Loss (RSL) [23]:

$$L_s(x, y) = \frac{1}{K} (\alpha ([f(x)]_y - \beta)^2 + \sum_{i=1, i \neq y}^K ([f(x)]_i)^2) \quad (9)$$

where α, β are hyper-parameters. This is the regular MSE loss when $\alpha = \beta = 1$. Larger values of β reduce soft-max entropy in the output by scaling positive output class signals while keeping negative ones close to zero.

We use β to control the interaction between component models. For example, when tasks are highly dissimilar,

Table 3. Data-Incremental Learning on D-MIT-67, 4 tasks using ResNet-18. For fair comparison, we run our method under the settings proposed by [49]. TMC outperforms the best replay-free method [30] by 4.61%, and the best replay-based method [49] by 1.27%, while supporting parallel training across tasks (PL). If we sacrifice parallelization by initializing each model at task t with h_{t-1} (TMC-Seq), we further improve over TMC by 0.28%.

Method	PL?	Replay Buffer	Accuracy
SGD	✗	None	63.48
LwF[30]	✗	None	67.21
MAS[2]	✗	None	62.49
OSLA[45]	✗	None	64.08
EWC [26]	✗	500	63.68
DLCFT[49]	✗	500	70.55
TMC	✓	None	71.82
TMC-Seq	✗	None	72.10
Joint (Linear)	-	-	73.61
Joint (Paragon)	-	-	74.40

such as in task- and class-incremental continual learning, we use larger values of β to reduce interference of component models trained on other tasks upon composition. On the other hand, when tasks are highly synergistic such as in data-incremental learning, lower values of β can be used to encourage interaction in the final composed model.

4. Experiments

We describe comparison baselines in Sec. 4.1, and our main results in Sec. 4.2, with additional studies in Sec. 4.3.

Our method is replay-free. As such, its computational cost is at most half of what it would be if it used a replay buffer since replay-based methods typically sample separate batches of equal size from both the current task and replay buffer at each iteration [7, 6, 10, 46, 9]. Nonetheless, we compare our method to both replay-based and replay-free methods in our baselines. While replay-based methods have thus far outperformed replay-free ones, ours is uniformly more accurate than replay-based methods tested, in addition to replay-free ones.

We evaluate our method on Caltech-256 [20], MIT-67 [42], and OxfordPets [38]. For continual fine-tuning, we further split each training dataset into multiple disjoint subsets, each representing a single task, for fair comparison across experiments. We prefix sequential datasets split into tasks with disjoint labels (for task- and class-incremental learning) with C-, and random splits using a standardized random seed (for data-incremental learning) with D-. We set $\alpha = 1, \beta = 25$ for the former, $\alpha = 1, \beta = 5$ for the latter, and $\alpha = 1, \beta = 5$ for TME. Results in all tables and figures are averaged across 3 runs.

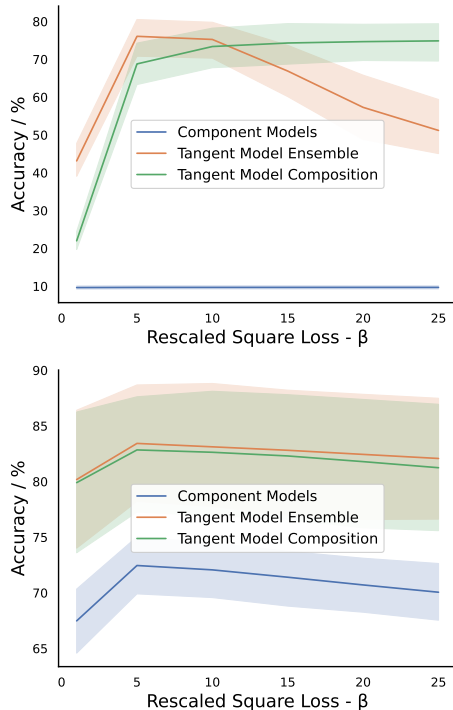


Figure 3. Ablation on β averaged across (**Top:**) C-Caltech-256, C-MIT-67, and C-OxfordPets and (**Bottom:**) D-Caltech-256, D-MIT-67, and D-OxfordPets, split into 10 disjoint tasks. For class-incremental datasets, larger values of β result in better generalization of the final composed model. For data-incremental datasets where tasks are synergistic, larger values of β instead harm the accuracy of both the component models and the composed model.

