
ON UNDERSTANDING OF THE DYNAMICS OF MODEL CAPACITY IN CONTINUAL LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

The core issue in continual learning (CL) is balancing catastrophic forgetting of prior knowledge with generalization to new tasks, otherwise, known as the stability-plasticity dilemma. We argue that the dilemma is akin to the capacity (the networks’ ability to represent tasks) of the neural network (NN) in the CL setting. Within this context, this work introduces “CL’s effective model capacity (CLEMC)” to understand the dynamical behavior of stability-plasticity balance point in the CL setting. We define CLEMC as a function of the NN, the task data, and the optimization procedure. Leveraging CLEMC, we demonstrate that the capacity is non-stationary and regardless of the NN architecture and optimization method, the network’s ability to represent new tasks diminishes if the incoming tasks’ data distributions differ from previous ones. We formulate these results using dynamical systems’ theory and conduct extensive experiments to complement the findings. Our analysis extends from a small feed-forward (FNN) and convolutional networks (CNN) to medium sized graph neural networks (GNN) to transformer-based large language models (LLM) with millions of parameters.

1 INTRODUCTION

Humans can easily adapt to multiple tasks. However, when neural networks (NN) seek to mimic this behavior [50], they exhibit a phenomenon known as catastrophic forgetting, where the model forgets older tasks while learning new ones [50]. This well recorded issue is seen irrespective of the NN architecture, from simple linear adaptive systems [33] to massive large language models [47, 38]. The field of artificial intelligence that studies this phenomenon is known as continual learning (CL).

In recent years, numerous studies in CL [12, 51, 29, 6, 46] have shown that the core issue behind CL is a trade-off between forgetting prior information (catastrophic forgetting) and learning new information (generalization), also known as the stability-plasticity dilemma. This trade-off captures the relationship between data and the optimization procedure, but that is only part of the picture. Independent lines of inquiry have also shown that over-parameterized NNs play a crucial role in achieving optimal performance in the CL paradigm [40, 19, 18]. While, [41] study the role of optimization characteristics, [25] study the learnability of CL problem when subsequent distributions are overlapping. Although, all of these works provide different but overlapping insights, they look at the different sides of the problem such as model and data in [31] or the model and optimization procedure in [25, 41, 6] and do not consider the complex interplay between the model/optimization/tasks.

In this work, we aim to provide holistic insights into this interplay and establish a foundational understanding of the effect of NN capacity (stability plasticity balance point) on CL optimization in the presence of a series of tasks. To this end, we extend the definition of capacity from [42] to the CL paradigm, describing capacity (Def 2) as the effect of network architecture, hyperparameters, and weights measured through the cost function. We then elucidate the connection between capacity and the stability-plasticity dilemma (Lemma 1 and Fig. 1) and show that stability plasticity balance point is akin to the networks ability to represent tasks. With this theoretical framework, we show that the smallest possible change in the network’s capacity is a function of changes in weights and tasks (Theorem 1). Our main results (Theorems 2, 3) demonstrate that capacity, and by extension, the balance point, is non-stationary in the CL setting. The key conclusion of this work is:

"The CL capacity of a model is a function of the interplay between model, data, and the optimization procedure. Moreover, regardless of the type of NN, optimization procedure, or task, the network eventually becomes unsuitable for representing the tasks if each subsequent task differs from the previous one even by a small constant."

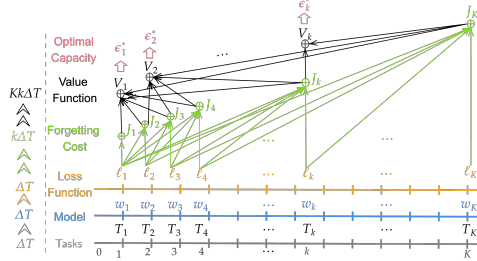


Figure 1: An illustration of how the forgetting cost, value function and CL effective capacity is calculated.

show that the capacity of a feed-forward NN (FNN) diverges even when the two major classes of CL methods, *experience replay* and *regularization approaches* are utilized. In the second, we extend the study to a standard convolutional NN (CNN) with the Omniglot dataset [3]. We also show results with a graph neural network in third case study and finally, develop a detailed study using large language models (LLM) to demonstrate our results. Our findings confirm that our theoretical results hold even when we scale from a simple FNN to a 134 million parameter LLM. Proofs complete with all assumptions are provided in the supplementary files.

2 RELATED WORKS

Starting from [17] in 1999 to [18] in 2024, numerous works have attempted to model/reduce catastrophic forgetting in neural networks. A simple taxonomy of recent published works reveals four categories: regularization-based [4, 27, 44, 29], model architecture-based [2, 10, 11, 14, 20, 30, 52], experience replay-based [7, 21, 28, 39, 53] and other optimization approaches for CL efficiency [1, 49, 54, 56]. This huge body of work is focused on improving empirical performance.

On the other hand, empirical attempts to study the characteristics of the CL problem have been made as well [12, 22, 32, 38]. For instance, [12] study the loss of plasticity in CL whereas [22, 32] study a phenomenon known as stability gap frequently observed in CL methods. The empirical investigative studies cover a wide range of neural network architectures as well, going from FNN/CNN in [12, 22, 32] to large language models in [38, 47]. Despite such a huge body of literature, there have only been a few attempts to study CL from a theoretical standpoint. The key reason behind this is that the NN learning problem in CL domain is rather complex to study requiring stringent assumptions that are scarcely held in practice. This is clearly seen from the few approaches that do theoretically analyze the problem. For instance, works in [18, 19, 13] study the effect of over parameterization and task similarity on forgetting with a linear model under two tasks. On the other hand [34] and [36] study the complete CL problem with a linear two layer NN. To the best of our knowledge, the only approach that does not make either a two task assumption or assume linearity of the model is [25] but instead focuses on the class incremental setting.

To obviate the necessity for such assumptions and provide a general framework to analyze CL, we take a Lyapunov analysis standpoint, a tool that has been used in the control literature [5]. In contrast with the existing literature, we analyze the CL problem through a dynamic programming-driven optimal control point of view following the perspective from [46]. The only assumptions required are twice differentiability and Lipschitz continuity of the loss function- two very practical assumptions in the NN learning domain and our analysis extends to a series of tasks. In a similar vein to [40] we also perform Taylor series approximation to get this differential equation characterization, however, our theoretical analysis easily extends from a simple FNN to a llm- a very novel contribution to the CL literature. To the best of our knowledge there has been no theoretical study, where the analysis considers a dynamical behavior of the CL problem that extends across FNN/CNN/GNN and LLM.

