

---

# Wide Neural Networks Forget Less Catastrophically

---

Seyed Iman Mirzadeh<sup>\*1</sup> Arslan Chaudhry<sup>2†</sup> Dong Yin<sup>2</sup>  
Huiyi Hu<sup>2</sup> Razvan Pascanu<sup>2</sup> Dilan Gorur<sup>2</sup> Mehrdad Farajtabar<sup>2†</sup>

## Abstract

A primary focus area in continual learning research is alleviating the “catastrophic forgetting” problem in neural networks by designing new algorithms that are more robust to the distribution shifts. While the recent progress in continual learning literature is encouraging, our understanding of what properties of neural networks contribute to catastrophic forgetting is still limited. To address this, instead of focusing on continual learning algorithms, in this work, we focus on the model itself and study the impact of “width” of the neural network architecture on catastrophic forgetting, and show that width has a surprisingly significant effect on forgetting. To explain this effect, we study the learning dynamics of the network from various perspectives such as gradient orthogonality, sparsity, and lazy training regime. We provide potential explanations that are consistent with the empirical results across different architectures and continual learning benchmarks.

## 1. Introduction

Machine learning is relying more and more on training large models on large static datasets to reach impressive results (Kaplan et al., 2020; Lazaridou et al., 2021; Hombaiyah et al., 2021). However, the real world is changing over time, and new information is becoming available at an unprecedented rate (Lazaridou et al., 2021; Hombaiyah et al., 2021). In such real-world problems, the learning agent is exposed to a continuous stream of data, with potentially changing data distribution, and it has to absorb new information efficiently while not being able to iterate on previous data as freely as wanted due to time, sample, compute, privacy, or environmental complexity issues (Parisi


<sup>\*</sup>Work done during an internship at DeepMind. <sup>†</sup>Equal advising. <sup>1</sup>Washington State University <sup>2</sup>DeepMind. Correspondence to: Seyed Iman Mirzadeh <seyediman.mirzadeh@wsu.edu>, Mehrdad Farajtabar <farajtabar@google.com>.

et al., 2019). To overcome these inefficiencies, fields, such as *continual learning* (CL) (Ring, 1995) or lifelong learning (Thrun, 1995) are gaining a lot of attention recently. One of the key challenges in continual learning models is the abrupt erasure of previous knowledge, referred to as *catastrophic forgetting* (CF) (McCloskey and Cohen, 1989).

Alleviating catastrophic forgetting has attracted a lot of attention lately, and many interesting solutions are proposed to partly overcome the issue (e.g., Toneva et al., 2019; Nguyen et al., 2019; Hsu et al., 2018; Li et al., 2019; Wallingford et al., 2020). These solutions vary in degree of complexity from simple replay-based methods to complicated regularization or network expansion-based methods. Unfortunately, however, there is not much fundamental understanding of the intrinsic properties of neural networks that affect continual learning performance through catastrophic forgetting or forward/backward transfer (Mirzadeh et al., 2020b). Moreover, most of the continual learning works focus on the algorithmic side of alleviating the forgetting rather than the architecture being used throughout the learning experience. This motivated our work to study one of the most important aspects of neural network architectures: the structure of the neural networks (e.g., width and depth).

In contrast to the orthodox importance of the depth for network performance, such as in AlexNet (Krizhevsky et al., 2012), Inception (Szegedy et al., 2016), ResNets (He et al., 2016) or GPT family of models (Brown et al., 2020), we study how the width of the network affects the continual learning performance – aligned with the recent research on infinite width networks and the neural tangent kernel (NTK) regime (Jacot et al., 2018; Lee et al., 2019; Arora et al., 2019; Chizat et al., 2019). We empirically demonstrate that increasing the width alone reduces catastrophic forgetting significantly, while it’s not the case for depth.

Concretely, Figures 1 and 2 depict our findings on two popular continual learning benchmarks. Figure 1 shows that by increasing the width of a 2-layer MLP from 32 to 2048 the amount of forgetting of task 1 (after learning the 5th task) decreases from 62 to 48 percent. On Split CIFAR-100, by multiplying the width of a WideResNet-10 by 8 we can decrease forgetting of the first task (after learning the 20th task) from 42 to 31 percent.



(a) MLP on Rotated MNIST (b) WRN on Split CIFAR-100

Figure 1. Amount of forgetting w.r.t. increasing width. Wider networks forget less.


The boost in performance gained by only increasing the width is comparable to employing alternative methods specially developed to tackle CF with moderately sized replay buffers or regularization schemes. For example, it can be seen that in Figure 1a, wider networks without replay achieve comparable or better performance than the network with width 32 and employing replay buffers of sizes of 125 and 250. A similar level of performance improvement can be observed in Figure 1b, where the baseline WideResNet (width factor equal to 1) is trained with replay buffer sizes of 100 and 400. The ability of the wider networks to be less forgetful points towards looking at forgetting through the lens of over-parametrization (Kaplan et al., 2020; Lazaridou et al., 2021; Hombaiah et al., 2021).

It is noteworthy that increasing width is not aimed to achieve state-of-the-art performance in alleviating forgetting. Our goal is not to compete with computationally cheaper and performant rehearsal or regularization-based methods (Chaudhry et al., 2018; Farajtabar et al., 2020; Mirzadeh et al., 2021) but to provide a perspective on how over-parametrization and especially the width of the neural network can affect CF. One may wonder that a similar over-parametrization achieved through depth might have a similar positive effect on continual learning performance. Figure 2 shows that this is not the case and that achieving over-parametrization through depth has no or even negative effect on mitigating forgetting.

**Contributions.** We study the implications of different neural network structures and over-parametrization (i.e., width and depth) for continual learning. We show that increased width can alleviate the forgetting and provide simplistic theoretical and empirical justifications such as lazy training regime and orthogonalization that can explain the benefits of using wide neural networks in continual learning.

## 2. Related Work

In continual learning (Ring, 1995), also referred to as lifelong learning (Thrun, 1995), an agent, constrained in time and memory, receives a sequence of tasks. The goal of the agent is to remember the useful knowledge of the past tasks to solve new tasks efficiently. Catastrophic forget-



(a) MLP on Rotated MNIST (b) WRN on Split CIFAR-100

Figure 2. Amount of forgetting w.r.t. increasing depth.

ting (McCloskey and Cohen, 1989; McClelland et al., 1995; Goodfellow et al., 2013)—a sudden erasure of previous knowledge when exposed to new tasks—is identified as one of the key challenges in continual learning. Recently, a lot of algorithmic progress has been made in mitigating catastrophic forgetting in neural networks. The progress can broadly be classified into three categories.

**Regularization-based** methods (Rebuffi et al., 2016; Kirkpatrick et al., 2017; Zenke et al., 2017; Nguyen et al., 2018; Aljundi et al., 2018; Yin et al., 2020) attempt to identify important parameters or features of the past tasks, and modify the training objective of the new task such that those parameters or features do not drift significantly. While regularization-based methods provide strong results for a small number of tasks, the importance of parameters drifts when the number of tasks is large (Chaudhry et al., 2018; Titsias et al., 2019).

**Expansion-based** methods (Rusu et al., 2016; Fernando et al., 2017; Aljundi et al., 2017; Rosenbaum et al., 2018; Chang et al., 2018; Xu and Zhu, 2018; Ferran Alet, 2018; Li et al., 2019; Veniat et al., 2021) add new network modules or experts for new tasks. By construction, expansion-based methods can have zero-forgetting, but their memory complexity can increase super-linearly with the number of tasks, making them undesirable when the pool of tasks is large. We note that even though the expansion-based works may have the model as their main subject, the main objective of these works is designing algorithms that can efficiently use a model (e.g., if the model should be expanded or not), and they do not explicitly analyze the structure of the model (e.g., width). For instance, while Yoon et al. (2018) show that expanding the network capacity thorough width is helpful, they have not studied the impact of increasing capacity when the depth increases, nor why increasing the width is helpful. Overall, in this work, we are interested in understanding the impacts of network structure (e.g., width and depth) on continual learning performance by focusing on the model itself rather than the algorithm that works with the model.

Perhaps the strongest are the **replay-based** methods that store a small subset of data from previous tasks in the replay buffer, and either use it directly in the form of experience replay with new task (Shin et al., 2017; Kirichenko et al.,

2021; Rolnick et al., 2019; Riemer et al., 2018; Chaudhry et al., 2019; Chaudhry et al., 2020; Balaji et al., 2020) or use it as a set of optimization constraints (Lopez-Paz and Ranzato, 2017; Chaudhry et al., 2018; Farajtabar et al., 2020).

While these methods provide algorithmic improvements on top of standard neural network training (empirical risk minimization), there is still limited understanding as to what causes forgetting in the first place. Some works have looked at the forgetting from the lens of network representations. For example, Javed and White (2019); Beaulieu et al. (2020) showed that meta-learned representations forget less when trained continually. On the other hand, Mehta et al. (2021) showed that pre-trained over-parameterized networks forget considerably less than randomly initialized networks when trained continually. Similarly, few works have studied the forgetting from the neural network’s training dynamics perspective. Mirzadeh et al. (2021) showed that the optima of the multitask learning and continual learning solutions are linearly connected in the parameter space. Contrary to these works, we look at catastrophic forgetting from the lens of network architecture, specifically with respect to its width and depth and show that wider networks forget considerably less.