4.1. Composition Baselines

In each of the following sections, we compare against the following baselines for model aggregation:

Logits Ensemble (Ens-L): Averaging the output logits of non-linear component models.

Softmax Ensemble (Ens-SM): Averaging the softmax output of non-linear component models.

Weight composition of non-linear models (Soup): Averaging the weights of non-linear component models. This method is inspired by Model Soups [58] where weights of various non-linear models trained on the same task but with different hyper-parameters and augmentations are averaged to improve generalization at no additional inference cost.

4.2. Main Results

We show that TMC outperforms existing methods in all three continual fine-tuning settings [51], from hardest to easiest: Class-Incremental (Sec. 4.2.1), Data-Incremental (Sec. 4.2.2), and Task-Incremental (Sec. 4.2.3).

Table 4. **Task-Incremental Learning:** Comparison against existing methods using ImageNet pre-trained ResNet-50. We almost uniformly outperform all replay-based and replay-free methods that we compare against. In particular, we improve over the best replay-free (RF) method by 2.46%, and the best replay-based method by 1.38% [6].

Method / Buffer	C-Caltech-256			C-MIT-67			
	Tasks	5	10	20	5	10	20
ER [46] / 300		92.11 \pm 0.26	94.33 \pm 0.48	95.76 \pm 0.13	85.84 \pm 0.09	90.59 \pm 0.51	95.13 \pm 0.27
LWF [30] / None		93.73 \pm 0.12	95.31 \pm 0.08	96.29 \pm 0.05	87.55 \pm 0.20	90.82 \pm 0.04	95.68 \pm 0.58
oEWC [48] / None		83.97 \pm 0.71	76.20 \pm 16.07	91.49 \pm 0.30	67.31 \pm 3.22	68.45 \pm 4.93	84.74 \pm 1.96
DER++ [7] / 300		93.79 \pm 0.13	94.92 \pm 0.11	95.87 \pm 0.17	86.83 \pm 0.26	91.63 \pm 0.75	94.54 \pm 0.62
Co2L [10] / 300		34.51 \pm 3.05	39.62 \pm 1.70	45.94 \pm 1.47	39.32 \pm 20.47	64.40 \pm 1.31	51.50 \pm 2.21
ER-ACE [9] / 300		93.11 \pm 0.12	94.63 \pm 0.06	95.86 \pm 0.10	87.13 \pm 0.16	90.96 \pm 0.28	94.92 \pm 0.60
TWF (RF) [6] / None		94.45 \pm 0.05	95.64 \pm 0.06	96.36 \pm 0.12	86.76 \pm 0.25	88.27 \pm 0.54	94.74 \pm 0.35
TWF [6] / 300		94.38 \pm 0.08	95.74 \pm 0.17	96.59 \pm 0.07	88.22 \pm 0.41	92.13 \pm 0.26	95.67 \pm 0.46
TMC / None		94.33 \pm 0.03	96.03 \pm 0.02	97.22 \pm 0.03	91.27 \pm 0.21	94.45 \pm 0.13	97.71 \pm 0.13

4.2.1 Class-Incremental Learning

In class-incremental learning, datasets are partitioned so that label spaces of each task are disjoint.

