

Provable Lifelong Learning of Representations

Xinyuan Cao
Georgia Tech
xcao78@gatech.edu

Weiyang Liu
University of Cambridge & MPI-IS
wl396@cam.ac.uk

Santosh S. Vempala
Georgia Tech
vempala@gatech.edu

Abstract

In lifelong learning, tasks (or classes) to be learned arrive sequentially over time in arbitrary order. During training, knowledge from previous tasks can be captured and transferred to subsequent ones to improve sample efficiency. We consider the setting where all target tasks can be represented in the span of a small number of unknown linear or nonlinear features of the input data. We propose a lifelong learning algorithm that maintains and refines the internal feature representation. We prove that for any desired accuracy on all tasks, the dimension of the representation remains close to that of the underlying representation. The resulting sample complexity improves significantly on existing bounds. In the setting of linear features, our algorithm is provably efficient and the sample complexity for input dimension d , m tasks with k features up to error ϵ is $\tilde{O}(dk^{1.5}/\epsilon + km/\epsilon)$. We also prove a matching lower bound for any lifelong learning algorithm that uses a single task learner as a black box. We complement our analysis with an empirical study, including a heuristic lifelong learning algorithm for deep neural networks. Our method performs favorably on challenging realistic image datasets compared to state-of-the-art continual learning methods.

1 Introduction

Recent years have witnessed significant advances in both theory and practice of supervised learning. While a variety of techniques are available for learning individual target functions, much less is known about *continual* or *lifelong* learning, where the learner is adding new target functions to their repertoire. Inspired by how humans learn and transfer knowledge during their lifespan, lifelong learning has many applications in computer vision [PKP⁺19] and robotics [TM95].

A central idea for lifelong learning is to learn an *efficient representation* that facilitates the collection of target functions to be learned. For example, if deep feed-forward networks are being used for classification, the goal might be to learn a hidden layer whose outputs are relevant and useful features for the family of tasks. Building a classifier on top of them is relatively easy or less expensive than building one from the original input features. This representation itself is incrementally refined as more target functions are learned.

We consider a very general setting of task/class incremental learning, where new samples from different tasks/classes are presented sequentially over time. The goal of the learner is to maintain hypothesis functions that work for all tasks/classes encountered so far. We assume that all targets are simple functions of a bounded number of unknown *linear or nonlinear* features.

Prior work [BBV15] considered the task-incremental setting where the target functions are linear classifiers of the input that all lie in a common low-dimensional subspace. Under this assumption, a simple algorithm can be shown to learn a good representation of size comparable to the optimal one (*i.e.*, a basis of the common low-dimensional subspace). The algorithm proceeds as follows: maintain a small number of linear features; learn the next function as a linear function of the features; if

the error is too high, learn the new function directly on the input, and add it as a new feature. Under mild assumptions on the input distribution (log-concavity), with a suitable choice of error parameters, this algorithm is guaranteed to learn a small set of features that work well for all the target functions. More recent works [DHK⁺20, TJJ20, CLL21] focus on the sample complexity of multi-task learning under strong distributional assumptions on both the data and the tasks.

Our paper is motivated by the following questions:

- Can the theoretical guarantees for linear features be extended to a representation with nonlinear features?
- Does the refinement of the internal representation have provable benefits?
- What is the best possible sample complexity of lifelong learning?

Our work addresses these questions for both task-incremental learning (classification or regression) and class-incremental learning (where we do not have access to the task ID). Our analysis applies to a broad class of lifelong learning algorithms that dynamically change the network architecture. First, we analyze the setting where the underlying common features are nonlinear, which is considerably more general than previously considered. We prove that this natural lifelong learning algorithm is guaranteed to learn low-error targets creating only a small number of nonlinear features. Secondly, we propose a new algorithm, with a refinement step, and show that it improves the sample complexity using a new perspective on feature subspaces. The resulting sample complexity improves significantly on known bounds for the setting of linear features, and perhaps surprisingly, we show that it is the best possible in the setting of linear features, assuming that the lifelong learner has black-box access to a single task learner to any desired level of accuracy. We do this by constructing a hard distribution over tasks.

We conduct experiments on class-incremental learning using benchmark data sets and find that our proposed algorithm outperforms state-of-the-art continual learning algorithms.

1.1 Problem Settings

We consider m tasks (or m -class classification) where the tasks (classes) arrive sequentially over time. Let $X = \mathbb{R}^d$ be the input space and Y be the label space. We study a discriminative model, where the target function of each task can be learned using a linear combination of at most k linear/nonlinear features. The goal is to learn a hypothesis function with small generalization errors on all tasks.

Formally, the problem is associated with a distribution P over $X \times Y$, D is the marginal of P over X . The label for an input data point $\mathbf{x} \in \mathbb{R}^d$ is given by

$$\ell(\mathbf{x}) = \phi(\langle \mathbf{c}^*, \boldsymbol{\sigma}^*(\mathbf{x}) \rangle)$$

where $\boldsymbol{\sigma}^*(\mathbf{x}) = (\sigma_1^*(\mathbf{x}), \dots, \sigma_k^*(\mathbf{x}))^\top \in \mathbb{R}^k$ is a vector of unknown features, $\mathbf{c}^* \in \mathbb{R}^k$, $\phi(\cdot) : \mathbb{R} \rightarrow Y$ is the map to the label space. ($k \ll \min(m, d)$). Equivalently, we can view it as a two-layer network with k neurons in the hidden layer (Figure 1.1).

Our goal is to learn a good hypothesis function $\hat{\ell}(\cdot)$ parameterized by $(\mathbf{c}^*, \boldsymbol{\sigma}^*)$ with a small generalization error $err = \mathbb{P}_{(\mathbf{x}, y) \sim P} L(l(\mathbf{x}), \hat{l}(\mathbf{x}))$, where $L(\cdot, \cdot)$ is some loss function for the specific task.