A flurry of recent papers in deep learning theory has taken a huge step in understanding the training behavior of neural networks. In particular, Jacot et al. (2018) first introduced the concept of neural tangent kernel (NTK) by studying the training dynamics of neural networks with square loss in the infinite width or “over-parameterized” regime. The resulting neural network can be approximated by a linear model. Based on the NTK framework, other works (Allen-Zhu et al., 2019; Du et al., 2019; Zou et al., 2020) have proved the convergence of the SGD on various other settings and proved the NTK to be a powerful framework to study neural networks training dynamics. Later Chizat et al. (2019) developed a parallel idea under the framework called “lazy training” where the weights stay near their initial values during the training. In this regime, the system becomes almost linear, and the dynamics of SGD within this region can be characterized via properties of the associated NTK. In the context of continual learning, Doan et al. (2021) have studied forgetting of projection-based algorithms in the NTK regime using the similarity of the tasks.

Finally, several recent works have explored the intersection of neural network architectures and continual learning. Mirzadeh et al. (2022) provides a comprehensive analysis on the significant role of architectures and architectural decisions in continual learning. They show that that different architecture families have different learning and retention capabilities, and compare the merits and flaws of various architectures. For the pre-training setup, Ramasesh et al. (2022) studies the effect of scaling in continual learning and

show that large pre-trained models perform significantly better than the models trained from scratch in continual learning scenarios.

### 3. Main Experiments

In this section, we study the impact of width on both average accuracy and average forgetting. Further detailed experiments and results will be provided in Sections 4 and 5.

#### 3.1. Experimental Setup

The experimental setup, such as benchmarks, network architectures, continual learning setting (e.g., number of tasks, episodic memory size, and training epochs per task), hyper-parameters, and evaluation metrics are chosen to be similar to several other studies (Chaudhry et al., 2019; Farajtabar et al., 2020; Mirzadeh et al., 2021). To ensure that the reported results are not more favorable to a specific architecture, we use a grid of hyper-parameters, and for each model, we report the results with the best hyperparameters. For all experiments, we report the average and standard deviation over five runs with different random seeds for network initialization. In addition to random seed for initialization, for Split CIFAR-100, we use three different seeds for each run where each seed corresponds to the random order in which classes are selected for each task. We provide more details about our experimental design in Appendix A.

**Benchmarks.** We report our findings on two standard continual learning benchmarks: Rotated MNIST and Split CIFAR-100. In Rotated MNIST, each task is generated by the continual rotation of the MNIST images for degrees 0, 22.5, 45, 67.5, and 90, respectively, constituting 5 different tasks. For Split CIFAR-100 each task contains the data from 5 random classes (without replacement), resulting in 20 tasks. Finally, in both benchmarks, each model trains on each task for five epochs. We have selected Rotated MNIST as a domain-incremental benchmark and Split CIFAR-100 as a task-incremental benchmark to represent two common scenarios in continual learning.

**Networks.** For the Rotated MNIST benchmark, we use a two-layer MLP with ReLU non-linearities in each layer and vary the width. For the Split CIFAR-100 benchmark, we use a WideResNet (WRN) with the depth of 10 (Zagoruyko and Komodakis, 2016) and different width factors which multiply the width of the layers (i.e., the number of convolutional channels) in the original WideResNet by  $\{1, 2, 4, 8\}$ . All models are trained using the SGD optimizer, and the best result for each model is reported.

**Evaluations.** Following previous work (Chaudhry et al., 2018; Riemer et al., 2018; Mirzadeh et al., 2021), we report the following metrics:

(1) *Average Accuracy*: The average validation accuracy

Table 1. MLP on Rotated MNIST

Width	Average Accuracy	Average Forgetting	Learning Accuracy	Joint Accuracy
32	65.9 $\pm$ 1.00	36.9 $\pm$ 1.27	95.5 $\pm$ 0.85	91.2 $\pm$ 0.57
128	70.8 $\pm$ 0.68	31.5 $\pm$ 0.92	96.0 $\pm$ 0.90	93.4 $\pm$ 0.66
512	72.6 $\pm$ 0.27	29.6 $\pm$ 0.36	96.4 $\pm$ 0.73	94.1 $\pm$ 0.77
2048	75.2 $\pm$ 0.34	26.7 $\pm$ 0.50	96.6 $\pm$ 0.61	94.0 $\pm$ 0.45

after the model has been continually trained for  $T$  tasks, is defined as:

$$A_T = \frac{1}{T} \sum_{i=1}^T a_{T,i} \quad (3.1)$$

where,  $a_{t,i}$  is the validation accuracy on the dataset of task  $i$  after the model finished learning task  $t$ .

(2) *Average Forgetting*: The average forgetting is calculated as the difference between the peak accuracy and the final accuracy of each task, after the continual learning experience is finished. For a continual learning benchmark with  $T$  tasks, it is defined as:

$$F = \frac{1}{T-1} \sum_{i=1}^{T-1} \max_{t \in \{1, \dots, T-1\}} (a_{t,i} - a_{T,i}) \quad (3.2)$$

(3) *Learning Accuracy*: The accuracy for each task directly after it is learned. The learning accuracy provides a good representation of the plasticity of a model and can be calculated using:

$$LA_T = \frac{1}{T} \sum_{i=1}^T a_{i,i} \quad (3.3)$$

(4) *Joint Accuracy*: The accuracy of the model trained on the combined data of all tasks together.

### 3.2. Main Results

Tables 1 and 2 show the main results in terms of average accuracy, average forgetting, and learning accuracy for MLP on Rotated MNIST and WideResNet on Split CIFAR-100 with varying widths. It can be seen from the tables that wider networks improve both average accuracy and forgetting.

It is important to note that reduced forgetting is not due to stabilizing the model dogmatically at the cost of plasticity. Increased average accuracy means that the wider models not only remember previous tasks better, but keep their competitive performance on the current task too. This can be further consolidated by observing the learning accuracy (LA) of the wider networks.

Table 2. WideResNet-10 on Split CIFAR-100

Width	Average Accuracy	Average Forgetting	Learning Accuracy	Joint Accuracy
1	47.1 $\pm$ 2.60	37.3 $\pm$ 2.62	81.1 $\pm$ 4.37	81.4 $\pm$ 0.62
2	49.7 $\pm$ 1.51	34.9 $\pm$ 1.72	82.9 $\pm$ 5.17	83.6 $\pm$ 0.13
4	53.8 $\pm$ 2.74	33.8 $\pm$ 2.16	83.9 $\pm$ 6.53	84.4 $\pm$ 0.64
8	59.7 $\pm$ 2.33	29.4 $\pm$ 2.52	87.5 $\pm$ 4.15	84.8 $\pm$ 0.49

## 4. Analysis

In this section, we overview the reasons we believe contribute to the reduction in catastrophic forgetting when widening the network. Although wide networks have higher capacity, similar to how it has been shown that wider networks generalize well because of implicit regularization imposed by gradient descent (Jacot et al., 2018; Lee et al., 2019), we argue that the behavior of catastrophic forgetting cannot be reduced to capacity alone and we need to integrate learning dynamics. Here, we first construct a simple example for a regression problem with squared loss to demonstrate this point. Informally, we have the following claim.

**Claim 4.1** (informal). *Consider learning problems with input space  $\mathbb{R}^d$ , output space  $\mathbb{R}$ , and squared loss. Let  $\mathcal{F}_1$  be the class of linear models that maps the input to the output and  $\mathcal{F}_2$  be the class of two-layer **linear** networks (i.e., no nonlinear activation). Then, there exist two tasks such that when we train task 2 using gradient descent, if we use model class  $\mathcal{F}_1$ , the amount of forgetting for task 1 is strictly zero; whereas if we use model class  $\mathcal{F}_2$ , the amount of forgetting can be positive.*

We provide details of the construction of the two tasks in Appendix C. As we can see, the two model classes  $\mathcal{F}_1$  and  $\mathcal{F}_2$  have exactly the same capacity, i.e.,  $\mathcal{F}_1 = \mathcal{F}_2$ , since there is no nonlinearity in  $\mathcal{F}_2$ . However, due to the use of gradient descent algorithm, the amount of forgetting for task 1 is different for the two model classes when training on task 2. Moreover, in Appendix C, we show that if the size of a hidden layer in  $\mathcal{F}_2$  is not large enough, then the forgetting could be more severe. Intuitively, if the additional layer becomes a bottleneck that projects the input to a very low dimensional space, the information can be lost.

While this example is a specially constructed one, we believe it can be helpful towards understanding the benefits of wide neural networks in continual learning. In particular, it shows that we have to take the commonly used *gradient descent algorithm* into account when investigating this phenomenon. In Appendix B, we provide further empirical results regarding the significance of the width of the first layer. In the following, we dive deeper in this phenomenon and provide several potential explanations.