3 CONTINUAL LEARNING EFFECTIVE MODEL CAPACITY (CLEMC)

Let \mathbf{x} and \mathbf{y} be random variables corresponding to input and output probability spaces with support \mathcal{X} and \mathcal{Y} and $\mathcal{B}(\mathcal{X})$ and $\mathcal{B}(\mathcal{Y})$ representing the corresponding Borel algebras. Define \mathbf{t} as a random variable denoting the joint space of $\mathbf{x} \times \mathbf{y}$ with a model $f_{(\mathbf{w}, \mathbf{h})} : \mathcal{X} \rightarrow \mathcal{Y}$ being specified using weights \mathbf{w} and hyperparameters \mathbf{h} . Given compact sets \mathcal{W} over \mathbf{w} and \mathcal{H} over \mathbf{h} , the goal is to learn the weights by searching over the hypothesis space $\mathcal{f} = \{f_{(\mathbf{w}, \mathbf{h})}, \forall \mathbf{h} \in \mathcal{H}, \mathbf{w} \in \mathcal{W}\}$ through a loss function $\ell(\mathbf{t}; \mathbf{w}, \mathbf{h})$. In this context, we characterize the notion of capacity of the model as follows.

Effective Model Capacity: We will assume that $\ell(\mathbf{t}; \mathbf{w}, \mathbf{h})$ is continuous and twice differentiable over the support $\mathcal{X} \times \mathcal{Y}$ or \mathcal{X} , and the compact set \mathcal{W} . Under these assumptions, let $\ell_{\min} = \mathcal{O}(\mathcal{T}, \mathcal{H} \times \mathcal{W}) = \min_{\mathbf{h} \times \mathbf{w} \in \mathcal{H} \times \mathcal{W}} E_{\mathbf{t} \in \mathcal{T}}[\ell(\mathbf{t}; \mathbf{w}, \mathbf{h})]$ be the optimization procedure with \mathcal{T} being a dataset of samples \mathbf{t} with $\mathcal{T} \subset \mathcal{B}(\mathcal{T})$. Then, \mathcal{O} seeks to find the best hyperparameter/architecture configurations $\mathbf{h}^* \in \mathcal{H}$ and weights $\mathbf{w}^* \in \mathcal{W}$. Given this setting, we define the effective model capacity (the upper/lower bounds derived in the appendix) as the smallest achievable loss value using \mathcal{O} that remains unchanged even when additional data or training is used.

Definition 1 (Effective Model Capacity (EMC)). *Given $\mathcal{H} \times \mathcal{W}, \mathcal{T} \in \mathcal{B}(\mathcal{T})$ and $\mathcal{O}(\mathcal{T}, \mathcal{H} \times \mathcal{W})$, the EMC of the model f is given as*

$$\epsilon = \min_{\mathcal{T} \in \mathcal{B}(\mathcal{T})} [\mathcal{O}(\mathcal{T}, \mathcal{H} \times \mathcal{W})] = \min_{\mathcal{T} \in \mathcal{B}(\mathcal{T})} \left[\min_{\mathbf{h} \times \mathbf{w} \in \mathcal{H} \times \mathcal{W}} E_{\mathbf{t} \in \mathcal{T}}[\ell(\mathbf{t}; \mathbf{w}, \mathbf{h})] \right] \quad (\text{EMC})$$

Def EMC takes an approximation error perspective (as in [43]), however, unlike [43], (EMC) depends on the optimization procedure, the model performance and the dataset. It is also similar to the capacity definition in [42], with the key distinction being that [42] focuses on the number of data points that are properly represented by the model. However, this way of defining capacity is often inadequate because numerical superiority over samples alone (without considering the data distribution characteristics) doesn't ensure model usefulness [15]. Since, a CL problem requires careful attention to the distribution characteristics, we define capacity through the forgetting loss.

Characterizing the CL Balance Point: CL involves learning a sequence of tasks indexed by $k \in [0, K], K \in \mathbb{N}$, where a task k is represented by its dataset $\mathcal{T}(k)$. The collection of all tasks until k can then be denoted as $\mathcal{T}_{(k)} = \{\mathcal{T}(0), \mathcal{T}(1), \dots, \mathcal{T}(k)\}$ with $\mathcal{T}_{(k)}$ being the cumulative support. Given a feasible weight set $\mathcal{W}_{(k)}$, hyperparameter set \mathcal{H} , and loss function $\ell_{(k)}(\mathbf{t}_{(k)}; \mathbf{w}_{(k)}, \mathbf{h})$, the model at k is denoted by $f_{(\mathbf{w}_{(k)}, \mathbf{h})}$, the goal of CL is to maintain memory of all observed tasks, then, the CL forgetting cost for the interval $k = [0, k]$ is given as

$$\min_{\mathbf{w}_{(k)} \in \mathcal{W}_{(k)}} J(\mathcal{T}_{(k)}; \mathbf{w}_{(k)}, \mathbf{h}) = \min_{\mathbf{w}_{(k)} \in \mathcal{W}_{(k)}} \sum_{i=0}^k \gamma^{(i)} \left[E_{\mathbf{t} \in \mathcal{T}(i)}[\ell(\mathbf{t}; \mathbf{w}_{(k)}, \mathbf{h})] \right], \mathcal{T}(i) \in \mathcal{T}_{(k)}, (J_F)$$

where, $\gamma^{(k)}$ ensures boundedness of $J(\mathcal{T}_{(k)}; \mathbf{w}_{(k)}, \mathbf{h})$ (see [46], Lemma 1). The forgetting cost formulation in (J_F) is the standard in the CL literature [41] but, has two key limitations [6, 19] that we highlight using the following illustrative example.

Example 1. *Consider three learning tasks with feasible regions $\mathcal{W}_1, \mathcal{W}_2$, and \mathcal{W}_3 , centered at ideal solutions $\mathbf{w}_1^*, \mathbf{w}_2^*$, and \mathbf{w}_3^* . The naive cost setup in (J_F) ignores the following interactions.*

Sequential Optimization: *Solving the first task (attaining \mathbf{w}_1^*) means the second task must start from \mathbf{w}_1^* . Therefore, \mathbf{w}_1^* and its distance from $\mathcal{W}_1 \cap \mathcal{W}_2$ (the feasible region all solutions that work on both tasks 1 and 2) determines how close we can get to \mathbf{w}_2^* . In general, as the optimal solution for tasks $[0, k-1]$ is used as the starting point for task k , the feasible region of the previous tasks has an influence on the subsequent task [13][Theorem 3.1].*

Influence of future tasks: *If the second task induces a significant deviation from \mathbf{w}_1^* , large forgetting is seen (see [13], Figure 1). Conversely, if the new task has no influence, there's no generalization.*

It is clear with this example that each tasks' solution has an influence on the future task and at the same time, future tasks performance dictates how well the the model can do on the present tasks. That is, there is an interplay between future tasks and the present task. Mathematically, a complete CL [46] characterization must therefore consider both the sequential optimization over tasks as well as how each tasks' solution impacts future tasks. For a fixed $\mathbf{h} \in \mathcal{H}$, the complete CL problem (an illustration of this cost function definition is given in Fig. 1) is

$$V_{(k)}^{(*)} = \min_{\mathbf{w}_{(i)} \in \{\mathcal{W}_{(i)}, i=k, k+1, \dots, K\}} \sum_{i=k}^K [J(\mathcal{T}_{(i)}; \mathbf{w}_{(i)}, \mathbf{h})]. \quad (\text{CL})$$