We use a similar model with multi-task learning, where all tasks share the same low-dimensional feature subspace. However, it is different from multi-task learning, which has T_1 source tasks to learn all-at-once and use the features learned to solve the target tasks. The assumption is made

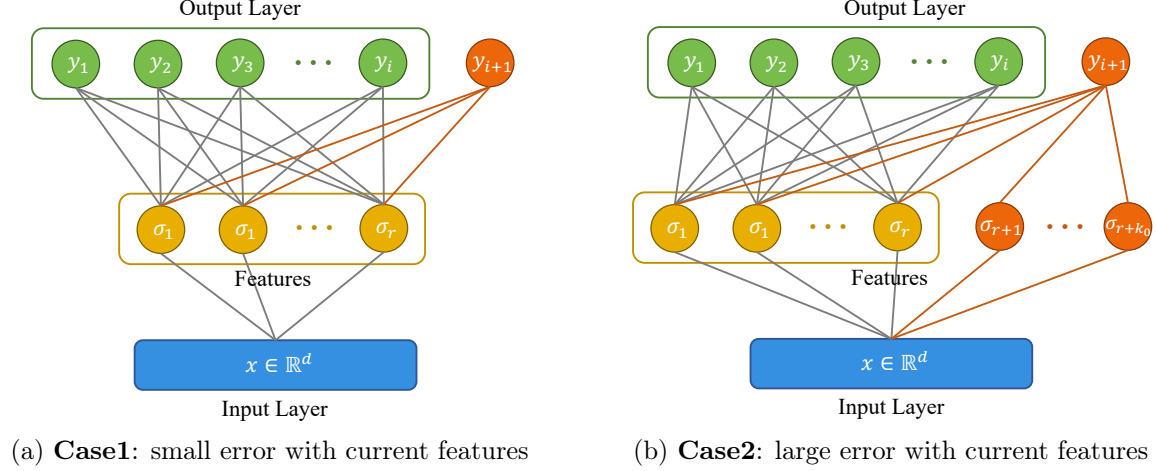


Figure 1.1: An illustration of LLL. Given a new task y_{i+1} , the algorithm tries to learn the task with existing features $\sigma_1, \dots, \sigma_r$. If the error is small (case1, Sub-figure (a)), then it moves to the next task. Otherwise (case2, Sub-figure (b)), it learns a new set of features $\sigma_{r+1}, \dots, \sigma_{r+k_0}$ and a linear combination of all features; for linear features, it learns a single new feature σ_{r+1} .

there that the features of the target task are covered by all features that have been learned. Instead, lifelong learning algorithms learn all tasks sequentially, with no prior knowledge of the incoming tasks during training.

Here we focus on the task-incremental learning of binary classification tasks. Extensions to task-incremental learning of linear regression and multi-class classification tasks are given in Appendix B.

Let $\mathbf{X} = \mathbb{R}^d$ be the input space, $\mathbf{Y} = \{\pm 1\}$ be the label space. For any task $i \in [m]$, any sample (\mathbf{x}, y) drawn from P satisfies $y = l_i(\mathbf{x}) = \text{sign}(\langle \mathbf{c}_i^*, \boldsymbol{\sigma}^*(\mathbf{x}) \rangle)$, where the features are $\boldsymbol{\sigma}^*(\mathbf{x}) = \mathbf{W}^* \mathbf{x}$ in the linear case and $\boldsymbol{\sigma}^*(\mathbf{x}) = f(\mathbf{W}^* \mathbf{x})$ in the nonlinear case. Here $f(\cdot)$ is a nonlinear activation function, *e.g.* ReLU. $\mathbf{W}^* \in \mathbb{R}^{k \times d}$, $\mathbf{c}_i^* \in \mathbb{R}^k$ etc. Specifically, in the linear case, for each task i , we equivalently have $y = \text{sign}(\langle \mathbf{a}_i^*, \mathbf{x} \rangle)$, where $\mathbf{a}_i^* = \mathbf{W}^{*\top} \mathbf{c}_i^* \in \mathbb{R}^d$. WLOG we assume that each \mathbf{a}_i^* is a unit vector, *i.e.*, $\|\mathbf{a}_i^*\|_2 = 1, \forall i \in [m]$. The generalization error is defined as $err = \mathbb{P}_{(\mathbf{x}, y) \sim P}(l_i(\mathbf{x}) \neq \hat{l}_i(\mathbf{x}))$.

1.2 Main Results

In all the results and analysis, we only have Assumption 1, which as Lemma 1 from [BBV15]) asserts, is satisfied by all log-concave distributions after an affine transformation. This class includes many common distributions, such as Gaussian, Uniform and Gamma distributions [LV07].

Assumption 1. (Data Distribution Assumption) Let $\theta(\cdot, \cdot)$ denote the angle between two vectors. We assume that there exist universal constants $c_2 > c_1 > 0$ s.t., for any unit vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^d$,

$$c_1 \theta(\mathbf{u}, \mathbf{v}) \leq \mathbb{P}_{\mathbf{x} \sim D}(\text{sign}(\mathbf{u} \cdot \mathbf{x}) \neq \text{sign}(\mathbf{v} \cdot \mathbf{x})) \leq c_2 \theta(\mathbf{u}, \mathbf{v})$$

Our theoretical upper bounds are summarized below. The detailed statements for these results appear as Theorem 3, Theorem 4 and Theorem 5 in Section 3. These results are based on the algorithms described in Section 2, called Basic Lifelong Learning (LLL) and Lifelong Learning with Representation Refinement (LLL-RR).

Theorem 1 (Summary of Upper Bounds). *Consider the lifelong learning setting of input dimension d , m tasks with k common features. The basic lifelong learning algorithm achieves a target error of ϵ on all tasks with sample complexity $\tilde{O}(dk^{1.5}/\epsilon + km/\epsilon)$ for linear features and a factor of k higher for nonlinear features. With representation refinement using at most $2k$ features, the sample complexity is $\tilde{O}(dk^{1.5}/\epsilon + km/\epsilon)$. In the linear setting refinement runs in polynomial-time.*

This raises the question of whether there exist algorithms with better sample complexity. We show that the answer is NO in a general sense. We assume that we have black-box access to a single-task learner that works as follows: it takes as input labeled examples and a target accuracy ϵ , and outputs some feasible solution with error at most ϵ . Then we show that any lifelong learning algorithm that achieves ϵ error for all tasks needs $\Omega(dk^{1.5}/\epsilon + km/\epsilon)$ samples.