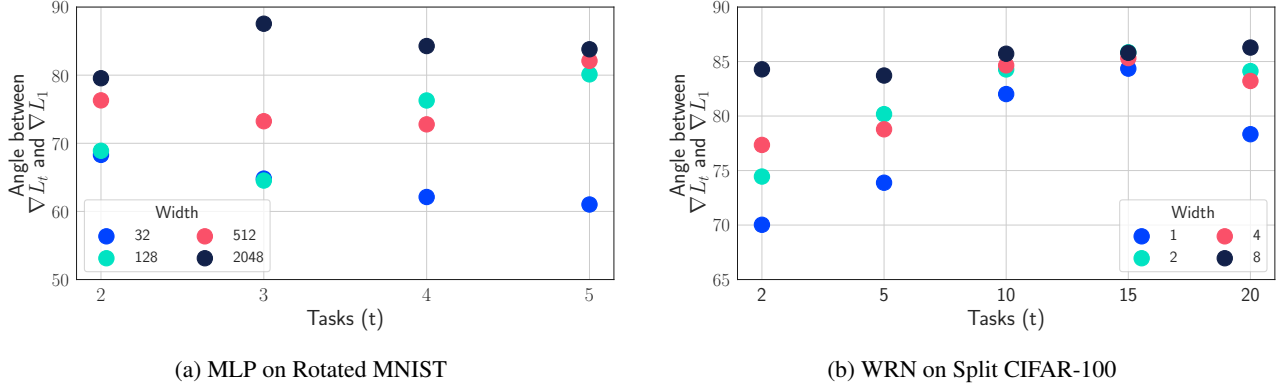


Figure 3. The angle between the gradients at the optimum of task 1,  $\nabla L_1(w_1^*)$ , and between the optimum of subsequent tasks  $t$ ,  $\nabla L_t(w_t^*)$ , for  $t > 1$ . Wider networks remain closer to 90 degree (more orthogonal).

#### 4.1. Increased Gradient Orthogonality

Earlier works in continual learning (Farajtabar et al., 2020; Chaudhry et al., 2018; Lopez-Paz and Ranzato, 2017; Riemer et al., 2018) have focused on explicitly forcing the gradients of the new tasks to be orthogonal to those of the previous ones to reduce forgetting among the tasks. Here, we want to look at how the gradient directions of different tasks evolve as the continual learning experience progresses while the networks’ width increases. We first establish why orthogonalizing the gradients of different tasks is a good continual learning strategy.

Let us consider a two-task continual learning problem, and study the amount of forgetting for task 1 induced by one gradient descent step on task 2. Let  $\mathcal{W} \subseteq \mathbb{R}^p$  be the model parameter space and  $L_t(w) : \mathcal{W} \mapsto \mathbb{R}$  be the training loss function of the  $t$ -th task. Assuming that all the loss functions are differentiable, according to mean value theorem, we have the following simple observation.

**Claim 4.2.** Let  $w' = w - \eta \nabla L_2(w)$ . Then there exists  $\xi \in [0, 1]$  such that

$$L_1(w') - L_1(w) = -\eta \langle \nabla L_1(w - \xi \eta \nabla L_2(w)), \nabla L_2(w) \rangle. \quad (4.1)$$

As we can see, the left hand side of Eq. 4.1 is the increase in the loss of task 1, and the right hand side is related to the inner product of the gradient of  $L_1$  at  $w - \xi \eta \nabla L_2(w)$ , and the gradient of  $L_2$  at  $w$ . Therefore, if the gradients of the two tasks are more orthogonal to each other, then when training on new tasks, the model forgets less about previous task. Formally, we have the following direct corollary.

**Claim 4.3.** Suppose that for any  $w_1, w_2 \in \mathcal{W}$  with  $\|w_1 - w_2\|_2 \leq \eta \sup_{w \in \mathcal{W}} \|\nabla L_2(w)\|_2$ , we have

$$|\langle \nabla L_1(w_1), \nabla L_2(w_2) \rangle| \leq \epsilon / \eta,$$

for some  $\epsilon > 0$ , and let  $w' = w - \eta \nabla L_2(w)$ . Then, we have  $L_1(w') - L_1(w) \leq \epsilon$ .

Hence, the one-step forgetting is small if we have near-orthogonal gradients for the two tasks. The next question is whether widening the network increases the gradient orthogonality. Empirically, to study the directions of the gradients for different tasks, in Figure 3, we plot the angle between the gradient vector at the minimum of the first task and that of the subsequent tasks for various width factors. For this figure, we concatenate the gradients from all the layers into one vector. It can be seen from the figure that the gradients at the minimum of subsequent tasks become more orthogonal to the gradients at the minimum of the first task as we increase the width. This trend also holds when the gradients of other tasks are used as reference (*i.e.*) the gradients at the minimum of tasks  $> 2$  also become orthogonal to the gradients at the minimum of task 2. Refer to the appendix for these results. Since the gradients of different tasks generally become more orthogonal for the wider networks, the reduction in forgetting can be attributed to this orthogonalization.

#### 4.2. Increased Gradient Sparsity

Another scenario where forgetting may be reduced is that during the training of a new task, we only need to update a small number of model parameters, and thus the network is stable and forgets less about the previous tasks. In fact, according to Claim 4.2, we can establish the following simple corollary.

**Claim 4.4.** Suppose that  $\|\nabla L_2(w)\|_2 \leq B$  for all  $w \in \mathcal{W}$ , and let  $b(w) := \sup_{\|\tilde{w} - w\|_2 \leq \eta B} \|\nabla L_1(\tilde{w})\|_\infty$ . Let  $w' = w - \eta \nabla L_2(w)$ . Then we have

$$L_1(w') - L_1(w) \leq \eta b(w) \|\nabla L_2(w)\|_0,$$

where  $\|\cdot\|_0$  denotes the number of non-zero elements in a vector.

Therefore, increased gradient sparsity may lead to less forgetting. Figure 4 confirms this hypothesis by demonstrating

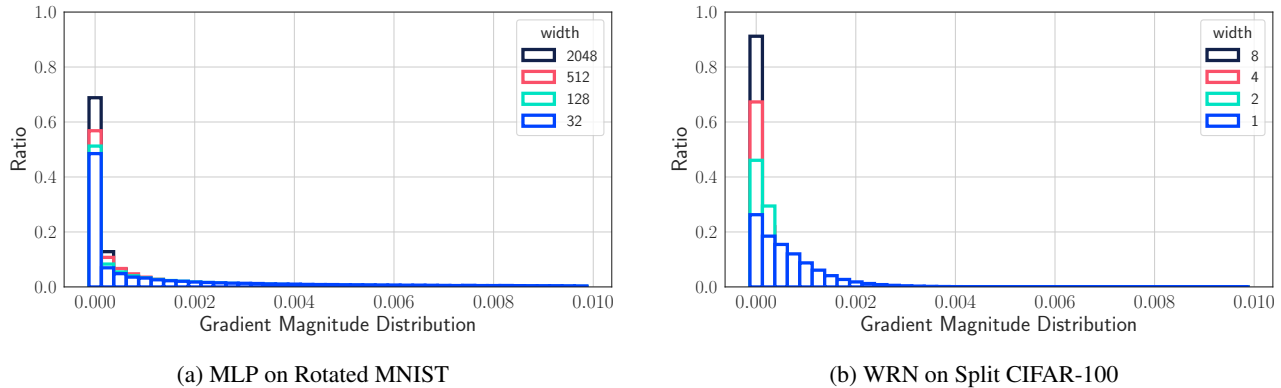


Figure 4. Histogram of the magnitude of the elements of the gradient vector. Wider networks have sparser gradients (the histogram is more skewed towards the left).

the histogram of the absolute value of the entries of the gradient. The gradient norm is constructed by calculating the gradients of all layers on the empirical loss of task 2, after learning task 1 and before starting task 2. From the figure, it can be seen that the wider networks have sparser gradients. This implies a smaller number of parameters needs to change to adapt to new tasks resulting in a reduced forgetting on already learned tasks.

### 4.3. Lazy Training Regime

In the previous two sections, we showed that the amount of forgetting induced by a single gradient descent step on a new task can be reduced by increased gradient orthogonality/sparsity. In this section, we take a different perspective and study the entire training process of a task. In the following, we denote by  $w_t^{\text{in}}$  and  $w_t^*$  the initial and end model parameters for task  $t$ , respectively. Note that in the continual learning scenario, the initializer for a subsequent task is in fact the solution to the previous task. Thus,  $w_t^{\text{in}} = w_{t-1}^*$ .

Let us consider the amount of forgetting of task  $t - 1$  at the end of the training of task  $t$ , measured by the increase in loss function  $L_{t-1}$ , i.e.,  $L_{t-1}(w_t^*) - L_{t-1}(w_{t-1}^*)$ . Suppose that the loss function is differentiable, by mean value theorem, there exists  $\xi \in [0, 1]$  such that

$$\begin{aligned} & L_{t-1}(w_t^*) - L_{t-1}(w_{t-1}^*) \\ &= \nabla L_{t-1}((1 - \xi)w_{t-1}^* + \xi w_t^*)^\top (w_t^* - w_{t-1}^*), \end{aligned}$$

which yields the following upper bound on the amount of forgetting:

$$\begin{aligned} & L_{t-1}(w_t^*) - L_{t-1}(w_{t-1}^*) \\ & \leq D_t \sup_{\{w: \|w - w_{t-1}^*\|_2 \leq D_t\}} \|\nabla L_{t-1}(w)\|_2, \end{aligned} \quad (4.2)$$

where  $D_t := \|w_t^* - w_{t-1}^*\|_2$  is the distance the the model parameter moves during the training of task  $t$ . This result


suggests the following intuitive claim: The forgetting may be less severe if the model moves a shorter distance  $D_t$  in the parameter space.