The optimization problem in (CL) provides the value function, where previous tasks are perfectly remembered (optimizing the sum of forgetting loss, (J_F)) and future tasks will be perfectly learnt (for task k , optimizing also for $[k+1, \dots, K]$ via successive update of model weights). That is, given a starting weight set $\mathbf{w}_1^* \in \mathcal{W}_1$, the solution to the CL problem with K expected tasks is $\{\mathbf{w}_1^* \in \mathcal{W}_1, \mathbf{w}_2^* \in \mathcal{W}_1 \cap \mathcal{W}_2, \mathbf{w}_3^* \in \mathcal{W}_1 \cap \mathcal{W}_2 \cap \mathcal{W}_3 \dots \mathbf{w}_K^* \in \cap_{k=1}^K \mathcal{W}_k\}$ and $V_{(k)}^{(*)}$ is the total cost (corresponding to the balance point). Naturally, the value of ℓ_{min} (see Def 1) corresponding to each of these \mathbf{w}_i^* , $i = 1, 2, 3, \dots, K$ describes how well the model performs at the respective stages of the CL problem and therefore (summation of the losses) quantifies capacity in the CL setting. We now extend Def 1 with above notion to define effective model capacity for a CL problem.

CL Effective Model Capacity and Balance Point

Definition 2 (Effective Model Capacity for CL (CLEMC)). *Fix $k \in [0, K]$, dataset $\mathcal{T}_{(k)}$ and choose $\mathcal{H} \times \mathcal{W}_{(k)}$. Provided with $\mathcal{O}_{(k)}(\mathcal{T}_{(k)}, \mathcal{H} \times \mathcal{W}_{(k)})$. The EMC at any k for $\mathbf{h} \in \mathcal{H}$ is given as $\max_{\forall k \in [k]} \min_{\mathbf{w}_{(k)} \in \mathcal{W}_{(k)}, \mathcal{T}_{(k)} \in \mathcal{B}(\mathcal{T}_{(k)})} J(\mathcal{T}_{(k)}, \mathbf{w}_{(k)}, \mathbf{h})$. Then, CLEMC is*

$$\epsilon_{(k)}^{(*)} = \sum_{i=k}^K \max_{\forall j \in [0, i]} \min_{\mathcal{T}_{(j)} \in \mathcal{B}(\mathcal{T}_{(j)})} \min_{\mathbf{w}_{(j)} \in \mathcal{W}_{(j)}} J(\mathcal{T}_{(j)}, \mathbf{w}_{(j)}, \mathbf{h}). \quad (\text{CLEMC})$$

Def (2) closely relates to the forgetting loss, as capacity at each k is defined by the highest forgetting loss in the interval $[0, k]$. For example, in a three-task scenario, the capacity at each task k is determined by the task it forgets the most. If the model learns ten tasks, then we obtain an EMC corresponding to each task, then, the $\epsilon_{(k)}^{(*)}$ is the sum of individual task capacities. Since the individual task capacities are proportional to the loss function, perfect representation of the underlying tasks is implied by $\epsilon_{(k)}^{(*)} = 0$ and representation (and capacity) gets poorer and poorer as $\epsilon_{(k)}^{(*)}$ increases.

Notably, $\epsilon_{(k)}^{(*)}$ measures the models' CL performance. Similar to (CLEMC), the measure of models' performance has also been defined proportional to the value of the forgetting loss. For instance, [25][Def 3.1] defines learnability as the gap between empirical risk and the smallest risk in the hypothesis space, but without the minimization over different data samples. Furthermore, [24][Theorem 1] suggests that necessary and sufficient conditions for good CL are proportional to effective learning on prior tasks, defined through the forgetting loss. In contrast with the above, where just loss on the prior tasks is considered, in Def (2), both future tasks and bias due to subsequent solution are also considered. The relationship between (CL) and (CLEMC) is formalized in the next lemma.

Lemma 1. *Consider $k \in [0, K]$ and define (J_F) with $\mathbf{w}_{(k)} \in \mathcal{W}_{(k)}, \mathbf{h}_{(k)} \in \mathcal{H}$. Assume $\forall k \in \mathbf{k}, \mathbf{h}_{(k)} = \mathbf{h}$ and define (CL) and (2). Then, it follows that*

$$\epsilon_{(k)}^{(*)} - \epsilon_{(k+1)}^{(*)} = \max_{\forall j \in [0, k]} \left\{ \min_{\mathcal{T}_{(j)} \in \mathcal{B}(\mathcal{T}_{(j)})} \left\{ \min_{\mathbf{w}_{(i)} \in \mathcal{W}_{(i)}, i=j, j+1, \dots, K} \right\} \right\} \\ + \sum_{p=0}^j \left\langle \nabla_{\mathcal{T}_{(p)}} V_{(j)}^{(*)}, \Delta \mathcal{T}_{(p)} \right\rangle$$

Lemma 1 quantifies the first order change in (CLEMC) as a function of the change in value function (cost corresponding to balance point). In essence, Lemma 1 implies that the change in the total effective capacity at the addition of a new task is proportional to the largest amount of infinitesimal change introduced by any task k to the balance point. In particular, for a network with maximum capacity, $\Delta \epsilon_{(k)}^{(*)}$ will be zero and no new task can change the value function or effect the balance point.

This means that the model can represent all the tasks in the interval $[0, K]$ perfectly. On the other hand, for a model with minimum capacity, there is a substantial change in the cost corresponding to the balance point which implies that the model cannot represent any task reasonably. In Lemma 1, the change in CL effective model capacity is due to two terms. The first term denotes the smallest change introduced by the weights (across all tasks) to the value function. The second term systematically quantifies the change introduced by each of the tasks observed till now (indicated by the sum). We can now intuitively summarize the main result of our paper.

If each subsequent task is different than the previous task, the cumulative change in tasks, $\Delta \mathcal{T}_{(p)}$, is going to lead to deteriorating capacity. In particular, the change in $\Delta \mathcal{T}_{(p)}$, is going to drive a change in weights, $\Delta \mathbf{w}_{(i)}$, which in turn drives a change in capacity. This interplay is going to accumulate as the number of tasks increases and lead to deteriorating capacity.

4 ANALYSIS

In this section, we perform a two-fold analysis to prove our main idea, ‘‘capacity diverges with if tasks change constantly’’. First, we formally prove this result. Later, we demonstrate experimentally, that the capacity diverges irrespective of the architecture or the data. An experimentally inclined reader can safely skip the theoretical analysis section and get the main conclusions experimentally, however we recommend reading through this section to get insights into why this divergence occurs.

4.1 THEORETICAL ANALYSIS

We begin by deriving a lower bound on the first difference of $\epsilon_{(k)}^{(*)}$ and then analyse the impact of the independent terms of the bound on the effective capacity.