Theorem 2 (Lower Bound). *Suppose that a lifelong learner has black-box access to a single task learner that takes an error parameter ϵ as input and is allowed to return any vector that is within distance ϵ of the true target unit vector, using $\Theta(d/\epsilon)$ samples in \mathbb{R}^d . Then, there exists a distribution of m tasks, $m = 2^{\Theta(k)}$ such that for any lifelong learning algorithm, WHP, the total number of samples required to learn all m tasks up to error ϵ is $\Omega(dk^{1.5}/\epsilon + km/\epsilon)$.*

Our contributions can be summarized as follows:

Sample complexity. We bound the sample complexity of lifelong learning for both the linear and nonlinear cases. In the linear case, our bound for the lifelong learning is $\tilde{O}(dk^{1.5}/\epsilon + km/\epsilon)$. This improves the dependence on both k and ϵ compared to past work [BBV15], which proved a bound of $\tilde{O}(dk^2/\epsilon^2 + km/\epsilon)$. It also improves existing theoretical results for multi-task learning [DHK⁺20, TJJ20], where the best sample complexity is $\tilde{O}(dk^2/\epsilon + km/\epsilon)$. Moreover, this bound is the best possible up to logarithmic factor for any lifelong learning algorithm.

Representation refinement. We propose and analyze the step of sample-free representation refinement in the lifelong learning setting. Specifically, this step aims to reduce the dimension of the feature subspace while keeping the subspace close to the true one. In the linear setting, we provide an algorithmically efficient approach via an SDP relaxation. To the best of our knowledge, we provide the first provable bound for representation refinement.

Proof techniques. Our analysis is based on geometric insights. The test error translates to the distance between the target and learned vectors. To show that our learned feature subspace is close to the true one, we consider the set of candidate k -dimensional subspaces. We would like to show that the measure of this set decreases rapidly during learning. Instead, we identify the set of well-approximated vectors by our current learned subspace and show that the set grows at a geometric rate until it includes all vectors in the true subspace.

Empirical results. We evaluate our lifelong learning algorithms on standard benchmarks and compare them with state-of-the-art methods, demonstrating their practice efficiency. We also perform simulations for the setting of linear features, and exhibit results that match our theoretical bounds.

1.3 Related work

Lifelong learning [TM95] aims to solve different tasks arriving in a stream, where knowledge from current and previous tasks is re-used in subsequent tasks to improve efficiency and sample complexity. Early works found that lifelong learning can encounter Catastrophic Forgetting (CF) [MC89], especially when using back-propagation [Rat90]. That is, the performance on old tasks can drop dramatically after learning a new task. There are three main approaches to addressing this problem: *adding a regularization term* [LH17, KPR⁺17], *freezing the network from previous tasks and adding branches to new tasks* [XZ18, RRD⁺16, LLRS19, LLL⁺21, YYLH17] and *replaying previous tasks' exemplars* [RKSL17]. Our work is closest to the second approach in that we dynamically change the architecture to overcome CF. Although we do not know the number of tasks in advance, we prove that our algorithm has a small model size and efficient sample complexity.

Despite a vast literature on lifelong learning methods, theoretical investigations are relatively few. [YFL⁺20] studies the optimization and generalization properties of the regularization-based method by analyzing the loss landscape. [BDS20, DBM⁺21] analyze the generalization of the OGD algorithm [FAML20] through NTK [JGH18]. [BBV15] gives an upper bound on the architecture size when we grow the network when training binary classifiers. We improve their bounds by getting nearly tight sample complexity in the linear case and generalize the approach to the nonlinear regime.

Two topics closely related to lifelong learning are meta-learning and transfer learning. There is a line of work where all tasks approximately [FAL17, KBT19, BKT19, DSCP19] or conditionally [WDC20, DPC20, DPC21] share a common representation. However, our work focuses on the setting where all tasks share one common low-dimensional representation. [CLL21] shows the benefits of task-specific fine-tuning, which is fundamentally different from our refinement step. Our refinement step aims to reduce the representation dimension with slight information loss and help to improve the sample complexity of subsequent tasks. This procedure needs no additional data.

There are other works with similar settings to ours where all tasks share one representation. [Bax97] bounds the sample complexity to achieve low average error from a Bayesian/information-theoretic point of view. We compare our results with recent work [DHK⁺20, TJJ20, BBV15] in Table 1 in the linear case. Previous works on multi-task learning [DHK⁺20, TJJ20] need $\tilde{O}(dk^2/\epsilon)$ or more samples from previous tasks to learn the hidden features, while our algorithm needs $\tilde{O}(dk^{1.5}/\epsilon)$ samples. After that, each new task can be learned up to ϵ error with $\tilde{O}(k/\epsilon)$ samples. These results illustrate the efficiency of lifelong learning compared to all-at-once training. Our analysis generalizes to the setting of nonlinear features, *e.g.*, if labels are generated by a two-layer neural network, with k hidden units. We prove our lifelong learning algorithm is sample efficient with only Assumption 1, which is minimal compared to related work.

Notation. We use bold upper-case letters to refer to matrices (*e.g.* \mathbf{X}) and bold lower-case letters to refer to vectors (*e.g.* \mathbf{x}). We use $[m] = \{1, 2, \dots, m\}$. We use \tilde{O} to hide polylogarithmic factors. O, Ω, Θ are standard notations for order of growth. For any two vectors \mathbf{x}, \mathbf{y} , let $\theta(\mathbf{x}, \mathbf{y})$ be the angle between them. The angle between a vector \mathbf{x} and a subspace \mathbf{U} is defined as $\theta(\mathbf{x}, \mathbf{U}) = \min_{\mathbf{u} \in \mathbf{U}} \theta(\mathbf{x}, \mathbf{u})$. For two subspaces \mathbf{U}, \mathbf{V} , define $\theta(\mathbf{U}, \mathbf{V}) = \max_{\mathbf{u} \in \mathbf{U}} \theta(\mathbf{u}, \mathbf{V})$. Thus $\theta(\mathbf{U}, \mathbf{V}) \leq \alpha$ iff for all $\mathbf{u} \in \mathbf{U}, \exists \mathbf{v} \in \mathbf{V}$ s.t. $\theta(\mathbf{u}, \mathbf{v}) \leq \alpha$. We define the distance from a vector \mathbf{u} to a subspace \mathbf{F} as the orthogonal distance: $d(\mathbf{u}, \mathbf{F}) = \min_{\mathbf{v} \in \mathbf{F}} \|\mathbf{u} - \mathbf{v}\|_2$. For a distribution D and two vectors \mathbf{u}, \mathbf{v} , we define $d_D(\mathbf{u}, \mathbf{v}) = \mathbb{P}_{\mathbf{x} \sim D}(\text{sign}(\mathbf{u} \cdot \mathbf{x}) \neq \text{sign}(\mathbf{v} \cdot \mathbf{x}))$.