Interestingly, for overparametrized neural networks, the distance that the model parameter moves during SGD training can indeed be small. In their seminal papers, Jacot et al. (2018) and Lee et al. (2019) connect learning in neural networks to kernel methods and show that at the infinite width limit, the kernel doesn't change throughout the training, and the neural network evolves as a linear model under SGD training. Moreover, Chizat et al. (2019) show that these neural networks operate in lazy regime where  $\|w^* - w^{\text{in}}\|_2$  is very small. Since wider networks are closer to the infinite width limit, the moving distance  $D_t$  may be smaller for wider networks and thus the forgetting is less severe.


To confirm this intuition, in Figure 5, we plot the distance between the optimum of the first task and that of the subsequent tasks as the continual learning experience progress. In the figure, each task is trained for 5 epochs. From the figure, it can be seen that, for both MLPs and ResNets, the relative distance of the parameters for the wider networks is substantially smaller than that of the less wide networks. This confirms that lazy training of wider networks is indeed beneficial for continual learning and reduces forgetting significantly. In the same vein, Mirzadeh et al. (2020b;a) observed that changing the learning regime (i.e., imposing a learning rate decay, decreasing batch size, employing dropout, etc.) also leads to a small change in parameters from the optimum of previous tasks resulting in less forgetting. However, their work was concerned about the optimization properties of neural networks rather than architectural ones.

### 4.4. The Case for Depth

In Section 3, we saw that the width of the network has a beneficial effect when it comes to catastrophic forgetting. We do not see a similar monotonic reduction in forgetting



(a) MLP on Rotated MNIST



(b) WRN on Split CIFAR-100

Figure 5. The distance between  $w_1^*$ , solution of task 1, and  $w_t^*$ , solution of the  $t$ -th task, for  $t > 1$  over the course of training. Wider networks remain closer to the solution of the first task.

when we increase the depth of the network in our experiments. In fact, we observe that with increasing depth, the forgetting increases, as shown in Figure 2. We now provide an intuitive analysis for why and when the network depth hurts the performance of a continual learning model.

Inspired by the exploding gradients phenomena in recurrent neural networks (Bengio et al., 1994; Pascanu et al., 2013), we show that deeper networks may cause larger gradient updates in the earlier layers. Larger gradients mean that the solution for the subsequent tasks will not be found in the vicinity of the current task, which would entail increased forgetting. To establish the larger gradients for deeper networks, we use the standard argument (Bengio et al., 1994) for exploding gradients. Suppose that the network consists of  $K$  feed-forward layers parameterized by weight matrices  $\{W_i\}_{i=1}^K$  and non-linearities  $\sigma(\cdot)$ . For ease of exposition, we ignore the biases and normalization layers in the network. The gradient on any intermediate layer  $l$  can be written as:

$$\begin{aligned} \frac{\partial L}{\partial h_l} &= \frac{\partial L}{\partial h_K} \cdot \frac{\partial h_K}{\partial h_l} = \frac{\partial L}{\partial h_K} \cdot \prod_{i=l}^{K-1} \frac{\partial h_{i+1}}{\partial h_i} \\ &= \frac{\partial L}{\partial h_K} \cdot \prod_{i=l}^{K-1} \Sigma_{i+1} W_{i+1}, \end{aligned}$$

where  $\Sigma_i$  is the diagonal Jacobian of the non-linearity at the  $i$ -th layer. Note, if a ReLU non-linearity is used, then  $\|\Sigma_i\| = 1$ , unless all the activations in the layer are negative.

The norm of the gradient on layer  $l$  can then be bounded as:


$$\begin{aligned} \left\| \frac{\partial L}{\partial h_l} \right\| &= \left\| \frac{\partial L}{\partial h_K} \cdot \prod_{i=l}^{K-1} \Sigma_{i+1} W_{i+1} \right\| \\ &\leq \left\| \frac{\partial L}{\partial h_K} \right\| \cdot \prod_{i=l}^{K-1} \|\Sigma_{i+1}\| \|W_{i+1}\| \\ &\leq \left\| \frac{\partial L}{\partial h_K} \right\| \cdot \prod_{i=l}^{K-1} \|W_{i+1}\|, \end{aligned}$$

where in the second line, we use sub-ordinate and sub-multiplicative properties of induced matrix norms, and in the third line, we assumed that ReLU activation is used and thus  $\|\Sigma_i\| \leq 1$ . This analysis shows that the gradient norm on the earlier layers is proportional to the product of matrix norms of higher layer weights. For a 2-norm, the gradient magnitude could be bounded by the product of largest singular values ( $\Lambda_i$ ) of the intermediate weight matrices:


$$\left\| \frac{\partial L}{\partial h_l} \right\| \leq \left\| \frac{\partial L}{\partial h_K} \right\| \cdot \prod_{i=l}^{K-1} \|\Lambda_{i+1}\|. \quad (4.3)$$

Although, for  $\Lambda_i > 1$  this bound is unusable, as it is an upper bound which grows exponentially, and becomes more and more loose, with  $K$  (the number of layers) increasing, however, it does potentially point to a scenario when the gradients on the earlier layers could increase with the depth. This *can* happen when the singular values of all the intermediate weight matrices are greater than 1, *i.e.*,  $\Lambda_i > 1, \forall i \in \{l, \dots, K-1\}$ . In the supplementary Figure 11, we show that this indeed is the case for the datasets and architectures considered in this work. A large gradient on earlier layers can then contribute to a larger forgetting of previous tasks.

Figure 6a shows the norm of the gradients of layer 1 for different MLP depths. It can be seen from the figure that the gradient norm on the earlier layers increases with the depth. For example, when the network is trained for task 5, the gradient norm on layer 1 for an 8-layer network is almost three times as that of the 2-layer network. In contrast, as depicted in Figure 6b, increasing the width has a minimal or even decreasing effect on gradient norm. We note here that the deep learning community has developed numerous strategies to avoid exploding gradients (Goodfellow et al., 2016), and that extensively studying those is not the purpose here. We use the exploding gradient analysis to understand the negative effect of depths in our experiments.



(a) Gradient norm for different depths



(b) Gradient norm for different widths

Figure 6. MLP on Rotated MNIST: The gradient norm in the first layer (i.e.,  $\left\|\frac{\partial L_1}{\partial h_1}\right\|$ ) grows faster in deeper networks (a) while this is not the case when the width increases (b).

Table 3. Comparing shallow and wide vs. deep and thin MLP models on Rotated MNIST. Each row in a subdivision of the table has roughly a comparable number of parameters. Shallower and wider networks have higher average accuracy and smaller forgetting than deeper and thinner networks.

Width	Depth	Parameters	Average Accuracy	Average Forgetting	Joint Accuracy
128	8	217.35 K	68.9 ± 1.07	35.4 ± 1.34	94.1 ± 0.73
256	2	269.32 K	<b>71.1 ± 0.43</b>	<b>31.4 ± 0.48</b>	93.9 ± 0.65
256	8	664.08 K	70.4 ± 0.61	32.1 ± 0.75	94.78 ± 0.67
512	2	669.70 K	<b>72.6 ± 0.27</b>	<b>29.6 ± 0.36</b>	94.08 ± 0.77

In conclusion, we have observed that wider networks have sparser gradients with more orthogonal gradients across tasks. In addition, the training dynamics of wide networks become more similar to the lazy training regime. Finally, the gradient norm in wide and shallow models does not increase as fast as in deeper and thinner models.

## 5. Additional Experiments

This section provides further empirical results on the role of width and depth in continual learning. More specifically, fixing the number of parameters, the wider and shallower models, outperform deeper and thinner ones. Moreover, we show that the benefits of increasing the width can be complementary to the benefits of algorithms.

### 5.1. Width versus Depth

We have already seen in Figure 1 and 2 that increasing width is beneficial while increasing depth has no such effect. However, it is interesting to answer the following question in a more direct manner: *Given a fixed budget on the total number of parameters, how would a wider and shallower model perform compared to a thinner and deeper one?*

Table 4. Comparing shallow and wide vs. deep and thin WRN models on Split CIFAR-100. Each row in a subdivision of the table has roughly a comparable number of parameters. Shallower and wider networks have higher average accuracy and smaller forgetting than deeper and thinner networks.

Depth	Width	Params	Average Accuracy	Average Forgetting	Joint Accuracy
10	4	1.69 M	<b>53.8 ± 2.74</b>	<b>33.8 ± 2.16</b>	83.4 ± 0.64
28	2	1.61 M	46.6 ± 2.56	37.1 ± 2.47	83.6 ± 0.54
10	8	3.72 M	<b>59.7 ± 2.33</b>	<b>29.4 ± 2.52</b>	84.8 ± 0.49
16	4	3.24 M	50.1 ± 2.59	37.0 ± 2.77	85.1 ± 0.45
28	3	3.58 M	49.4 ± 1.82	36.2 ± 1.98	84.7 ± 0.92

To this end, we compare several architectures on both Rotated MNIST and Split CIFAR-100 benchmarks. Table 3 and 4 show that on both benchmarks, wider and shallower models outperform their thinner and deeper counterparts. Note that to make sure that the number of parameters are roughly the same, we have added some new models (e.g., MLP with a width of 256).