Theorem 1. *The first difference in CLEMC given by Lemma 1 is lower bounded as*

$$\epsilon_{(k)}^{(*)} - \epsilon_{(k+1)}^{(*)} \geq \max_{\forall j \in [0, k]} \left\{ \sum_{p=0}^j \gamma^{(p)} \min_{\mathbf{T}(p) \in \mathcal{B}(\mathcal{T}(p))} \left\| \nabla_{\mathbf{w}_{(j)}^{(*)}} E_{\mathbf{t} \in \mathcal{T}(p)} [\ell(\mathbf{t}; \mathbf{w}_{(j)}^{(*)}, \mathbf{h})] \right\| \left\| \Delta \mathbf{w}_{(j)}^{(*)} \right\| \right. \quad (\text{LB}) \\ \left. + \sum_{p=0}^j \gamma^{(p)} \min_{\mathbf{T}(p) \in \mathcal{B}(\mathcal{T}(p))} \left\| \nabla_{\mathbf{T}(p)} \sum_{i=j}^K E_{\mathbf{t} \in \mathcal{T}(p)} [\ell(\mathbf{t}; \mathbf{w}_{(i)}^{(*)}, \mathbf{h})] \right\| \left\| \Delta \mathbf{T}(p) \right\| \right\}$$

It is straightforward to see that this lower bound in Theorem 1 is zero, given no change in tasks ($\Delta \mathbf{T}(p)$) or the weights ($\Delta \mathbf{w}_{(j)}^{(*)}$). However, in practice each time a task j is introduced to the CL problem, there is a change in the value function. This change is an accumulation of the impact of the new task j , on all the prior tasks that in the interval $[0, j]$ ($\sum_{p=0}^j$ at the outer of the two terms in (LB) accumulates this change). For each task p in this sum, (LB) is a function of two key terms, (I) ‘‘the norm of the gradient of the value function with respect to the solution of the CL problem at p^{th} task’’ and (II) ‘‘the norm of the change in the value function due to change in the data at the p^{th} task.’’ We now study the effects of each of these terms below. For the following analysis, we will define

$$\max_{\forall j \in [0, p]} \left\{ \left\| \Delta \mathbf{w}_{(p)}^{(*)} \right\| \right\} = d_{\text{MAX}}^p(\mathbf{w}^{(*)}) \text{ and } \max_{\forall j \in [0, p]} \left\{ \min_{\mathbf{T}(p) \in \mathcal{B}(\mathcal{T}(p))} \left\| \Delta \mathbf{T}(p) \right\| \right\} = d_{\text{MAX}}^p(\mathbf{T}).$$

(I)-Effect of weight update on Capacity: To illustrate the effect of weight update, we assume that experience replay (ER)-driven CL methods define either (i) a forgetting cost using all the available tasks, and/or (ii) utilize a regularizer on top of the forgetting cost [9]. We further assume that, at each task k the weights are updated for a total of I steps. Under these assumptions, we show that for both settings (i) and (ii) above, the effective capacity diverges. We now state the following theorem.

Theorem 2. *Let $\sum_{p \in [0, K]} d_{\text{MAX}}^p(\mathbf{T}) = c$ and $\mathbf{w}_{(k)}^{(i)}$ is updated for I steps at each k using the loss $\ell(\mathbf{t}_{(k)}^{(i)}; \mathbf{w}_{(k)}^{(i)}, \mathbf{h})$, then $\epsilon_{(k)}^{(*)} - \epsilon_{(K)}^{(*)}$ diverges for (i) $d_{\text{MAX}}^p(\mathbf{w}^{(*)}) \geq I\alpha_{(\text{MIN})}[L]$ and $d_{\text{MAX}}^p(\mathbf{w}^{(*)}) \geq I\alpha_{(\text{MIN})}[L + \beta L_{\mathcal{R}}]$, where $L, L_{\mathcal{R}}$ are Lipschitz constants, β is the penalty parameter, and $\alpha_{(\text{MIN})}$ is the smallest learning rate.*

Theorem 2 demonstrates an important and novel result in the CL literature. In an essence, for any standard CL algorithm in the literature with standard gradient driven optimization regime, capacity will diverge as long as the each subsequent tasks keeps accumulating constant albeit small differences. Therefore, all known CL algorithms have the potential to result in a model does not represent all the tasks reasonably. Moreover, this is uncontrollable because the tasks are unknown in prior.

(II)-Effect of Tasks on Capacity To demonstrate the effect of tasks on capacity.

Theorem 3. *Let the maximum change in subsequent tasks and weights be given by $d_{\text{MAX}}^p(\mathbf{T}) = c, \forall p$ and $\sum_{p \in [0, K]} d_{\text{MAX}}^p(\mathbf{w}^{(*)}) = \zeta, \zeta \in \mathbb{R}$ respectively. Then $\lim_{K \rightarrow \infty} [\epsilon_{(k)}^{(*)} - \epsilon_{(K)}^{(*)}] \rightarrow \infty$.*

Theorem 3 shows that when $d_{\text{MAX}}^p(\mathbf{T}) = c$ and the cumulative change (change across all tasks at all learning instances) in the weights is a small constant, the model becomes unsuitable to represent the tasks. The impact of task similarity on CL has also been studied in [36, 13, 25, 18]. In contrast with Theorem 3, [36, 13, 18] study the impact for a linear classifier. In particular, [18][Theorem 3] shows a monotonic decrease in forgetting cost as a function of similarity. For a two task case, Theorem 3 indicates the same result in [18][Theorem 3] as similar tasks will result in no change in capacity. At first, Theorem 3 might appear contradictory to [25][Theorem 3.7], however, our result actually aligns with [25][Theorem 3.7]. Note that in the case when the overlap between distributions will keep decreasing, the loss function will proportionately increase and the risk gap will diverge.

4.2 EXPERIMENTAL ANALYSIS

In this section, we aim to substantiate the theoretical results and to that end, we develop an array of experiments where we show that capacity diverges with respect to change in tasks irrespective of the type and scale of the model. We emphasize that, *this work does not present a new method nor does it pertain to demonstrating a new way of doing CL*, but, the goal is to elucidate how the shift in the data-distribution affects the neural network model. To illuminate on this perspective, we build our experiments on popular neural network architecture, namely: feed forward NN (FNN), convolutional NN (CNN), graph NN (GNN) and a transformer-based model. We argue that, for any particular model, the phenomenon of deteriorating capacity as observed on one dataset does translate to other datasets as well because the divergence of capacity is the function of how the NN model react to the shift in the data distribution. Therefore, we choose datasets that are easier to analyze but still relevant in the continual learning paradigm, both in the supervised and the unsupervised learning paradigm. In particular, we utilize a feedforward NN with a synthetic sine wave dataset [23], a convolutional NN with the Omniglot dataset [3] and a transformer-driven large language model (LLM) on a trillion (T) tokens dataset provided by RedPajama [8]. We execute FNN/CNN/GNN experiments using the JAX library and we utilize pytorch for the LLM experiments.