Method	Assumptions (linear)	Feature Dim	Total Samples
Few-shot learning [DHK ⁺ 20] [†]	• Sub-gaussian input	k	$\tilde{O}(\frac{dk^2}{\epsilon} + \frac{km}{\epsilon})$
Meta learning [TJJ20] [†]	• Sub-gaussian input	k	$\tilde{O}(\frac{dk^2}{\epsilon} + \frac{km}{\epsilon})$
LLL [BBV15] [★]	• Log-concave input	k	$\tilde{O}(\frac{\epsilon}{\epsilon^2} + \frac{km}{\epsilon})$
LLL (our paper) [★]	• Well-spread input (Assumption 1)	$2k \log(\frac{\log k}{\epsilon})$	$\tilde{O}(\frac{dk^{1.5}}{\epsilon} + \frac{km}{\epsilon})$
LLL-RR (our paper) [★]	• Well-spread input (Assumption 1)	$2k$	$\tilde{O}(\frac{dk^{1.5}}{\epsilon} + \frac{km}{\epsilon})$

Table 1: Comparison of different transfer learning algorithms in the linear setting. †: the method trains all source tasks all at once and then uses the representation to train the target tasks. ★: all source tasks and target tasks are learned sequentially.

2 Algorithms

We study three algorithms: the *basic lifelong learning* algorithm (basic LLL) in Section 2.1, *lifelong learning with representation refinement* algorithm (LLL-RR) in Section 2.2, and *heuristic lifelong learning* algorithm (H-LLL) in Section 2.3. We prove guarantees for basic LLL and LLL-RR in Section 3. We show that H-LLL for deep neural networks (Section 2.3) outperforms state-of-the-art continual learning algorithms in Section 5.2.

2.1 Basic Lifelong Learning

Our algorithm maintains a set of features $\sigma_1(\cdot), \dots, \sigma_r(\cdot)$ while tasks are presented incrementally. As is shown in Figure 1.1, when the next task, say $(i+1)$ -th task arrives, the algorithm first tries to learn a new linear combination y_{i+1} of existing features using examples from the current task. If the best such combination has a low error, it records the linear combination parameters and moves on to the next task. If the error is higher than a threshold ϵ , then it learns a new set of features $\sigma_{r+1}, \dots, \sigma_{r+k_0}$ and a new linear combination of them with error up to ϵ_{acc} . Denote \tilde{k} as the number of steps that the algorithm learns new features. Let k_0 be the number of features learned at one time, it is a constant dependent on whether the features are linear or not. We describe the algorithm in Algorithm 1.

Algorithm 1 Basic Lifelong learning Algorithm (Basic LLL)

Input: d, m, k , labeled examples of m tasks, threshold parameters ϵ_{acc}, ϵ .

The algorithm maintains a set of features $\sigma_1(\cdot), \dots, \sigma_r(\cdot)$ along training. When task $i+1$ arrives,

- Use the data from the $(i+1)$ -th task, attempt to learn the linear function $\tilde{\mathbf{c}}_{i+1}$ using the current features $\boldsymbol{\sigma}(\cdot) = (\sigma_1(\cdot), \dots, \sigma_r(\cdot))^\top$.
- Check whether the hypothesis $\mathbf{x} \mapsto \text{sign}(\tilde{\mathbf{c}}_{i+1}^\top \boldsymbol{\sigma}(\mathbf{x}))$ has error less than ϵ .
 1. If yes, record the linear combination parameters $\tilde{\mathbf{c}}_{i+1}$.
 2. Otherwise, learn a new set of features $\boldsymbol{\sigma}'(\cdot) = (\sigma_{r+1}(\cdot), \dots, \sigma_{r+k_0}(\cdot))^\top$ and a linear function $\tilde{\mathbf{c}}_{i+1}$ such that the predictor $\mathbf{x} \mapsto \text{sign}(\tilde{\mathbf{c}}_{i+1}^\top \boldsymbol{\sigma}'(\mathbf{x}))$ has error less than ϵ_{acc} . Update the representation $\boldsymbol{\sigma}(\cdot) = (\sigma_1(\cdot), \dots, \sigma_{r+k_0}(\cdot))^\top$.

return m predictors: $\mathbf{x} \rightarrow \text{sign}(\tilde{\mathbf{c}}_i^\top \boldsymbol{\sigma}(\mathbf{x}))$, where $\boldsymbol{\sigma}(\mathbf{x}) = (\sigma_1(\mathbf{x}), \dots, \sigma_{\tilde{k}k_0}(\mathbf{x}))^\top, 1 \leq i \leq m$.

The algorithm works for both linear and nonlinear features. For linear features, if a new target function does not have a good representation as a combination of the features learned so far, the new

target is itself a new feature since everything is linear (Algorithm 1, [BBV15]), so k_0 above is 1. For nonlinear features, when the current representation is not good enough, we can learn a set of $k_0 \leq k$ nonlinear features with low error since each task corresponds to a target with at most k features. Here we assume that a single such combination can be learned efficiently (*i.e.*, a neural network with a small, single hidden layer) [BHL19]. Section 3 proves that the number of features to be learned can be upper bounded by $\tilde{O}(kk_0)$. We give the full guarantees for this basic LLL algorithm in Theorem 3.