### 5.2. Interacting with CL Algorithms

We have seen that increasing width automatically reduces forgetting without the need to employ any continual learning algorithm. One may wonder if the effect of width on performance remains additive even when a specialized CL algorithm is employed.

In Figure 7, we compare two replay-based methods, ER (Riemer et al., 2018) and A-GEM (Chaudhry et al., 2018), for various width factors using both MLPs and WideResNets. We use the replay buffer of sizes 125 (MNIST) and 100 (CIFAR-100) for both ER and A-GEM. The figure shows that, similar to vanilla fine-tuning, by increasing width, average accuracy improves, and forgetting gets reduced across the board. Hence, we conclude that the benefit of increasing the width is additive to the benefits brought by continual learning algorithms.






Figure 7. Average accuracy (blue) and forgetting (red) of CL methods for varying widths. Similar to the naive fine-tuning case, in the presence of other CL algorithms, wider networks perform better.

In Appendix B.7, we provide additional results for more algorithms such as Mode Connectivity SGD (Mirzadeh et al., 2021), LwF (Li and Hoiem, 2018), and EWC (Kirkpatrick et al., 2017) on Rotated MNIST. We refer the reader to (Mirzadeh et al., 2022) for additional results on other benchmarks and architectures such as CNNs, ResNets (He et al., 2016), and Vision Transformers (Dosovitskiy et al., 2021).

## 6. Discussion and Conclusion

In this paper, we studied a previously unexplored connection between the width of the network and catastrophic forgetting (CF)—an important bottleneck for continual learning (CL) systems. Through different experiments, we hypothesize that high orthogonality between the gradients of different tasks induced by the wider networks, increased gradient sparsity, and a lazy training regime, i.e., the decreased distance of updated parameters from their initialization are the main factors alleviating catastrophic forgetting. We highlighted the contrasts between over-parametrization with increasing width to increasing depth and argued that increasing depth might have negative consequences for forgetting while acknowledging its importance for learning representations.

The scope of current work is to show a surprising and, at the same time, very intuitive correlation between width and forgetting. One should be cautious of utilizing increasing

width to tackle catastrophic forgetting as it is undoubtedly an inefficient decision in terms of computation time, energy and cost. By showing the connection between width and CF, our purpose here is to enhance our basic understanding of how over-parametrization, both in terms of width and depth, interacts with catastrophic forgetting. The extent of over-parametrization should be an important factor when designing continual learning benchmarks.

We hope this work acts as a stepping stone for studying the interaction between the inherent properties of neural networks and catastrophic forgetting. In addition to varying architecture sizes, modern neural networks employ several training tricks that are not well studied with respect to catastrophic forgetting. Understanding these effects, such as activation functions, normalization schemes, skip connections, etc., remains an interesting future work. Very recently, Mirzadeh et al. (2022) have studied the significance of architectures in continual learning. They have extended our main results for other architectures (e.g., Vision Transformers (Dosovitskiy et al., 2021)) on large-scale benchmarks. In addition, they show that the role of architecture in continual learning can be as important as the algorithms, which suggests the intersection of continual learning and neural network architecture is an interesting and under-explored research direction.

## Acknowledgments

The authors are grateful to Amal Rannen-Triki, Timothy Nguyen, Hooman Shahrokhi, and Anonymous Reviewers for their valuable comments and feedback on this work.

## References

- Aljundi, R., Babiloni, F., Elhoseiny, M., Rohrbach, M., and Tuytelaars, T. (2018). Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 139–154.
- Aljundi, R., Chakravarty, P., and Tuytelaars, T. (2017). Expert gate: Lifelong learning with a network of experts. In *CVPR*, pages 7120–7129.
- Allen-Zhu, Z., Li, Y., and Liang, Y. (2019). Learning and generalization in overparameterized neural networks, going beyond two layers. *Advances in neural information processing systems*.
- Arora, S., Du, S. S., Hu, W., Li, Z., Salakhutdinov, R., and Wang, R. (2019). On exact computation with an infinitely wide neural net. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 8141–8150.
- Balaji, Y., Farajtabar, M., Yin, D., Mott, A., and Li, A.

- (2020). The effectiveness of memory replay in large scale continual learning. *arXiv preprint arXiv:2010.02418*.
- Beaulieu, S., Frati, L., Miconi, T., Lehman, J., Stanley, K. O., Clune, J., and Cheney, N. (2020). Learning to continually learn. In *ECAI 2020 - 24th European Conference on Artificial Intelligence*.
- Bengio, Y., Simard, P., and Frasconi, P. (1994). Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on neural networks*, 5(2):157–166.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020*.
- Chang, M., Gupta, A., Levine, S., and Griffiths, T. L. (2018). Automatically composing representation transformations as a means for generalization. In *ICML workshop Neural Abstract Machines and Program Induction v2*.
- Chaudhry, A., Gordo, A., Dokania, P. K., Torr, P. H. S., and Lopez-Paz, D. (2021). Using hindsight to anchor past knowledge in continual learning. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021*.
- Chaudhry, A., Khan, N., Dokania, P. K., and Torr, P. H. (2020). Continual learning in low-rank orthogonal subspaces. In *Advances in Neural Information Processing Systems*.
- Chaudhry, A., Ranzato, M., Rohrbach, M., and Elhoseiny, M. (2018). Efficient lifelong learning with a-gem. In *International Conference on Learning Representations*.
- Chaudhry, A., Rohrbach, M., Elhoseiny, M., Ajanthan, T., Dokania, P. K., Torr, P. H. S., and Ranzato, M. (2019). On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486*.
- Chizat, L., Oyallon, E., and Bach, F. (2019). On lazy training in differentiable programming. In *NeurIPS 2019-33rd Conference on Neural Information Processing Systems*, pages 2937–2947.
- Doan, T., Bannani, M. A., Mazouze, B., Rabusseau, G., and Alquier, P. (2021). A theoretical analysis of catastrophic forgetting through the ntk overlap matrix. In *International Conference on Artificial Intelligence and Statistics*.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. (2021). An image is worth 16x16 words: Transformers for image recognition at scale. In *9th International Conference on Learning Representations, ICLR, 2021*.
- Du, S., Lee, J., Li, H., Wang, L., and Zhai, X. (2019). Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, pages 1675–1685. PMLR.
- Farajtabar, M., Azizan, N., Mott, A., and Li, A. (2020). Orthogonal gradient descent for continual learning. In *International Conference on Artificial Intelligence and Statistics*, pages 3762–3773. PMLR.
- Fernando, C., Banarse, D., Blundell, C., Zwols, Y., Ha, D., Rusu, A. A., Pritzel, A., and Wierstra, D. (2017). Pathnet: Evolution channels gradient descent in super neural networks. *arXiv preprint arXiv:1701.08734*.
- Ferran Alet, Tomas Lozano-Perez, L. P. K. (2018). Modular meta-learning. *arXiv preprint arXiv:1806.10166v1*.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep learning*. MIT press.
- Goodfellow, I. J., Mirza, M., Xiao, D., Courville, A., and Bengio, Y. (2013). An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Hombaiah, S. A., Chen, T., Zhang, M., Bendersky, M., and Najork, M. (2021). Dynamic language models for continuously evolving content. In *KDD '21: The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.
- Hsu, Y.-C., Liu, Y.-C., and Kira, Z. (2018). Re-evaluating continual learning scenarios: A categorization and case for strong baselines. *arXiv preprint arXiv:1810.12488*.
- Jacot, A., Gabriel, F., and Hongler, C. (2018). Neural tangent kernel: convergence and generalization in neural networks. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 8580–8589.
- Javed, K. and White, M. (2019). Meta-learning representations for continual learning. In *Advances in Neural Information Processing Systems*, pages 1820–1830.
- Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., and Amodei, D. (2020). Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Kirichenko, P., Farajtabar, M., Rao, D., Lakshminarayanan, B., Levine, N., Li, A., Hu, H., Wilson, A. G., and Pascanu, R. (2021). Task-agnostic continual learning with