Comparison against Composition Baselines: We compare against model composition baselines in Table 1. While TME yields the best generalization performance, this comes at a cost of $\mathcal{O}(T)$ inference time. On the other hand, TMC achieves an $\mathcal{O}(1)$ inference time with only a 1.73% trade-off in generalization accuracy on average, while outperforming weight averaging, logit ensembling, and even soft-max ensembling of non-linear models by an average of 7.83%, 11.37%, and 6.08% respectively, even though the latter two both require $\mathcal{O}(T)$ inference time.

Comparison against existing methods: In Table 2, we compare against recent continual fine-tuning methods including both replay-free and replay-based methods. Here, replay-free methods [30, 48] perform poorly in the class-incremental setting relative to stronger replay-based methods [46, 9, 7, 6]. However, we show that not only do we improve over the best replay-free method by 14.48%, but also over the best replay-based method by 1.53% [6], corresponding to a 17.59% relative error reduction towards the paragon performance of joint training (Table 1).

Comparison against EFCIL methods: In Table 5, we compare against EFCIL methods on C-TinyImageNet [27]. The original EFCIL methods are evaluated under a different setting (denoted †), where the zeroth task contains 50% of the dataset. Hence, we further modified two of the best performing EFCIL methods, SSRE [62] and FeTrIL [40] to use the pre-trained ImageNet initialization, and run them under our setting of uniformly sharded data. TMC not only outperforms all methods ran under our harder setting, by an average of 15.3% over the next best method [40], but even yields equal or better performance when compared to all methods ran under †.

4.2.2 Data-Incremental Learning

Also referred to as domain-incremental learning, this setting assumes that tasks are split in a non-stratified manner, where each task has the same output space [51, 49].

Comparison against Composition Baselines: In Table 1, we compare against model composition baselines. Similar to the class-incremental setting, TME outperforms TMC by an average of 0.66% at a cost of $\mathcal{O}(T)$ inference time. TMC uses $\mathcal{O}(1)$ inference time and outperforms non-linear composition, logit ensemble, and even soft-max ensemble by an average of 7.04%, 6.32%, and 5.21% respectively.

Comparison against existing methods: In Table 3, we benchmark our model against results obtained by [49] on a pre-trained ResNet-18 model using the D-MIT-67 dataset split into 4 disjoint subsets. We improve over the best replay-free method [30] by 4.61%, and the best replay-based method [49] by 1.27%, corresponding to a 64% and 33% relative error reduction towards the joint training paragon respectively. TMC-Seq, which uses sequential initialization by initializing each component model h_t with h_{t-1} instead of h_0 , further improves over TMC by 0.28% but at the cost of parallelizability.

4.2.3 Task-Incremental Learning

In task-incremental learning, the output spaces of different tasks are disjoint, but the source task of each sample is known at inference time. Thus, we can restrict the model predictions to that of the given task during evaluation.

Comparison against Composition Baselines: Due to the disjoint output spaces of each task, ensembling outputs of specialist component models trained on individual tasks and restricting them to a target task essentially only considers the predictions of the specialist model trained on the target task (since, in general, specialist models are not effective discriminants for tasks other than those they are trained on). Hence, the independence of output spaces renders en-

Table 5. Comparison against EFCIL methods using either ImageNet pre-trained ResNet-50, or under the original EFCIL setting where the zeroth task contains 50% of the dataset (denoted †). All experiments are ran on C-TinyImageNet. *Italics* denote the best method under †, and **bold** denotes the best method under our harder setting. Under the former, TMC performs equally or better than all methods. Under the latter, TMC outperforms the best method by an average of 15.3%.

Tasks	PASS† [61]	IL2A† [60]	SSRE† [62]	FeTrIL† [40]	SSRE[62]	FeTrIL[40]	TMC
5	49.6	47.3	50.4	54.8	33.7	44.1	57.6
10	47.3	44.7	48.9	53.1	25.6	36.1	53.1

sembles of specialist models highly effective. In contrast, composing non-linear models in weight space does not ensure minimal interference across different tasks. Indeed, we show in Table 1 that TME yields the best results, while the non-linear composition approach yields the worst in task-incremental learning. On the other hand, due to the equivalence of TMC to (logit) ensembling of linearized models, TMC outperforms not only weight averaging, but also logit and soft-max ensembles of non-linear models, by 3.78%, 1.83%, and 1.79% respectively.