2.2 Lifelong Learning with Representation Refinement

Similar to the basic LLL algorithm, LLL-RR also expands the feature space gradually. Whenever we learn a new task i , we attempt to learn it using the current representation and check whether a linear combination exists with an error less than ϵ . If yes, we record the classifier for the current task and move to the next one. Otherwise, we learn a new classifier for the current task with error at most ϵ_{acc} , via new features; we then do a step of *representation refinement* on all the features learned so far. The refinement step can also be done when the number of features grows above a threshold rather than every time a new task is learned to high accuracy. The formal description of LLL-RR is given in Appendix A.

Refinement algorithm. Denote $\tilde{\mathbf{w}}_1, \dots, \tilde{\mathbf{w}}_{(\hat{k}+1)k_0}$ as all the features learned so far. The goal of refinement is to find a minimal dimensional feature subspace that is within distance ϵ_{acc} to all learned features. We *minimize the dimension* of feature subspace while keeping it close to the original representation by solving the optimization problem (2.1). This problem is NP-hard, but we provide an efficient approximation algorithm for the linear case (and practical implementation for the general case in Section 2.3).

Algorithm 2 Representation Refinement (RR)

Input: All features learned so far $\tilde{\mathbf{w}}_1, \dots, \tilde{\mathbf{w}}_{(\hat{k}+1)k_0}$, and the desired feature subspace dimension k .

Solve the following optimization problem, and get the solution \mathbf{V}' .

$$\begin{aligned} \min_{\mathbf{V}} \quad & \dim(\mathbf{V}) \\ \text{s.t.} \quad & d(\tilde{\mathbf{w}}_i, \mathbf{V}) \leq \epsilon_{acc}, \quad \text{for } 1 \leq i \leq (\hat{k} + 1)k_0 \end{aligned} \tag{2.1}$$

return *Refined representation* \mathbf{V}' .

The refinement step is provably beneficial to the total sample complexity. Theorem 4 guarantees that lifelong learning algorithm with representation refinement (LLL-RR) can be ended with learning new features in $\tilde{O}(k)$ steps. The analysis is shown in Section 3.

Linear features. We provide an efficient implementation for the linear case by using a Semi-definite Programming (SDP) relaxation (2.2) and then applying Principal Component Analysis (PCA) to round the SDP solution. The relaxation from (2.1) to (2.2) is natural. The positive semi-definite (PSD) matrix \mathbf{X} represents the projection matrix to \mathbf{V}^\perp , the complement of the subspace \mathbf{V} . It is a relaxation since \mathbf{X} might have fractional eigenvalues between 0 and 1. We describe the formal algorithm in Algorithm 3. As is proved in Theorem 5 (Section 3), if the optimal dimension of the feature subspace is k , this linear case implementation will output a $(2k - 1)$ -dimensional subspace \mathbf{V}' with $d(\tilde{\mathbf{w}}_i, \mathbf{V}') \leq \sqrt{2}\epsilon_{acc}, \forall i \in [\hat{k} + 1]$. Consequently, LLL-RR terminates with feature dimension $O(k)$.

Algorithm 3 Representation Refinement (RR) Implementation in Linear Case

Input: All features learned so far $\tilde{\mathbf{w}}_1, \dots, \tilde{\mathbf{w}}_{\hat{k}+1}$, and the desired feature subspace dimension k .

1. Solve the following SDP, and get the solution \mathbf{X}^*, t^* .

$$\begin{aligned} \min_{\mathbf{X}, t} \quad & t \\ \text{s.t.} \quad & \tilde{\mathbf{w}}_i^\top \mathbf{X} \tilde{\mathbf{w}}_i \leq t, 1 \leq i \leq \hat{k} + 1 \\ & 0 \preceq \mathbf{X} \preceq I \\ & \text{Tr}(\mathbf{X}) = d - k \end{aligned} \tag{2.2}$$

2. Do the singular value decomposition $\mathbf{X}^* = \sum_{i=1}^d \lambda_i \mathbf{u}_i \mathbf{u}_i^\top$, where $0 \leq \lambda_1 \leq \dots \leq \lambda_d \leq 1$.

return *Refined representation* $\mathbf{V}' = \text{span}(\mathbf{u}_1, \dots, \mathbf{u}_{2k-1})$.

2.3 A Lifelong Learning Heuristic for Deep Neural Networks

In order to apply our basic LLL algorithm to deep neural networks, we propose a heuristic lifelong learning (H-LLL) algorithm. The intuition of our LLL algorithm is to build an expandable and dynamic representation that can adapt to incoming tasks/classes without sacrificing the quality for previous tasks/classes. Following this intuition, we propose to learn a separate encoder for each task. We observe the training data \mathcal{D}_i for the i -th task and the memory buffer \mathcal{M}_i for the previous tasks. The memory buffer is constructed based on herding selection [Wel09, RKSL17]. H-LLL works iteratively in two phases. First, H-LLL learns the representation with a separate encoder f_i in the i -th task, while the other encoders $f_j, j < i$ are frozen during the training in the i -th task. Second, H-LLL finetunes the last classifier layer using the memory buffer \mathcal{M}_i and the current task data \mathcal{D}_i . These two steps are iterated as the training proceeds. We take the i -th task as an example. Since we train a separate encoder f_i for the i -th task, the representation of a sample \mathbf{x} (by the end of the i -th task) is constructed by concatenating all the learned features: $\mathbf{v}_i(\mathbf{x}) = \{f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_i(\mathbf{x})\}$ where \mathbf{v}_i denotes the representation after learning the i -th task. The training uses cross-entropy loss on both the memory buffer \mathcal{M}_i and the current dataset \mathcal{D}_i :

$$\mathcal{L} = -\frac{1}{|\mathcal{M}_i \cup \mathcal{D}_i|} \sum_{i=1}^{|\mathcal{M}_i \cup \mathcal{D}_i|} \log \left(\text{SoftMax} \left(\mathbf{W}_{\text{cls}}^\top \mathbf{v}_i(\mathbf{x}) \right) \right) \tag{2.3}$$

where \mathbf{W}_{cls} is the weight of the last classifier layer. After training of the representation is completed, we follow [YXH21] and re-train the classifier layer with a heated-up softmax [ZYK⁺18] and a balanced finetuning method [CMJG⁺18]. Note that, for each encoder $f_j, \forall j$, we can parameterize it with any neural network. In this paper, we use ResNet-18 as $f_j, \forall j$.