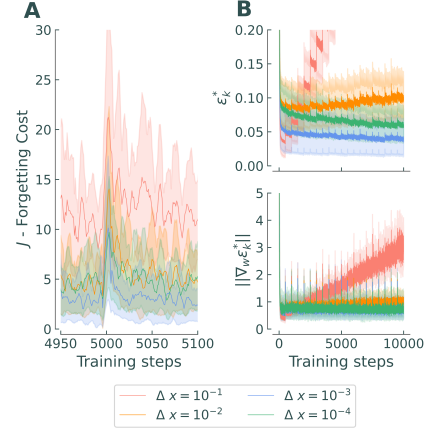


Figure 2: A: Forgetting cost with ER; B: (top) capacity; (bottom) the gradient of capacity with respect to weights as a function of training steps.

Case Study 1: Feed-forward NNs

Setup: For this experiment, we generate a total of twenty tasks, where each task is comprised of sine waves, generated by increasing the value of amplitude and frequency by a quantity $\|\Delta T(p)\|$ to indicate distribution shift. For analysis, we observe the trend of $\epsilon_k^{(*)}$ (capacity) for two standard methods in CL: Experience Replay (ER) shown in Fig. 2 and regularized ER shown in Fig. 3. We simulate four versions of this twenty task CL problem by choosing different values of Δx , i.e. $\|\Delta T(p)\| \in \{10^{-01}, 10^{-02}, 10^{-03}, 10^{-04}\}$ and learn twenty tasks for a total of 10 repetitions using mean squared error (MSE) as a cost function and 500 epochs per task.

Analysis of CL using ER: In panel A of Fig. 2, we plot MSE (averaged across 10 repetitions) and standard deviation (represented using a light region). We first note that, for any new task (we choose a random task at the middle of the learning process to illustrate this), there is an instantaneous increase in the forgetting cost which is then minimized by the optimizer, a phenomenon known as stability gap ([22]) in the CL literature. The more different

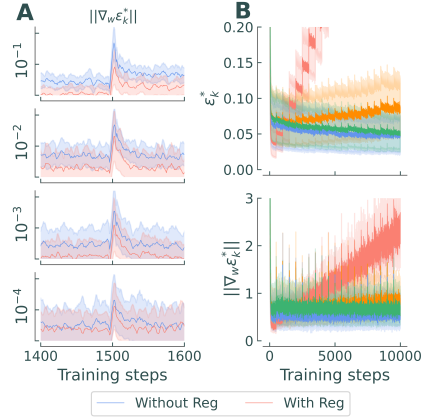


Figure 3: A: Forgetting cost with ER and L_2 regularization; B: capacity; (bottom) the gradient of capacity with respect to weights as a function of training steps under L_2 regularization.

the new task is, the larger the stability gap. Moreover, smaller the value of $\|\Delta T(p)\|$, the closer to zero, the forgetting cost appears to be. Our theoretical result in Theorem 3 precisely indicates that each small change in the task leads to a proportional change in the forgetting cost and by extension, the capacity. We see this trend also in Fig. 2, panel B, where we plot the change in capacity, $\epsilon_k^{(*)}$ with respect to update instances. For each new task, the same behavior as panel A is observed. Similar to Panel A, the capacity of the network gets worse proportional to $\|\Delta T(p)\|$ (a conclusion from Theorem 3). Moreover, the simulation also substantiates that vanilla ER, which is supposed to compensate for task differences does demonstrate this deteriorating behavior. This is observed in the top panel where the capacity curve deteriorates for change in Δx , (blue is poorer than green, which

worse than blue, orange is poorer and the red curve simply blows up) – an expected result shown in Theorem 2.

Analysis of CL using regularized ER: We plot the mean forgetting cost as a function of the training instances done using L_2 regularization in panels A and B in Fig.3. We expect regularization to improve model performance [35, 26] as the model can now adapt better data change (due to incoming tasks). In other words, the slope of curves observed in panels A and B of Fig. 2, should be smaller when a regularizer is utilized. Fig.3, Panel A, reinforces the observation that regularization applied to ER improves the slope of the capacity, i.e., $\epsilon_{(k)}^{(*)}$ deteriorates slower when regularization is added to forgetting loss (blue curve is better than the orange curve). This is true for all the values of $\|\Delta T(p)\|$. A curious thing to note is that, for smaller values of $\|\Delta T(p)\|$, $\nabla_w \epsilon_{(k)}^{(*)}$ values are rather close to each other. This is expected because, when the difference between two tasks is small, regularization only slightly improves the performance of ER. Moreover, as shown in Theorem 2, in spite of regularization, with a large enough $\|\Delta T(p)\|$, capacity increases drastically (the red curve corresponding to 10^{-01} increase very fast as seen in Fig.3, Panel B.

Case Study 2: Convolutional NNs

Setup: We now use the Omniglot dataset [9, 3]. Note that Omniglot dataset is a commonly used in continual, meta continual learning problems because of the presence of large numbers of tasks in opposition to MNIST and CIFAR, that are mostly image recognition datasets. We create a total of 10 classes and sequentially expose the CNN to one class at a time under the incremental class learning paradigm [37].

Analysis: Overall, all the conclusions from the previous case study does carry forward. The stability gap [22] phenomenon is seen in Fig. 4, Panel A. Similar to Fig. 2, the size of the jump in forgetting cost (Panel A, top) keeps on increasing indicating that, for each subsequent change in the tasks, the cost keeps growing and subsequently capacity becomes poor, i.e. the model is not capable of representing the underlying distribution properly. This was our main contention in Theorems 3 and 2. Although, deteriorating capacity was easier to observe in the synthetic dataset, we show that, for a real world benchmark CL problem, the theoretical results are indeed valid. Moreover in Fig. 4, Panel B, we observe the comparison with and without the regularization. We do not observe any significant impact of regularization on training behavior. However, even with regularization, the forgetting cost indeed grows and the model becomes unusable, a result clearly observed from Theorem 2.

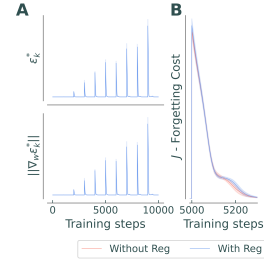


Figure 4: Panel A: (top) Capacity and (bottom) the gradient of capacity; Panel B: Forgetting cost; as a function of training instances.

Case Study 3: Graph NN

Setup: We generate a total of 10 tasks using the PyTorch geometric library [16] with each task comprising of 4 randomly sampled classes from a 10-class classification problem. The key feature of this synthetic data is that both the node and edge features change. We serially feed these tasks to the graph neural network and train the model for a total of 500 epochs.