- hybrid probabilistic models. In *ICML Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models*.
- Kirkpatrick, J. N., Pascanu, R., Rabinowitz, N. C., Veness, J., and et. al. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences of the United States of America*, 114 13:3521–3526.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105.
- Lazaridou, A., Kuncoro, A., Gribovskaya, E., Agrawal, D., Liska, A., Terzi, T., Gimenez, M., d’Autume, C. d. M., Ruder, S., Yogatama, D., et al. (2021). Pitfalls of static language modelling. *arXiv preprint arXiv:2102.01951*.
- Lee, J., Xiao, L., Schoenholz, S., Bahri, Y., Novak, R., Sohl-Dickstein, J., and Pennington, J. (2019). Wide neural networks of any depth evolve as linear models under gradient descent. *Advances in neural information processing systems*, 32:8572–8583.
- Li, X., Zhou, Y., Wu, T., Socher, R., and Xiong, C. (2019). Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In *Proceedings of the 36th International Conference on Machine Learning, ICML, Proceedings of Machine Learning Research*.
- Li, Z. and Hoiem, D. (2018). Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40:2935–2947.
- Lopez-Paz, D. and Ranzato, M. (2017). Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, pages 6467–6476.
- McClelland, J. L., McNaughton, B. L., and O’Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological review*, 102(3):419.
- McCloskey, M. and Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of Learning and Motivation*, 24:109–165.
- Mehta, S. V., Patil, D., Chandar, S., and Strubell, E. (2021). An empirical investigation of the role of pre-training in lifelong learning. *ICML CL Workshop*.
- Mirzadeh, S. I., Chaudhry, A., Yin, D., Nguyen, T., Pascanu, R., Gorur, D., and Farajtabar, M. (2022). Architecture matters in continual learning. *ArXiv*, abs/2202.00275.
- Mirzadeh, S. I., Farajtabar, M., and Ghasemzadeh, H. (2020a). Dropout as an implicit gating mechanism for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 232–233.
- Mirzadeh, S. I., Farajtabar, M., Gorur, D., Pascanu, R., and Ghasemzadeh, H. (2021). Linear mode connectivity in multitask and continual learning. In *International Conference on Learning Representations*.
- Mirzadeh, S. I., Farajtabar, M., Pascanu, R., and Ghasemzadeh, H. (2020b). Understanding the role of training regimes in continual learning. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020*.
- Nguyen, C. V., Achille, A., Lam, M., Hassner, T., Mahadevan, V., and Soatto, S. (2019). Toward understanding catastrophic forgetting in continual learning. *arXiv preprint arXiv:1908.01091*.
- Nguyen, C. V., Li, Y., Bui, T. D., and Turner, R. E. (2018). Variational continual learning. In *6th International Conference on Learning Representations, ICLR 2018*.
- Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., and Wermter, S. (2019). Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71.
- Pascanu, R., Mikolov, T., and Bengio, Y. (2013). On the difficulty of training recurrent neural networks. In *International conference on machine learning*, pages 1310–1318. PMLR.
- Ramasesh, V. V., Lewkowycz, A., and Dyer, E. (2022). Effect of scale on catastrophic forgetting in neural networks. In *International Conference on Learning Representations*.
- Rebuffi, S.-A., Kolesnikov, A. I., Sperl, G., and Lampert, C. H. (2016). icarl: Incremental classifier and representation learning. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5533–5542.
- Riemer, M., Cases, I., Ajemian, R., Liu, M., Rish, I., Tu, Y., and Tesauro, G. (2018). Learning to learn without forgetting by maximizing transfer and minimizing interference. In *International Conference on Learning Representations*.
- Ring, M. B. (1995). *Continual learning in reinforcement environments*. PhD thesis, University of Texas at Austin, TX, USA.

- Rolnick, D., Ahuja, A., Schwarz, J., Lillicrap, T. P., and Wayne, G. (2019). Experience replay for continual learning. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019*.
- Rosenbaum, C., Klinger, T., and Riemer, M. (2018). Routing networks: Adaptive selection of non-linear functions for multi-task learning. In *International Conference on Learning Representations*.
- Rusu, A. A., Rabinowitz, N. C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., and Hadsell, R. (2016). Progressive neural networks. *arXiv preprint arXiv:1606.04671*.
- Shin, H., Lee, J. K., Kim, J., and Kim, J. (2017). Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems*, pages 2990–2999.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. (2016). Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826.
- Thrun, S. (1995). A lifelong learning perspective for mobile robot control. In *Intelligent robots and systems*, pages 201–214. Elsevier.
- Titsias, M. K., Schwarz, J., Matthews, A. G. d. G., Pascanu, R., and Teh, Y. W. (2019). Functional regularisation for continual learning using gaussian processes. *arXiv preprint arXiv:1901.11356*.
- Toneva, M., Sordani, A., Combes, R. T. d., Trischler, A., Bengio, Y., and Gordon, G. J. (2019). An empirical study of example forgetting during deep neural network learning. In *7th International Conference on Learning Representations, ICLR 2019*.
- Veniat, T., Denoyer, L., and Ranzato, M. (2021). Efficient continual learning with modular networks and task-driven priors. In *ICLR*.
- Wallingford, M., Kusupati, A., Alizadeh-Vahid, K., Walsman, A., Kembhavi, A., and Farhadi, A. (2020). In the wild: From ml models to pragmatic ml systems. *ArXiv*, abs/2007.02519.
- Wortsman, M., Ramanujan, V., Liu, R., Kembhavi, A., Rastegari, M., Yosinski, J., and Farhadi, A. (2020). Supermasks in superposition. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020*.
- Xu, J. and Zhu, Z. (2018). Reinforced continual learning. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018*.
- Yin, D., Farajtabar, M., Li, A., Levine, N., and Mott, A. (2020). Optimization and generalization of regularization-based continual learning: a loss approximation viewpoint. *arXiv preprint arXiv:2006.10974*.
- Yoon, J., Yang, E., Lee, J., and Hwang, S. J. (2018). Lifelong learning with dynamically expandable networks. In *Sixth International Conference on Learning Representations. ICLR*.
- Zagoruyko, S. and Komodakis, N. (2016). Wide residual networks. In *British Machine Vision Conference 2016*. British Machine Vision Association.
- Zenke, F., Poole, B., and Ganguli, S. (2017). Continual learning through synaptic intelligence. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3987–3995. JMLR.
- Zou, D., Cao, Y., Zhou, D., and Gu, Q. (2020). Gradient descent optimizes over-parameterized deep relu networks. *Machine Learning*, 109(3):467–492.

---

# Wide Neural Networks Forget Less Catastrophically: Supplementary Material

---

## A. Experimental Setup Details

In this section, we discuss our experimental setup, including the design choices for benchmarks, architectures, and hyper-parameters.

### A.1. Design Choices

#### BENCHMARKS

We have chosen the rotated MNIST benchmark since it is a common benchmark in the research literature and, unlike the permuted MNIST, has a meaningful shift across tasks. The number of 5 tasks on MNIST is also typical in the literature (Farajtabar et al., 2020; Mirzadeh et al., 2020b). While the rotation of 10 degrees per task is more common in the literature, we have increased the rotation to 22.5 degrees per task to make the benchmark more difficult and the performance gap between different models become more visible (Mirzadeh et al., 2021).

While the rotated MNIST benchmark is domain-incremental, we have chosen the split CIFAR-100, which is a common task-incremental benchmark in the literature. Moreover, the split CIFAR-100 has 20 tasks, which helps support our arguments when the number of tasks increases. Our goal was to show our results in various setups such as domain versus task incremental, small versus large number of tasks, and single-head versus multi-head learning.

#### ARCHITECTURES

On the MNIST benchmark, MLP is a popular architecture in the literature (Wortsman et al., 2020; Kirkpatrick et al., 2017). However, on the split CIFAR-100, ResNet-18 is a more popular choice (Chaudhry et al., 2021; Mirzadeh et al., 2021). However, WideResNets (WRN) have two essential properties that are more suitable for this study: first, they share similar architectural properties with ResNets (e.g., skip-connections, batch normalization layers, etc.), and second, they have the nice property of ease of scaling both width and depth.

We note that for the experiments that involved increasing depth, we have not chosen the thinnest models as a starting point. For instance, in Fig. 2, we have used MLP with a width of 128 and WideResNet with a width of 2. The reason is to make sure that the models are not too small, so by increasing the depth, the number of parameters increases in a meaningful manner.

### A.2. Hyperparameters

For all experiments, we use a grid to ensure that one specific set of hyper-parameters is not favoring a specific model. Moreover, the training regime plays an important role in continual learning performance (Mirzadeh et al., 2020b).

1. learning rate: [0.001, 0.01 (MNIST), 0.05 (CIFAR), 0.1]
2. batch size: [16, 32 (CIFAR), 64 (MLP)]
3. SGD momentum: [0.0 (MNIST), 0.8 (CIFAR)]
4. weight decay: [0.0 (MNIST), 0.0001 (CIFAR)]


For each set of hyper-parameters, we run the model using 5 different seeds for random initialization of the models. Moreover, to ensure that the specific ordering of tasks in the split CIFAR-100 is not playing a role in the results, we also use 3 additional seeds to shuffle the ordering of tasks.

## B. Additional Results


### B.1. Evolution of the Average Accuracy Throughout the Learning Experience

While tables 1 and 2 provide the “final” average accuracy across different widths, it will be interesting to also compare the “evolution” of the average accuracy in throughout of learning experience. Fig. 8a and 8b show this during the learning.

There are a few interesting observations we can make from these results. First, the wider models perform consistently better during the learning experience. Second, the performance gap between wide and thin models is not large initially; however, as learning progresses, the gap increases. Finally, it shows that the decrease in the average accuracy due to forgetting still happens for wide models. This points out to our discussion in Sec. 6 that even though increasing width has beneficial effects, it is not the ultimate solution to continual learning. However, it will be interesting to understand the reasons behind these beneficial effects and use them in future works to design better continual learning algorithms, and this is the main objective of this study.



(a) MLP on Rotated MNIST



(b) WideResNet on Split CIFAR-100

Figure 8. Evolution of the average accuracy during the continual learning experience: while the gap is not initially large, it grows as the number of tasks increases.