3 Theoretical Guarantees

Here we state the main theorems for the basic LLL algorithm and LLL-RR algorithm, bounding the representation size and complexity. Here our algorithm and analysis apply for both linear and nonlinear features. For nonlinear features, we consider the kernel induced by them. These features live in a potentially infinite-dimensional space (or exponential in d dimensional space if, *e.g.*, the input is from the Boolean hypercube).

The theorem for the basic lifelong learning algorithm is stated as follows.

Theorem 3 (Basic LLL). *Consider the lifelong learning setting of input dimension d , m tasks with k common features. Let $\epsilon_{acc} = \frac{\epsilon}{c\sqrt{k}}$ for a sufficiently small constant $c > 0$. Under Assumption 1, the basic LLL algorithm, learns new features at most $\tilde{k} = O(k \log(\log(k)/\epsilon))$ times and the dimension of the learned feature space is $O(k \log(\log(k)/\epsilon))$ for linear features and $O(k^2 \log(\log(k)/\epsilon))$ for nonlinear features. The total number of labeled examples to learn all tasks to within error ϵ is $O(\frac{dk^{1.5}}{\epsilon} \log(\frac{\log(k)}{\epsilon}) \log(\frac{k}{\epsilon}) + \frac{km}{\epsilon} \log(\frac{\log(k)}{\epsilon}) \log(\frac{1}{\epsilon})) = \tilde{O}(dk^{1.5}/\epsilon + km/\epsilon)$ for linear features and a factor of k higher for nonlinear features.*

Our main result analyzes the lifelong learning algorithm with representation refinement.

Theorem 4 (LLL with Representation Refinement). *Consider the lifelong learning setting of input dimension d , m tasks with k common features. Suppose that the algorithm has access to an oracle that gives a constant-factor approximation of Optimization Problem 2.1. Set $\epsilon_{acc} = \frac{\epsilon}{c\sqrt{k}}$ for a sufficiently small constant $c > 0$. Under Assumption 1, the LLL-RR algorithm learns at most $O(k \log(\log(k)/\epsilon))$ new features, and the dimension of the feature space is $O(k)$. The total number of labeled examples to learn tasks to within error ϵ is $O(\frac{dk^{1.5}}{\epsilon} \log(\frac{\log(k)}{\epsilon}) \log(\frac{k}{\epsilon}) + \frac{km}{\epsilon} \log(\frac{1}{\epsilon})) = \tilde{O}(dk^{1.5}/\epsilon + km/\epsilon)$.*

In the linear setting, we provide an efficient implementation of the constant-factor approximation oracle in Algorithm 3 with the following guarantee.

Theorem 5 (Approximation). *In the linear case, for the optimization problem (2.1), if there exists a subspace \mathbf{V}^* of k dimension with $d(\tilde{\mathbf{w}}_i, \mathbf{V}^*) \leq \epsilon_{acc}, \forall i \in [\hat{k} + 1]$, then for any constant $c > 1$, we can get a $(ck - 1)$ -dimensional subspace solution with approximation factor $\sqrt{1 + \frac{1}{c-1}}$ in maximum distance. Specifically, let $c = 2$, the output of Algorithm 3, \mathbf{V}' , is a $(2k - 1)$ -dimensional subspace s.t. $d(\tilde{\mathbf{w}}_i, \mathbf{V}') \leq \sqrt{2}\epsilon_{acc}, \forall i \in [\hat{k} + 1]$.*

In Appendix C, we give another approach for analyzing the basic LLL algorithm with weaker guarantees, similar to the proof in [BBV15].

3.1 Proof of Guarantees for Basic LLL algorithm and LLL-RR algorithm

We will prove the guarantees for LLL-RR algorithm (Theorem 4) first, and then the proof of Theorem 3 follows easily.

Proof idea and plan. Our proof is based on geometry. Firstly, we show that Assumption 1 guarantees that the distance between hypothesis vectors approximates the test error within a constant factor. As is shown in Figure 3.1, for any target \mathbf{a}_i that has a large error based on the previous features, $d(\mathbf{a}_i, \mathbf{V}_{i-1}) \geq \epsilon$ neglecting constants. Learning \mathbf{a}_i accurately will help reduce the angle between the feature subspace of the algorithm \mathbf{V}_i and the true feature subspace \mathbf{V}^* . To quantify the improvement in the angle between the feature subspaces, we construct a convex set whose volume grows at a geometric rate. This leads to the upper bound on the number of new features.

To be more specific, denote $i_1, \dots, i_{\tilde{k}}$ as the indices of tasks where we learn new features. At step $i_{\tilde{k}}$, we construct a set $Y_{i_{\tilde{k}}}$ of all possible subspaces that are feasible solutions to the refinement optimization problem (2.1). Let $X_{i_{\tilde{k}}}$ be the set of vectors in the unit ball in the true k -dimensional feature subspace that are within distance $O(\epsilon/\sqrt{k})$ to all subspaces in $Y_{i_{\tilde{k}}}$. Then we can show that $X_{i_{\tilde{k}}}$ is a symmetric convex set. Clearly, the set $Y_{i_{\tilde{k}}}$ shrinks during training as we have more and more constraints in the optimization problem. Alongside, the volume of X_i increases exponentially. The learning procedure terminates when $X_{i_{\tilde{k}}}$ covers the ball $B_k(0, 1/2\sqrt{k})$, which means that a target function (unit vector) spanned by the true k features will have error $O(\epsilon)$ to the solution learned by

LLL-RR. In other words, the feature subspace we learn can solve all future tasks with small errors using only hypothesis vectors from the learned feature subspace.




Figure 3.1: Geometric illustration of the proof sketch. For any target \mathbf{a}_i that has more than ϵ error based on the previous feature subspace \mathbf{V}_{i-1} , the algorithm accurately learns a new feature within error ϵ_{acc} , and therefore pushes the new feature subspace \mathbf{V}_i towards the true one \mathbf{V}^* .