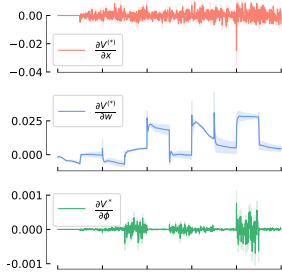


Figure 5: Effect of graph data on the weight updates

Analysis: In this study, we again show that, the larger the distribution shift across subsequent tasks, by virtue of change in the node features $\frac{\partial V^{(*)}}{\partial x}$ and edge features $\frac{\partial V^{(*)}}{\partial \theta}$, the larger the model has to adapt to this shift. This adaptation is reflected in large changes in model weights $\frac{\partial V^{(*)}}{\partial w}$. We observe this in Fig. 5, where the norm of the gradient used to change the tasks increases with each subsequent tasks. More specifically, where there is a large spike in the edge or node features (observed around 4000 step), there is a large update in the weights as well. Moreover, the size of the jump corresponding to weight updates increases with subsequent tasks.

Case Study 4: Transformer-based Large Language Models (8M and 134M parameters)

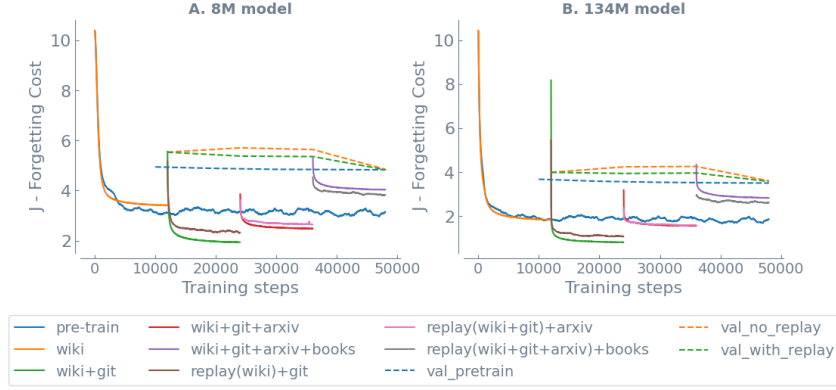


Figure 6: CL on language models demonstrate that forgetting cost increases as new tasks arrive both with and without ER. As expected, the 134M model has higher effective capacity than the 8M model.

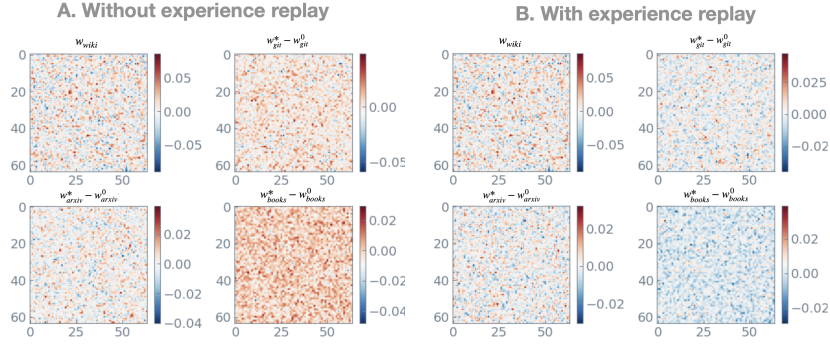


Figure 7: For a task k , the 64×64 heat map shows the difference in weights from the initial value, \mathbf{w}_k^0 , at the start of training to the final value, \mathbf{w}_k^* , at the end of CL training. The weights are randomly sampled from the MLP sublayers in the 8M parameter model. Task arrival order: wiki \rightarrow git \rightarrow arxiv \rightarrow books.

Setup: We utilize four sub-datasets (wiki \rightarrow git \rightarrow arxiv \rightarrow books) from the RedPajama 1T tokens dataset [8] for both pre- and continual pre-training. We use the LLama2 tokenizer [48] and decoder model architecture [48] to construct models with 8M and 134M parameters (details in Appendix). Pre-training was done with a batch size of approx. 4M tokens for 48K steps (about 200B tokens), and a 2K-step linear warmup. For CL, we conduct two experiments: one without ER, using data from only the current task, and another with ER, mixing 80% current task data with 20% from previous tasks (details on data mix in Appendix). Each task is trained for 12K steps (about 50B tokens), starting each new task from the previous task’s final checkpoint. Validation scores are computed on the C4-en validation set [45] using the final checkpoint for each task. We use identical hyper-parameter settings for both models and leverage PyTorch FSDP [55] on 64 A10 (40GB) GPUs.

Analysis: We compare the forgetting costs for continual pre-training with and without ER of 8M and 134M parameter models in Fig. 6. Pre-training costs for both models is also shown for reference.

8M model: Without ER, we see that the forgetting cost initially goes down for the second task (git) but then keeps increasing with the arrival of each new task (arxiv followed by books). This is an expected baseline result [47] and indicates forgetting. Even with ER, we observe an increase in forgetting cost as new tasks arrive. This is a consequence of Theorem 2, as the model needs to learn concepts from a mix of data from multiple tasks. The only exception occurs for the books task, where the cost observed with ER is lower than without ER. We attribute this to initialization bias (of the solution from the previous task) This can also be inferred from Theorem 3, where more similarity in task leads to better learning- an effect that has been shown theoretically in [36]. The forgetting cost on the validation set is lower with ER than without it due to improved generalization.

For reference, we add the pre-training cost curve where all tasks are available together. Initially, the learning objectives (both with and without ER) are relatively easier and therefore forgetting cost is

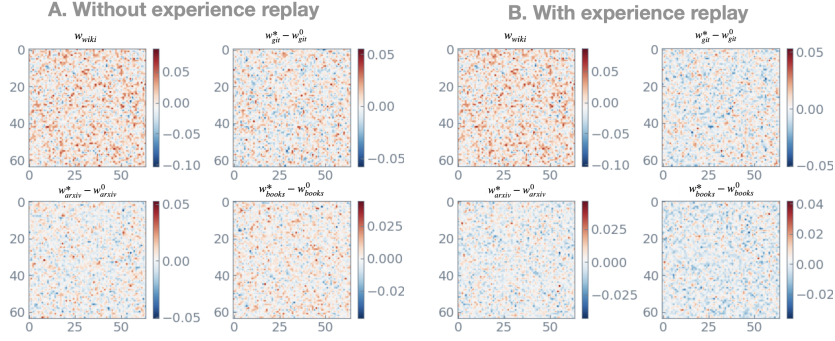


Figure 8: For a task k , the 64×64 heat map shows the difference in weights from the initial value, \mathbf{w}_k^0 , at the start of training to the final value, \mathbf{w}_k^* , at the end of CL training. The weights are randomly sampled from the MLP sublayers in the 134M parameter model. Task arrival order: wiki \rightarrow git \rightarrow arxiv \rightarrow books.

lower than the pre-training cost. However, as more tasks arrive the forgetting cost eventually becomes higher than the pre-training cost because the models keeps on forgetting even with ER (Theorem 2). The validation cost for pre-training model is always lower than both with and without ER indicating that the pre-trained model forgets less than the continual pre-trained model and generalizing better.