Comparison against existing methods: Since each component model is trained only on task t , by design, there is no interference from other tasks when training \tilde{h}_t . As such, the final composed model is equivalent to a specialist model when restricted to any task. However, in other state of the art continual learning methods such as those involving replay or loss regularization, this separation among trained for each task is violated. Hence, under the task-incremental setting, we show in Table 4 that TMC achieves the best performance across both replay-free and replay-based methods even without using any memory buffer, beating the next best method by 2.46% and 1.38% respectively [6].

4.3. Further studies

We present additional studies on the Rescaled Square Loss (Sec. 4.3.1), and compare against TMC-Seq to show that our method yields an effective and highly parallelizable training scheme for continual fine-tuning (Sec. 4.3.2).

4.3.1 Rescaled Square Loss

In Sec. 3.3, we showed that β can be used to control the amount of interaction between tasks. In Fig. 3, we empirically show that in the class-incremental setting where tasks are highly dissimilar, larger values of β result in a better generalist model. On the other hand, under the data-incremental setting where tasks are more similar, information from other tasks can benefit the learning of any given task. Here, larger values of β instead harm generalization due to lowering the accuracy of each component model.

4.3.2 Convexity and Parallelizability

Given a sequence of tasks $1 \dots T$ and a pre-trained model h_0 , there are at least two natural ways to learn the compo-

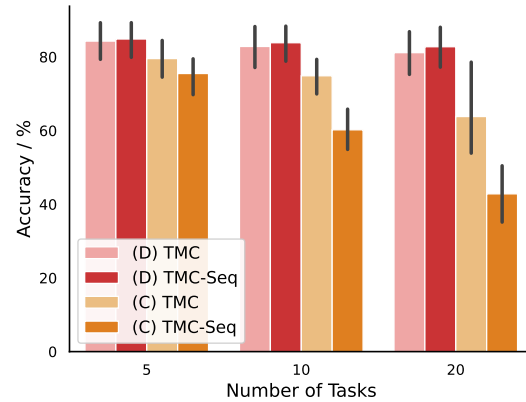


Figure 4. Accuracy of composed models, where component models are initialized with h_0 (TMC) and h_{t-1} (TMC-Seq). In the data-incremental setting (D), TMC-Seq outperforms TMC, while in the class-incremental setting (C), TMC performs significantly better. All results are averaged across D/C-Caltech-256, D/C-MIT-67, D/C-OxfordPets (we exclude C-OxfordPets in the 20-task setting since it does not contain enough classes).

nent model \tilde{h}_t : (1) Train \tilde{h}_t with h_{t-1} as initialization, and (2) Train \tilde{h}_t with h_0 as initialization. The first option (TMC-Seq) is most common among continual learning methods which impose regularization in the weights between the current and previous tasks or optimize using memory buffers consisting of examples obtained from previous tasks. As such, these methods are restricted to a sequential learning framework, and hence cannot be parallelized across tasks. In Sec. 3, we showed that, thanks to the convexity of the loss landscape, both options (1) and (2) are equivalent under our framework when we assume comparable generalization performance among global optima.

In practice, for massively over-parameterized networks, such assumptions do not hold. However, we show empirically in Fig. 4 that under the data-incremental setting, compared to TMC-Seq, TMC only incurs a small loss in generalization performance to enable massive parallelization across tasks. This can potentially provide orders of magnitude speed-ups compared to existing continual learning methods. Moreover, under the class-incremental setting, TMC greatly outperforms TMC-Seq, showing that the pre-trained initialization yields solutions of better generalization than h_{t-1} . We attribute this to the fact that the hypothesis space of models which are globally optimal for the training data is vast and yields varying generalization.