We will prove this step by step. Lemma 1 bridges the test error to the distance metric. Lemma 2, Lemma 3 and Corollary 1 show that the convex hull of true feature vectors is contained in the set X_{i_k} . Lemma 4 carefully analyzes the maximum volume ellipsoid in X_{i_k} , whose volume grows exponentially. Based on these facts, we can bound the number of new features and prove Theorem 4.

We will first prove the theorem for the linear setting. Then it can be naturally extended to the nonlinear case when we consider the kernel of the features, regarding each feature as a potentially infinite-dimensional vector.

Let $\mathbf{X} = \mathbb{R}^d$, for each task i , there exists unit length \mathbf{a}_i such that all (\mathbf{x}, y) drawn from P satisfies $\text{sign}(\langle \mathbf{a}_i, \mathbf{x} \rangle) = y$. Let $\mathbf{A} \in \mathbb{R}^{m \times d}$, rows of which are \mathbf{a}_i^\top . Since the parameters \mathbf{a}_i lie in some k -dimensional subspace with $k \ll \min(m, n)$, there exists $\mathbf{W} \in \mathbb{R}^{k \times d}$, $\mathbf{C} \in \mathbb{R}^{m \times k}$ such that $\mathbf{A} = \mathbf{C}\mathbf{W}$. Rows $\mathbf{w}_1^\top, \dots, \mathbf{w}_k^\top$ can be seen as k linear meta-features that are sufficient to learn m tasks. In each step when the current feature subspace cannot achieve low error, we learn new features. Then we take the refinement step to keep a minimal dimensional subspace that is close to all current features $\tilde{\mathbf{w}}_1, \dots, \tilde{\mathbf{w}}_k$.

Lemma 1. *Given two unit vectors \mathbf{u}, \mathbf{v} and a distribution D . If D satisfies Assumption 1, then there exist nonzero constants c' and c'' such that $c' \|\mathbf{u} - \mathbf{v}\|_2 \leq d_D(\mathbf{u}, \mathbf{v}) \leq c'' \|\mathbf{u} - \mathbf{v}\|_2$.*

Proof. By Assumption 1, $\exists c_1, c_2$ s.t. $c_1 \theta(\mathbf{u}, \mathbf{v}) \leq d_D(\mathbf{u}, \mathbf{v}) \leq c_2 \theta(\mathbf{u}, \mathbf{v})$. Using the Taylor expansion of cosine function, we know $1 - x^2/2 \leq \cos(x) \leq 1 - x^2/2! + x^4/4! \leq 1 - 11x^2/24$. Since $\|\mathbf{u} - \mathbf{v}\|_2^2 = 2 - 2\cos(\theta(\mathbf{u}, \mathbf{v}))$, we have $\sqrt{\frac{11}{12}}\theta(\mathbf{u}, \mathbf{v}) \leq \|\mathbf{u} - \mathbf{v}\|_2 \leq \theta(\mathbf{u}, \mathbf{v})$. Choose $c' = c_1, c'' = \sqrt{\frac{12}{11}}c_2$, we get the results proved. \square

Lemma 2. *Let S be a set of subspaces. Let $X = \{\mathbf{x} \in B_k(0, 1) | d(\mathbf{x}, \mathbf{Y}) \leq r, \forall \mathbf{Y} \in S\}$. Here $B_k(0, 1)$ is the unit ball in k -dimension. Then the set X is a symmetric convex set.*

Proof. For any $\mathbf{x} \in X, \forall \mathbf{Y} \in S$, since $d(\mathbf{x}, \mathbf{Y}) = d(-\mathbf{x}, \mathbf{Y})$, we have $-\mathbf{x} \in X$. So S is symmetric about the origin.

For any $\mathbf{x}_1, \mathbf{x}_2 \in X$, for any $\mathbf{Y} \in S$, we have $d(\mathbf{x}_1, \mathbf{Y}) \leq r, d(\mathbf{x}_2, \mathbf{Y}) \leq r$. Let \mathbf{P} be the projection matrix of \mathbf{Y} ,

$$\begin{aligned} & (\|\mathbf{P}\mathbf{x}_1\|_2 + \|\mathbf{P}\mathbf{x}_2\|_2)^2 - \|\mathbf{P}(\mathbf{x}_1 + \mathbf{x}_2)\|_2^2 \\ &= \mathbf{x}_1^\top \mathbf{P}^\top \mathbf{P} \mathbf{x}_1 + \mathbf{x}_2^\top \mathbf{P}^\top \mathbf{P} \mathbf{x}_2 + 2\|\mathbf{P}\mathbf{x}_1\|_2 \|\mathbf{P}\mathbf{x}_2\|_2 - (\mathbf{x}_1 + \mathbf{x}_2)^\top \mathbf{P}^\top \mathbf{P} (\mathbf{x}_1 + \mathbf{x}_2) \\ &= \|\mathbf{P}\mathbf{x}_1\|_2 \|\mathbf{P}\mathbf{x}_2\|_2 - (\mathbf{P}\mathbf{x}_1)^\top (\mathbf{P}\mathbf{x}_2) \geq 0 \end{aligned}$$

So we have

$$d\left(\frac{\mathbf{x}_1 + \mathbf{x}_2}{2}, \mathbf{Y}\right) = \left\| \mathbf{P} \left(\frac{\mathbf{x}_1 + \mathbf{x}_2}{2} \right) \right\|_2 \leq \left\| \mathbf{P} \left(\frac{\mathbf{x}_1}{2} \right) \right\|_2 + \left\| \mathbf{P} \left(\frac{\mathbf{x}_2}{2} \right) \right\|_2 = \frac{1}{2}d(\mathbf{x}_1, \mathbf{Y}) + \frac{1}{2}d(\mathbf{x}_2, \mathbf{Y}) \leq r$$

For a fixed \mathbf{Y} , $\{\mathbf{x} \in B_k(0, 1) | d(\mathbf{x}, \mathbf{Y}) \leq r\}$ is closed and thus convex. Therefore

$$X = \bigcap_{\mathbf{Y} \in S} \{\mathbf{x} \in B_k(0, 1) | d(\mathbf{x}, \mathbf{Y}) \leq r\}$$

is a convex set. □

Lemma 3. For any $\hat{k} \in [\tilde{k}]$, let

$$Y_{i_{\hat{k}}} := \{\mathbf{V} | d(\tilde{\mathbf{w}}_{i_j}, \mathbf{V}) \leq c_1 \epsilon_{acc}, \forall j \leq \hat{k}\},$$

$$X_{i_{\hat{k}}} := \{\mathbf{x} \in B_k(0, 1) | d(\mathbf{x}, \mathbf{V}) \leq (c_1 + \frac{1}{c'}) \epsilon_{acc}, \forall \mathbf{V} \in Y_{i_{\hat{k}}}\},$$

where c_1, c' are small constants. Then for any $j \geq \hat{k}, \pm \mathbf{a}_{i_{\hat{k}}} \in X_{i_j}$.