We note that while wider models learn task 1 with higher accuracy, the gap between wide and thin models grows as the number of tasks increases. For instance, in Fig. 8b, the accuracy of WRN-10-8 on task 1 is roughly 4% higher than WRN-10-1. However, when the learning ends, the accuracy gap between WRN-10-8 and WRN-10-1 grows to more than 12%. While part of this gap can be explained by the fact that WRN-10-8 is a better learner, we note that the majority of the gap comes from the fact that WRN-10-8 has substantially smaller forgetting, as shown in Table. 2.

### B.2. Varying Width

In Claim 4.1 in Sec. 4, we have theoretically studied the importance of width for a simple case. Here, we provide additional empirical evidence for the Rotated MNIST benchmark where we use a 3-layer MLP model and measure continual learning metrics for different widths (64, 256, 1024) at each layer.


As shown in Table. 5, we show that the width of the first layer plays a more important role compared to the second and third layers, which is consistent with our analysis. While the models that have a width of 1024 in their first layer have more parameters, we note that we have already discussed in Sec. 5 that simply increasing the number of parameters is not necessarily helpful for reducing the forgetting.

### B.3. Lazy Training Regime in Deeper Models


In Sec. 4.3, we have discussed the role of parameter changes in the network. In Fig. 5a, we have provided empirical evidence each parameter of wide models changes less during the learning. However, one important argument could be that the norm of parameters increases as the number of parameters increases, and this might be the reason behind the result that the relative distance of wider models is smaller. However, we show in Fig. 9b that this is not the case, and the relative distance for deeper models does change as depth increases. Fig. 9a is similar to Fig. 5a and was put here for the ease comparison.

Table 5. Performance metrics for different 3-layer MLP models on Rotated MNIST.

L1	L2	L3	Params	Average Accuracy	Average Forgetting	Learning Accuracy
64	64	64	59.2 K	67.9 $\pm$ 0.28	37.2 $\pm$ 0.56	97.2 $\pm$ 0.37
64	64	256	73.6 K	68.5 $\pm$ 0.99	37.1 $\pm$ 0.18	97.5 $\pm$ 0.67
64	64	1024	131.2 K	68.1 $\pm$ 0.64	37.2 $\pm$ 0.82	97.8 $\pm$ 0.28
64	256	64	84.0 K	67.7 $\pm$ 0.58	37.6 $\pm$ 0.79	97.7 $\pm$ 0.29
64	1024	64	183.1 K	68.2 $\pm$ 0.55	36.9 $\pm$ 0.72	97.8 $\pm$ 0.31
256	256	64	283.8 K	71.9 $\pm$ 0.61	32.8 $\pm$ 0.84	97.8 $\pm$ 0.28
256	256	256	335.1 K	71.8 $\pm$ 0.54	32.9 $\pm$ 0.63	98.1 $\pm$ 0.33
256	256	1024	540.2 K	71.4 $\pm$ 0.68	33.6 $\pm$ 0.83	98.0 $\pm$ 0.37
1024	64	64	874.2 K	73.1 $\pm$ 0.94	29.8 $\pm$ 1.10	97.1 $\pm$ 0.72
1024	64	1024	946.2 K	73.9 $\pm$ 0.69	29.4 $\pm$ 0.76	97.4 $\pm$ 0.52
1024	1024	64	1.92 M	74.3 $\pm$ 0.40	28.5 $\pm$ 0.51	97.1 $\pm$ 0.67
1024	1024	256	2.12 M	74.4 $\pm$ 0.35	28.4 $\pm$ 0.50	97.1 $\pm$ 0.66
1024	1024	1024	2.91 M	75.1 $\pm$ 0.41	27.8 $\pm$ 0.52	97.3 $\pm$ 0.60



(a) MLP on Rotated MNIST (different widths)



(b) MLP on Rotated MNIST (different depths)

Figure 9. The distance between  $w_1^*$ , solution of task 1, and  $w_t^*$ , solution of the  $t$ -th task, for  $t > 1$  over the course of training. Wider networks remain closer to the solution of the first task, but deeper models do not show the same property.

#### B.4. Detailed Results for Other Methods

In Fig. 7 of Sec. 5, we have shown the average forgetting and accuracy of the A-GEM and Experience Replay(ER) algorithms to illustrate that the gain due to increasing width also happens in the presence of other popular continual learning algorithms. Tables 6 and 7 show the exact numbers for the mentioned figures, and we can see that by increasing width, the average accuracy increases and average forgetting decreases consistently.

Moreover, we have included the data from tables 1 and 2 for ease of comparison. An interesting observation from these results could be that adding replay buffer or gradient manipulation by the A-GEM algorithm has an additive effect to the gain due to width.

#### B.5. Orthogonality Between Task 2 and Subsequent Tasks


To show that our conclusion in Section 4 about orthogonality of gradients holds for the next tasks, in this section, we provide a similar figure to Fig. 3 but for task 2. In Fig. 10, we plot the angle between the gradient vector at the minimum of the second task and that of the subsequent tasks for various width factors. We can see that for wider models, the trend of gradients being orthogonal continues after learning the second task as well.

Table 6. Rotated MNIST: the gain due to width for different CL algorithms.


Width	Method	Average Accuracy	Average Forgetting
32	Naive	65.9 $\pm$ 1.00	36.9 $\pm$ 1.27
128	Naive	70.8 $\pm$ 0.68	31.5 $\pm$ 0.92
512	Naive	72.6 $\pm$ 0.27	29.6 $\pm$ 0.36
2048	Naive	75.2 $\pm$ 0.34	26.7 $\pm$ 0.50
32	ER	68.9 $\pm$ 1.08	34.6 $\pm$ 1.29
128	ER	75.6 $\pm$ 0.55	27.7 $\pm$ 0.72
512	ER	77.3 $\pm$ 1.00	25.9 $\pm$ 1.28
2048	ER	78.6 $\pm$ 0.41	24.6 $\pm$ 0.45
32	A-GEM	70.2 $\pm$ 1.50	33.0 $\pm$ 1.91
128	A-GEM	74.1 $\pm$ 0.52	29.6 $\pm$ 0.73
512	A-GEM	76.8 $\pm$ 0.62	26.6 $\pm$ 0.74
2048	A-GEM	78.8 $\pm$ 0.34	24.3 $\pm$ 0.50

Table 7. Split CIFAR-100: the gain due to width for different CL Algorithms.

Width	Method	Average Accuracy	Average Forgetting
1	Naive	47.1 $\pm$ 2.60	37.3 $\pm$ 2.62
2	Naive	49.7 $\pm$ 1.51	34.9 $\pm$ 1.72
4	Naive	53.8 $\pm$ 2.74	33.8 $\pm$ 2.16
8	Naive	59.7 $\pm$ 2.33	29.4 $\pm$ 2.52
1	ER	51.6 $\pm$ 1.74	24.7 $\pm$ 0.86
2	ER	57.1 $\pm$ 0.98	23.1 $\pm$ 1.74
4	ER	59.0 $\pm$ 1.10	20.9 $\pm$ 2.45
8	ER	63.8 $\pm$ 0.37	18.8 $\pm$ 0.49
1	A-GEM	54.2 $\pm$ 0.93	21.5 $\pm$ 0.03
2	A-GEM	57.9 $\pm$ 1.32	20.0 $\pm$ 1.11
4	A-GEM	60.6 $\pm$ 0.24	18.4 $\pm$ 0.79
8	A-GEM	62.8 $\pm$ 0.24	17.7 $\pm$ 0.06



(a) MLP on Rotated MNIST



(b) WideResNet on Split CIFAR-100

Figure 10. The degree between the gradients of task 2 and the gradients of subsequent tasks

## B.6. Singular Values of Deeper Models

In Eq. (4.3) of Sec. 4.4, we have discussed that one condition that impacts our analysis, is the singular values of all intermediate layers. Fig. 11a shows the singular values for our shallowest model with depth of 2, and Fig. 11b shows the values for the deepest model. We can see that the top singular values are all greater than 1.

We note that to make the figures easier to read, we show the top-10 singular values at the end of tasks 1, 2, and 5.

Finally, we would like to provide some clarifications on our analysis in Sec. 4.4. We note that the bound provided for the gradient norm is unusable for  $\Lambda_i > 1$ , as it is an upper bound which grows exponentially – becomes more and more loose – with  $K$  (the number of layers). The purpose of providing the analysis is to potentially point to a scenario when the gradients on the earlier layers could increase with the depth. This *can* happen when the singular values of all the intermediate weight matrices are greater than 1, *i.e.*,  $\Lambda_i > 1, \forall i \in \{1, \dots, K-1\}$ . In Fig. 11, we empirically see that this indeed is the case for Rotated MNIST tested with varying MLP depths.

## B.7. Interacting With Additional Algorithms

In Sec. 5.2, we have shown that similar to the naive finetuning case, when using A-GEM (Chaudhry et al., 2018), and Experience Replay, wide networks perform better for both Rotated MNIST and Split CIFAR-100 benchmarks.

Here, we extend the result for other algorithms on Rotated MNIST (with 5 tasks). We use Mode Connectivity SGD (MC-SGD) (Mirzadeh et al., 2021), which is the state-of-the-art algorithm in continual learning, in addition to Learning without Forgetting (LwF) (Li and Hoiem, 2018) and Elastic Weight Consolidation (EWC) (Kirkpatrick et al., 2017) methods.






Figure 11. Top-10 singular values of each layer during the learning: top singular values are larger than 1.