Thus, regardless of the loss landscape convexity, convergence to each point in this space depends on the “direction of approach”. Hence, initialization at h_{t-1} introduces a bias towards the previous tasks. Due to the high level of dissimilarity between tasks in the class-incremental setting, this bias harms the final model generalization performance.

5. Discussion

Our approach enables incremental learning under the assumption that increment tasks are “close” to the pre-trained model. We ensure proximity by training increment models independently on the tangent plane of the pre-trained model. The advantage is that pre-trained tasks are not forgotten, and increment tasks can be learned independently, in parallel, and easily forgotten if needed.

Of course, this approach does not address the problem of incremental learning in full generality, when increment tasks can be arbitrary and arbitrarily unrelated. Nonetheless, we have shown that our method is competitive with existing methods in some of the most challenging settings, for instance when the hypothesis spaces of component tasks are disjoint. On the positive side, despite not requiring a replay buffer, our method outperforms replay-based ones under the assumptions in which we operate.

A limitation of our approach is that it requires computing the linear span of pre-trained models, which can be challenging. However, for common deep network architectures such as the ResNet [21] or Transformer [52, 15] family, the Jacobian-Vector product can be efficiently evaluated [39, 35, 1, 31], reducing the computational complexity of training the tangent model to that of the same magnitude as a linear classifier on network activations. Hence, inference time for the linearized network is at most double that of the original model.

Another limitation of our approach is assuming a first-order expansion around a pre-trained objective can well approximate the fine-tuned model. Such assumptions hold poorly when the pre-training and downstream tasks are highly unrelated. Nevertheless, an ImageNet [13] pre-trained model already yields a sufficiently good local approximation for various real-world datasets [1, 31].

While we mainly explored linear compositionality in the context of continual learning, we note that this framework can be easily generalized. We discussed possible applications towards federated learning and forgetting, since our method naturally yields a parallel training framework that fully compartmentalizes information obtained each task.

6. Acknowledgements

We would like to thank Aditya Golatkar, Albert Zhao, and the anonymous reviewers for their feedback on the initial version of the paper. This work was supported by ARO

W911NF-17-1-0304 and ONR N00014-22-1-2252.

References

- [1] Alessandro Achille, Aditya Golatkar, Avinash Ravichandran, Marzia Polito, and Stefano Soatto. Lqf: Linear quadratic fine-tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15729–15739, 2021. 1, 2, 4, 9
- [2] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 139–154, 2018. 3, 6
- [3] Rahaf Aljundi, Punarjay Chakravarty, and Tinne Tuytelaars. Expert gate: Lifelong learning with a network of experts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3366–3375, 2017. 3
- [4] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. *Advances in neural information processing systems*, 32, 2019. 3
- [5] Elie Bienenstock and Stuart Geman. Compositionality in neural systems. In *The handbook of brain theory and neural networks*, pages 223–226. 1998. 1
- [6] Matteo Boschini, Lorenzo Bonicelli, Angelo Porrello, Giovanni Bellitto, Matteo Pennisi, Simone Palazzo, Concetto Spampinato, and Simone Calderara. Transfer without forgetting. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIII*, pages 692–709. Springer, 2022. 3, 5, 6, 7, 8
- [7] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems*, 33:15920–15930, 2020. 3, 5, 6, 7
- [8] Pietro Buzzega, Matteo Boschini, Angelo Porrello, and Simone Calderara. Rethinking experience replay: a bag of tricks for continual learning. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 2180–2187. IEEE, 2021. 3
- [9] Lucas Caccia, Rahaf Aljundi, Nader Asadi, Tinne Tuytelaars, Joelle Pineau, and Eugene Belilovsky. New insights on reducing abrupt representation change in online continual learning. *arXiv preprint arXiv:2203.03798*, 2022. 3, 5, 6, 7
- [10] Hyuntak Cha, Jaeho Lee, and Jinwoo Shin. Co2l: Contrastive continual learning. In *Proceedings of the IEEE/CVF International conference on computer vision*, pages 9516–9525, 2021. 3, 5, 6, 7
- [11] Arslan Chaudhry, Marc Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. *arXiv preprint arXiv:1812.00420*, 2018. 3
- [12] Leshem Choshen, Elad Venezian, Noam Slonim, and Yoav Katz. Fusing finetuned models for better pretraining. *arXiv preprint arXiv:2204.03044*, 2022. 2
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image