134M model: We observe very similar behavior as the 8M model, with an increase in forgetting cost as new tasks arrive. However, owing to larger scale, the forgetting costs are lower indicating a higher effective capacity compared to the 8M model. This is an expected result, as a larger model is more resilient to small changes in the tasks as there are more number of parameters to help with adaptation.

Case Study 5: Visualization for deeper understanding of the impact of CL on the LLM models

Setup: We randomly sampled 64×64 parameters (2% of the MLP parameters for 8M and 0.007% for 134M) and tracked how their weights changed from the start (\mathbf{w}_k^0) to the end of training (\mathbf{w}_k^*) for each task k . We then correlated this with the forgetting cost (Fig. 6) which measures the weight changes caused by each task (the second term in (LB)). Note that, for this example, the last checkpoint from one task serves as the starting point for the next, i.e., $\mathbf{w}_k^0 = \mathbf{w}_{k-1}^*$. Although only a small sample of weights was used, repeated trials showed consistent trends.

Analysis: For the 8M model without ER, large weight changes (red in Fig. 7(A)) lead to capacity loss and increased forgetting. In the `arxiv` task, smaller changes (blue/red) show less learning and more forgetting correlating to the two terms in Lemma 1 where we quantify, how weight and task changes affect the balance point. Significant weight changes occur for the `git` task which effect the second term in Lemma 1 to balance generalization and forgetting by extension reducing forgetting (Fig. 7(A)). In contrast, with ER (Fig. 7(B)), weight changes between tasks are more controlled (more blue than red), reflecting how the two terms in Lemma 1 balance each other. For the `books` task, weight changes are minimal (more blue), indicating marginal model adjustment, higher forgetting, and poorer capacity because the first term no longer balances the second (as shown in Theorem 2).

For the 134M model, we observe similar trends in weight updates. Without ER (Fig. 8(A)), initial changes are slightly larger and continue to increase with each subsequent task. As with the 8M model, increased forgetting costs and significant parameter changes indicate that capacity limits the model’s representation capability. On the other hand, with ER (Fig. 8(B)), weight changes are more regularized (more blue than red) as prior tasks reduce the amount of change in the capacity.

5 CONCLUSION

We studied capacity in continual learning, focusing on the interplay between the model, tasks, optimization procedure, and their impact on the balance point. We introduce CL’s effective model capacity (CLEMC) and find that changes in CLEMC depend on the importance of each task, the cumulative weight changes at each task onset, and the cumulative task changes due to data distribution shifts. Our main conclusion is that even if each subsequent task is only slightly different from the previous one, the effective capacity eventually becomes small, rendering the model unusable.

REFERENCES

- [1] Step-size Optimization for Continual Learning.
- [2] Sayantan Auddy, Sebastian Bergner, and Justus Piater. Effect of Optimizer, Initializer, and Architecture of Hypernetworks on Continual Learning from Demonstration.
- [3] Shawn Beaulieu, Lapo Frati, Thomas Miconi, Joel Lehman, Kenneth O Stanley, Jeff Clune, and Nick Cheney. Learning to continually learn. In *ECAI 2020*, pages 992–1001. IOS Press, 2020.
- [4] Lorenzo Bonicelli, Matteo Boschini, Emanuele Frascaroli, Angelo Porrello, Matteo Pennisi, Giovanni Bellitto, Simone Palazzo, Concetto Spampinato, and Simone Calderara. On the Effectiveness of Equivariant Regularization for Robust Online Continual Learning.
- [5] Massimo Cencini and Francesco Ginelli. Lyapunov analysis: from dynamical systems theory to applications. *Journal of Physics A: Mathematical and Theoretical*, 46(25):250301, 2013.
- [6] Qi Chen, Changjian Shui, Ligong Han, and Mario Marchand. On the Stability-Plasticity Dilemma in Continual Meta-Learning: Theory and Algorithm.
- [7] Aristotelis Chrysakis and Marie-Francine Moens. Online Continual Learning from Imbalanced Data. In *Proceedings of the 37th International Conference on Machine Learning*, pages 1952–1961. PMLR.
- [8] Together Computer. Redpajama: An open source recipe to reproduce llama training dataset, 2023.
- [9] Nicholas Cummins, Brad Killen, Somayeh Bakhtiari Ramezani, Shahram Rahimi, Maria Seale, and Sudip Mittal. A Comparative Study of Continual, Lifelong, and Online Supervised Learning Libraries. 36.
- [10] Zachary A. Daniels, Jun Hu, Michael Lomnitz, Phil Miller, Aswin Raghavan, Joe Zhang, Michael Piacentino, and David Zhang. Efficient Model Adaptation for Continual Learning at the Edge.
- [11] Kamil Deja, Bartosz Cywiński, Jan Rybarczyk, and Tomasz Trzciński. Adapt & Align: Continual Learning with Generative Models Latent Space Alignment.
- [12] Shibhansh Dohare, Juan Hernandez-Garcia, Parash Rahman, Richard Sutton, and A. Rupam Mahmood. Loss of Plasticity in Deep Continual Learning.
- [13] Itay Evron, Edward Moroshko, Gon Buzaglo, Maroun Khriesh, Badea Marjeh, Nathan Srebro, and Daniel Soudry. Continual learning in linear classification on separable data. In *International Conference on Machine Learning*, pages 9440–9484. PMLR, 2023.
- [14] Yan Fan, Yu Wang, Pengfei Zhu, and Qinghua Hu. Dynamic Sub-graph Distillation for Robust Semi-supervised Continual Learning.
- [15] Alberto Fernández, Salvador Garcia, Francisco Herrera, and Nitesh V Chawla. Smote for learning from imbalanced data: progress and challenges, marking the 15-year anniversary. *Journal of artificial intelligence research*, 61:863–905, 2018.
- [16] Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric. *arXiv preprint arXiv:1903.02428*, 2019.
- [17] Robert M French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135, 1999.
- [18] Daniel Goldfarb, Itay Evron, Nir Weinberger, Daniel Soudry, and PAul HAnd. The joint effect of task similarity and overparameterization on catastrophic forgetting — an analytical model. In *The Twelfth International Conference on Learning Representations*, 2024.
- [19] Daniel Goldfarb and Paul Hand. Analysis of catastrophic forgetting for random orthogonal transformation tasks in the overparameterized regime. In *International Conference on Artificial Intelligence and Statistics*, pages 2975–2993. PMLR, 2023.