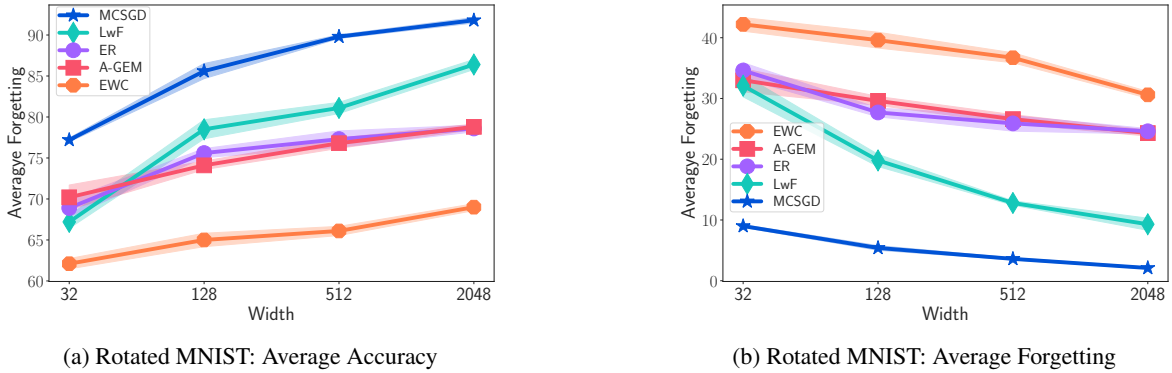


Figure 12. Rotated MNIST: Similar to the naive fine-tuning case, in the presence of other CL algorithms, wider networks perform better.

Fig. 12 shows that for all algorithms, by increasing the width, the average accuracy increases and average forgetting decreases. However, interestingly, the performance gain varies across algorithms, and methods such as MC-SGD and LwF benefit more from the increased width compared to other algorithms. An interesting future direction is to study this phenomenon in more detail.

### C. Task Construction in Claim 4.1

Consider a learning problem where the input space is  $\mathbb{R}^d$  and the output space is  $\mathbb{R}$ , and two model classes,  $\mathcal{F}_1, \mathcal{F}_2 \subseteq \{\mathbb{R}^d \mapsto \mathbb{R}\}$ , where  $\mathcal{F}_1$  is parameterized by  $w \in \mathbb{R}^d$ :

$$\mathcal{F}_1 := \{f_w : f_w(x) = \langle x, w \rangle, w \in \mathbb{R}^d\},$$

and  $\mathcal{F}_2$  is parameterized by  $U \in \mathbb{R}^{h \times d}$  and  $v \in \mathbb{R}^h$ :

$$\mathcal{F}_2 := \{f : f_{U,v}(x) = v^\top Ux, U \in \mathbb{R}^{h \times d}, v \in \mathbb{R}^h\}.$$

This means that  $\mathcal{F}_1$  is the set of linear models, and  $\mathcal{F}_2$  is the set of two-layer linear networks. Clearly, we have  $\mathcal{F}_1 = \mathcal{F}_2$ , i.e., they have the same capacity, since there is no non-linearity used in  $\mathcal{F}_2$ . For any model  $f \in \mathcal{F}_1$  or  $\mathcal{F}_2$ , and data point  $(x, y) \in \mathbb{R}^d \times \mathbb{R}$ , we use the quadratic loss function  $\ell(f; x, y) = \frac{1}{2}(y - f(x))^2$ .

Now consider two tasks with training data points  $\{(x_{1,i}, y_{1,i})\}_{i=1}^n$  and  $\{(x_{2,i}, y_{2,i})\}_{i=1}^n$ . We denote by  $X_t = [x_{t,i} \cdots x_{t,n}]^\top \in \mathbb{R}^{n \times d}$  and  $Y_t = [y_{t,i} \cdots y_{t,n}]^\top \in \mathbb{R}^n$  the data matrix and label vector for task  $t$  ( $t = 1, 2$ ), respectively. Define the training loss function:

$$L_t(f) = \frac{1}{n} \sum_{i=1}^n \ell(f; x_{t,i}, y_{t,i}), \quad t = 1, 2.$$

For either model class, the learner first trains the model with data for task 1, and obtains model  $f_1$  and then after training ends, the data for task 1 are all removed. The learner then trains the model with data for task 2. Let  $f_2$  be the model that the learner obtains after training on task 2. The *forgetting* metric is defined as

$$\text{fgt}(f_2, f_1) := L_1(f_2) - L_1(f_1).$$

We will consider the over-parameterized regime where  $d \gg n$ . We know that as long as  $X_1 X_1^\top$  is invertable, there exist models with zero loss on task 1. In this case, let us assume that the model is sufficiently trained on task 1, i.e., at the end of task 1, we have  $L_1(f_1) = 0$  for both model classes. Then we train the model on task 2 using gradient descent for  $T$  steps with learning rate  $\eta$ . Let  $w^0$  and  $w^T$  be the initial and end parameters for model class  $\mathcal{F}_1$  in task 2 training, respectively, and let  $(U^0, v^0), (U^T, v^T)$  be the initial and end parameters for model class  $\mathcal{F}_2$  in task 2 training, respectively. Then we have the following claim.

**Claim C.1** (Claim 4.1 restated). *There exist two tasks such that when we train task 2 using  $T$  steps of gradient descent, if we use model class  $\mathcal{F}_1$ , the amount of forgetting is strictly zero, i.e.,*

$$\text{fgt}(f_{w^T}, f_{w^0}) = 0;$$

*whereas if we use model class  $\mathcal{F}_2$ , the amount of forgetting can be positive, i.e.,*

$$\text{fgt}(f_{U^T, v^T}, f_{U^0, v^0}) \geq 0.$$

*In particular, if  $h \leq n$ , and the matrix  $X_1(U^0)^\top \in \mathbb{R}^{n \times h}$  has rank  $h$ , and  $v^T \neq v^0$ , then the forgetting for model class  $\mathcal{F}_2$  is positive, i.e.,  $\text{fgt}(f_{U^T, v^T}, f_{U^0, v^0}) > 0$ .*

*Proof.* We construct the two tasks in the following way. We assume that the input features of the two tasks satisfy

$$\langle x_{1,i}, x_{2,j} \rangle = 0 \quad \text{for all } 1 \leq i, j \leq n,$$

i.e.,  $X_1 X_2^\top = 0$ . This means that the data points in the two tasks are in two orthogonal subspaces.

When training on task 2, the learner use gradient descent algorithm with respect to the parameters. Here, we rewrite the loss function such that it becomes a function of the parameters. For model class  $\mathcal{F}_1$ , we let

$$L_2^1(w) := \frac{1}{2n} \|Y_2 - X_2 w\|_2^2,$$

and for function class  $\mathcal{F}_2$ , we let

$$L_2^2(U, v) := \frac{1}{2n} \|Y_2 - X_2 U^\top v\|_2^2.$$

Consider model class  $\mathcal{F}_1$ . With learning rate  $\eta$ , the parameter update rule for model class  $\mathcal{F}_1$  is

$$w^{t+1} = w^t - \frac{\eta}{n} (X_2^\top X_2 w^t - X_2^\top Y_2).$$

Since  $X_1 X_2^\top = 0$ , we know  $X_1(w^{t+1} - w^t) = 0$ . Let  $w^0$  be the model parameter that we have at the beginning of the training of task 2. Then after  $T$  steps training on task 2, we have  $X_1 w^T = X_1 w^0$ . Thus the forgetting for function class  $\mathcal{F}_1$  is zero, i.e.,

$$\text{fgt}(f_{w^T}, f_{w^0}) = 0. \tag{C.1}$$

Now consider model class  $\mathcal{F}_2$ . The parameter update rule is

$$\begin{aligned} U^{t+1} &= U^t - \frac{\eta}{n} (v^t (v^t)^\top U^t X_2^\top X_2 - v^t Y_2^\top X_2) \\ v^{t+1} &= v^t - \frac{\eta}{n} (U^t X_2^\top X_2 (U^t)^\top v^t - U^t X_2^\top Y_2). \end{aligned}$$

Suppose that the second task starts with  $f_1$  with parameters  $U^0$  and  $v^0$ . We assume that task 1 is well trained, so  $L_1(U^0, v^0) = 0$ . Since  $X_1^\top X_2 = 0$ , we know that  $X_1((U^{t+1})^\top - (U^t)^\top) = 0$ , thus  $X_1(U^t)^\top = X_1(U^0)^\top$  for any  $t$ . Therefore

$$\begin{aligned} \text{fgt}(f_{U^t, v^t}, f_{U^0, v^0}) &= \frac{1}{2n} \|X_1(U^t)^\top v^t - Y_1\|_2^2 \\ &= \frac{1}{2n} \|X_1(U^0)^\top v^t - Y_1\|_2^2 \\ &= \frac{1}{2n} \|X_1(U^0)^\top (v^t - v^0)\|_2^2 \geq 0. \end{aligned}$$

One can also easily observe that if  $h \leq n$ , and the matrix  $X_1(U^0)^\top \in \mathbb{R}^{n \times h}$  has rank  $h$ , and  $v^t \neq v^0$ , then  $\text{fgt}(f_{U^t, v^t}, f_{U^0, v^0}) > 0$ . □