- database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 4, 9
- [14] Thomas G Dietterich. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pages 1–15. Springer, 2000. 2, 4
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 9
- [16] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel Roy, and Michael Carbin. Linear mode connectivity and the lottery ticket hypothesis. In *International Conference on Machine Learning*, pages 3259–3269. PMLR, 2020. 4
- [17] Timur Garipov, Pavel Izmailov, Dmitrii Podoprikin, Dmitry P Vetrov, and Andrew G Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. *Advances in neural information processing systems*, 31, 2018. 2
- [18] Aditya Golatkar, Alessandro Achille, Avinash Ravichandran, Marzia Polito, and Stefano Soatto. Mixed-privacy forgetting in deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 792–801, 2021. 2
- [19] Raphael Gontijo-Lopes, Yann Dauphin, and Ekin D Cubuk. No one representation to rule them all: Overlapping features of training methods. *arXiv preprint arXiv:2110.12899*, 2021. 2
- [20] Gregory Griffin, Alex Holub, and Pietro Perona. Caltech-256 object category dataset. 2007. 6
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 9
- [22] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 831–839, 2019. 3
- [23] Like Hui and Mikhail Belkin. Evaluation of neural architectures trained with square loss vs cross-entropy in classification tasks. *arXiv preprint arXiv:2006.07322*, 2020. 5
- [24] Gabriel Ilharco, Mitchell Wortsman, Samir Yitzhak Gadre, Shuran Song, Hannaneh Hajishirzi, Simon Kornblith, Ali Farhadi, and Ludwig Schmidt. Patching open-vocabulary models by interpolating weights. *arXiv preprint arXiv:2208.05592*, 2022. 2
- [25] Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Averaging weights leads to wider optima and better generalization. *arXiv preprint arXiv:1803.05407*, 2018. 2
- [26] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017. 3, 6
- [27] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015. 7
- [28] Janghyeon Lee, Hyeong Gwon Hong, Donggyu Joo, and Junmo Kim. Continual learning with extended kronecker-factored approximate curvature. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9001–9010, 2020. 3
- [29] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021. 1, 2
- [30] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. 3, 5, 6, 7
- [31] Tian Yu Liu, Aditya Golatkar, and Stefano Soatto. Tangent transformers for composition, privacy and removal. *arXiv preprint arXiv:2307.08122*, 2023. 9
- [32] Tian Yu Liu, Aditya Golatkar, Stefano Soatto, and Alessandro Achille. Integral continual learning along the tangent vector field of tasks. *arXiv preprint arXiv:2211.13108*, 2022. 2, 3
- [33] Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7765–7773, 2018. 3
- [34] Michael Matena and Colin Raffel. Merging models with fisher-weighted averaging. *arXiv preprint arXiv:2111.09832*, 2021. 2
- [35] Fangzhou Mu, Yingyu Liang, and Yin Li. Gradients as features for deep representation learning. *arXiv preprint arXiv:2004.05529*, 2020. 2, 4, 9
- [36] Cuong V Nguyen, Yingzhen Li, Thang D Bui, and Richard E Turner. Variational continual learning. *arXiv preprint arXiv:1710.10628*, 2017. 3
- [37] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019. 3
- [38] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3498–3505. IEEE, 2012. 6
- [39] Barak A Pearlmutter. Fast exact multiplication by the hessian. *Neural computation*, 6(1):147–160, 1994. 4, 9
- [40] Grégoire Petit, Adrian Popescu, Hugo Schindler, David Picard, and Bertrand Delezoide. Fetrl: Feature translation for exemplar-free class-incremental learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3911–3920, 2023. 3, 7, 8
- [41] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 524–540. Springer, 2020. 5
- [42] Ariadna Quattoni and Antonio Torralba. Recognizing indoor scenes. In *2009 IEEE conference on computer vision and pattern recognition*, pages 413–420. IEEE, 2009. 6