-
- [20] Bing Han, Feifei Zhao, Wenxuan Pan, Zhaoya Zhao, Xianqi Li, Qingqun Kong, and Yi Zeng. Adaptive Reorganization of Neural Pathways for Continual Learning with Spiking Neural Networks.
- [21] Md Yousuf Harun, Jhair Gallardo, and Christopher Kanan. GRASP: A Rehearsal Policy for Efficient Online Continual Learning.
- [22] Md Yousuf Harun and Christopher Kanan. Overcoming the Stability Gap in Continual Learning.
- [23] Khurram Javed and Martha White. Meta-learning representations for continual learning. *Advances in neural information processing systems*, 32, 2019.
- [24] Gyuhak Kim, Changnan Xiao, Tatsuya Konishi, Zixuan Ke, and Bing Liu. Open-world continual learning: Unifying novelty detection and continual learning. *arXiv preprint arXiv:2304.10038*, 2023.
- [25] Gyuhak Kim, Changnan Xiao, Tatsuya Konishi, and Bing Liu. Learnability and algorithm for continual learning. In *International Conference on Machine Learning*, pages 16877–16896. PMLR, 2023.
- [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.
- [27] Yajing Kong, Liu Liu, Huanhuan Chen, Janusz Kacprzyk, and Dacheng Tao. Overcoming Catastrophic Forgetting in Continual Learning by Exploring Eigenvalues of Hessian Matrix. pages 1–15.
- [28] Yajing Kong, Liu Liu, Maoying Qiao, Zhen Wang, and Dacheng Tao. Trust-Region Adaptive Frequency for Online Continual Learning. 131(7):1825–1839.
- [29] Saurabh Kumar, Henrik Marklund, and Benjamin Van Roy. Maintaining Plasticity in Continual Learning via Regenerative Regularization.
- [30] Qingfeng Lan and A. Rupam Mahmood. Elephant Neural Networks: Born to Be a Continual Learner.
- [31] Qingfeng Lan and A Rupam Mahmood. Elephant neural networks: Born to be a continual learner. *arXiv preprint arXiv:2310.01365*, 2023.
- [32] Matthias De Lange, Guido M van de Ven, and Tinne Tuytelaars. Continual evaluation for lifelong learning: Identifying the stability gap. In *The Eleventh International Conference on Learning Representations*, 2023.
- [33] Dale A Lawrence, William A Sethares, and Wei Ren. Parameter drift instability in disturbance-free adaptive systems. *IEEE transactions on automatic control*, 38(4):584–587, 1993.
- [34] Sebastian Lee, Sebastian Goldt, and Andrew Saxe. Continual learning in the teacher-student setup: Impact of task similarity. In *International Conference on Machine Learning*, pages 6109–6119. PMLR, 2021.
- [35] John Lees-Miller, John Hammersley, and R Wilson. Theoretical maximum capacity as benchmark for empty vehicle redistribution in personal rapid transit. *Transportation Research Record: Journal of the Transportation Research Board*, (2146):76–83, 2010.
- [36] Sen Lin, Peizhong Ju, Yingbin Liang, and Ness Shroff. Theory on forgetting and generalization of continual learning. In *International Conference on Machine Learning*, pages 21078–21100. PMLR, 2023.
- [37] Vincenzo Lomonaco, Lorenzo Pellegrini, Andrea Cossu, Antonio Carta, Gabriele Graffieti, Tyler L Hayes, Matthias De Lange, Marc Masana, Jary Pomponi, Guido M Van de Ven, et al. Avalanche: an end-to-end library for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3600–3610, 2021.

-
- [38] Yun Luo, Zhen Yang, Xuefeng Bai, Fandong Meng, Jie Zhou, and Yue Zhang. Investigating Forgetting in Pre-Trained Representations Through Continual Learning.
- [39] Xiaoyue Mi, Fan Tang, Zonghan Yang, Danding Wang, Juan Cao, Peng Li, and Yang Liu. Adversarial Robust Memory-Based Continual Learner.
- [40] Seyed Iman Mirzadeh, Arslan Chaudhry, Dong Yin, Huiyi Hu, Razvan Pascanu, Dilan Gorur, and Mehrdad Farajtabar. Wide neural networks forget less catastrophically. In *International conference on machine learning*, pages 15699–15717. PMLR, 2022.
- [41] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Razvan Pascanu, and Hassan Ghasemzadeh. Understanding the role of training regimes in continual learning. *Advances in Neural Information Processing Systems*, 33:7308–7320, 2020.
- [42] Preetum Nakkiran, Gal Kaplun, Yamini Bansal, Tristan Yang, Boaz Barak, and Ilya Sutskever. Deep double descent: Where bigger models and more data hurt. *Journal of Statistical Mechanics: Theory and Experiment*, 2021(12):124003, 2021.
- [43] Partha Niyogi and Federico Girosi. On the relationship between generalization error, hypothesis complexity, and sample complexity for radial basis functions. *Neural Computation*, 8(4):819–842, 1996.
- [44] Sahil Nokhwal and Nirman Kumar. RTRA: Rapid Training of Regularization-based Approaches in Continual Learning.
- [45] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv e-prints*, 2019.
- [46] Krishnan Raghavan and Prasanna Balaprakash. Formalizing the generalization-forgetting trade-off in continual learning. *Advances in Neural Information Processing Systems*, 34:17284–17297, 2021.
- [47] Vinay Venkatesh Ramasesh, Aitor Lewkowycz, and Ethan Dyer. Effect of scale on catastrophic forgetting in neural networks. In *International Conference on Learning Representations*, 2021.
- [48] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models, 2023.
- [49] Lam Tran Tung, Viet Nguyen Van, Phi Nguyen Hoang, and Khoat Than. Sharpness and Gradient Aware Minimization for Memory-based Continual Learning. In *Proceedings of the 12th International Symposium on Information and Communication Technology*, pages 189–196. ACM, 2024.
- [50] Liyuan Wang, Xingxing Zhang, Hang Su, and Jun Zhu. A comprehensive survey of continual learning: Theory, method and application. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [51] Maorong Wang, Nicolas Michel, Ling Xiao, and Toshihiko Yamasaki. Improving Plasticity in Online Continual Learning via Collaborative Learning.

-
- [52] Fei Ye and Adrian G. Bors. Wasserstein Expansible Variational Autoencoder for Discriminative and Generative Continual Learning. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 18619–18629. IEEE.
- [53] Min Zeng, Wei Xue, Qifeng Liu, and Yike Guo. Continual Learning with Dirichlet Generative-based Rehearsal.
- [54] Shilin Zhang and Jiahui Wang. Density Distribution-based Learning Framework for Addressing Online Continual Learning Challenges.
- [55] Yanli Zhao, Andrew Gu, Rohan Varma, Liang Luo, Chien-Chin Huang, Min Xu, Less Wright, Hamid Shojanazeri, Myle Ott, Sam Shleifer, Alban Desmaison, Can Balioglu, Pritam Damania, Bernard Nguyen, Geeta Chauhan, Yuchen Hao, Ajit Mathews, and Shen Li. Pytorch fsdp: Experiences on scaling fully sharded data parallel. *Proc. VLDB Endow.*, 16(12):3848–3860, aug 2023.
- [56] Yuqing Zhao, Divya Saxena, and Jiannong Cao. AdaptCL: Adaptive Continual Learning for Tackling Heterogeneity in Sequential Datasets. pages 1–14.