- [43] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017. [3](#)
- [44] Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. Learning to learn without forgetting by maximizing transfer and minimizing interference. *arXiv preprint arXiv:1810.11910*, 2018. [3](#)
- [45] Hippolyt Ritter, Aleksandar Botev, and David Barber. On-line structured laplace approximations for overcoming catastrophic forgetting. *Advances in Neural Information Processing Systems*, 31, 2018. [3](#), [6](#)
- [46] Anthony Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 7(2):123–146, 1995. [5](#), [6](#), [7](#)
- [47] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016. [3](#)
- [48] Jonathan Schwarz, Wojciech Czarnecki, Jelena Luketina, Agnieszka Grabska-Barwinska, Yee Whye Teh, Razvan Pascanu, and Raia Hadsell. Progress & compress: A scalable framework for continual learning. In *International conference on machine learning*, pages 4528–4537. PMLR, 2018. [5](#), [7](#)
- [49] Hyounguk Shon, Janghyeon Lee, Seung Hwan Kim, and Junmo Kim. Dlcft: Deep linear continual fine-tuning for general incremental learning. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIII*, pages 513–529. Springer, 2022. [2](#), [3](#), [6](#), [7](#)
- [50] Doris Y Tsao, Winrich A Freiwald, Tamara A Knutsen, Joseph B Mandeville, and Roger BH Tootell. Faces and objects in macaque cerebral cortex. *Nature neuroscience*, 6(9):989–995, 2003. [1](#)
- [51] Gido M Van de Ven and Andreas S Tolias. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*, 2019. [3](#), [6](#), [7](#)
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [1](#), [9](#)
- [53] Johannes Von Oswald, Seijin Kobayashi, Joao Sacramento, Alexander Meulemans, Christian Henning, and Benjamin F Grewe. Neural networks with late-phase weights. *arXiv preprint arXiv:2007.12927*, 2020. [2](#)
- [54] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In *International Conference on Machine Learning*, pages 23318–23340. PMLR, 2022. [2](#)
- [55] 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. *arXiv preprint arXiv:2204.04799*, 2022. [1](#), [2](#)
- [56] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 139–149, 2022. [2](#), [3](#)
- [57] Mitchell Wortsman, Suchin Gururangan, Shen Li, Ali Farhadi, Ludwig Schmidt, Michael Rabbat, and Ari S Morcos. lo-fi: distributed fine-tuning without communication. *arXiv preprint arXiv:2210.11948*, 2022. [2](#)
- [58] Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, et al. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International Conference on Machine Learning*, pages 23965–23998. PMLR, 2022. [2](#), [5](#), [6](#)
- [59] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, et al. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7959–7971, 2022. [2](#)
- [60] Fei Zhu, Zhen Cheng, Xu-Yao Zhang, and Cheng-lin Liu. Class-incremental learning via dual augmentation. *Advances in Neural Information Processing Systems*, 34:14306–14318, 2021. [3](#), [8](#)
- [61] Fei Zhu, Xu-Yao Zhang, Chuang Wang, Fei Yin, and Cheng-Lin Liu. Prototype augmentation and self-supervision for incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5871–5880, 2021. [3](#), [8](#)
- [62] Kai Zhu, Wei Zhai, Yang Cao, Jiebo Luo, and Zheng-Jun Zha. Self-sustaining representation expansion for non-exemplar class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9296–9305, 2022. [3](#), [7](#), [8](#)