Proof. For $\forall \mathbf{V} \in Y_{i_{\hat{k}}}$, $d(\tilde{\mathbf{w}}_{i_{\hat{k}}}, \mathbf{V}) \leq c_1 \epsilon_{acc}$. Since we learn the feature vector $\tilde{\mathbf{w}}_{i_{\hat{k}}}$ within error ϵ_{acc} , by Lemma 1, we have $d(\mathbf{a}_{i_{\hat{k}}}, \tilde{\mathbf{w}}_{i_{\hat{k}}}) \leq \epsilon_{acc}/c'$. So $d(\mathbf{a}_{i_{\hat{k}}}, \mathbf{V}) \leq (c_1 + \frac{1}{c'}) \epsilon_{acc}$. Hence $\mathbf{a}_{i_{\hat{k}}} \in X_{i_{\hat{k}}}$. Because $d(\mathbf{x}, \mathbf{V}) = d(-\mathbf{x}, \mathbf{V})$, we also know $-\mathbf{a}_{i_{\hat{k}}} \in X_{i_{\hat{k}}}$. Also $Y_{i_1} \supseteq Y_{i_2} \supseteq \dots \supseteq Y_{i_{\hat{k}}}$, so $X_{i_1} \subseteq X_{i_2} \subseteq \dots \subseteq X_{i_{\hat{k}}}$. So for any $j \geq \hat{k}, \pm \mathbf{a}_{i_{\hat{k}}} \in X_{i_j}$. □

Corollary 1. For any $\hat{k} \in [\tilde{k}]$, $\text{conv}(\pm \mathbf{a}_{i_1}, \dots, \pm \mathbf{a}_{i_{\hat{k}}}) \subseteq X_{i_{\hat{k}}}$.

Proof. From Lemma 3, we know $\pm \mathbf{a}_{i_1}, \dots, \pm \mathbf{a}_{i_{\hat{k}}} \in X_{i_{\hat{k}}}$. Combined with Lemma 2, we get the corollary. □

Lemma 4 (Max Ellipsoid). Let $K \subset \mathbb{R}^k$ be a symmetric convex body and $E(K)$ be the maximum volume ellipsoid contained in K . For a vector \mathbf{u} on the boundary of $E(K)$, i.e. $\mathbf{u} \in \text{bd}(E(K))$, let $K' = \text{conv}(K, 2\sqrt{k}\mathbf{u}, -2\sqrt{k}\mathbf{u})$. Then,

$$\frac{\text{vol}(E(K'))}{\text{vol}(E(K))} \geq \frac{13}{10}.$$

Proof. Since $E(K) \subseteq K$, we have

$$K'' := \text{conv}(E(K), 2\sqrt{k}\mathbf{u}, -2\sqrt{k}\mathbf{u}) \subseteq K'$$

So it suffices to prove that

$$\frac{\text{vol}(E(K''))}{\text{vol}(E(K))} \geq \frac{13}{10}$$

First, let's assume that $E(K) = B_k(0, 1)$, *i.e.*, the unit ball around the origin, and the vector \mathbf{u} is $(1, 0, \dots, 0)^\top \in B_k(0, 1)$. By symmetry, we can assume the ellipsoid

$$E(K'') = \left\{ \mathbf{x} \left| \frac{x_1^2}{a^2} + \sum_{i=2}^k \frac{x_i^2}{b^2} \leq 1 \right. \right\}$$

We consider the two-dimensional slice first, where $\mathbf{u} = (1, 0)^\top$, $E(K'') = \{(x, y) | \frac{x^2}{a^2} + \frac{y^2}{b^2} \leq 1\}$. Direct calculation shows that

$$\forall (x, y) \in K'', y \leq -\frac{x}{\sqrt{4k-1}} + \frac{2\sqrt{k}}{\sqrt{4k-1}}$$

where $y = -\frac{x}{\sqrt{4k-1}} + \frac{2\sqrt{k}}{\sqrt{4k-1}}$ is the line tangent to the unit ball and go across the point $(\frac{1}{2\sqrt{k}}, \frac{\sqrt{4k-1}}{2\sqrt{k}})$ and the point $(2\sqrt{k}, 0)$ (see Figure 3.2). So for any point on the boundary of $E(K'')$, it also satisfies




Figure 3.2: Maximum volume ellipsoid

$$y^2 = b^2 \left(1 - \frac{x^2}{a^2} \right) \leq \left(-\frac{x}{\sqrt{4k-1}} + \frac{2\sqrt{k}}{\sqrt{4k-1}} \right)^2$$

Simplifying the inequality we get

$$\left(\frac{b^2}{a^2} + \frac{1}{4k-1} \right) x^2 - \frac{4\sqrt{k}}{4k-1} x + \frac{4k}{4k-1} - b^2 \geq 0$$

To ensure that the quadratic inequality holds, let the determinant equal zero, and we get $a^2 = 4k - b^2(4k-1)$. So,

$$\left\{ (x, y) \left| \frac{x^2}{4k - b^2(4k-1)} + \frac{y^2}{b^2} \leq 1 \right. \right\} \subseteq K'' \text{ for } b < 1.$$

By symmetry, in k dimensions, we get that $E_b \subseteq E(K'')$ for $b < 1$.

$$E_b = \left\{ \mathbf{x} \left| \frac{x_1^2}{4k - b^2(4k-1)} + \sum_{i=2}^k \frac{x_i^2}{b^2} \leq 1 \right. \right\}$$

The volume of this ellipsoid is

$$\text{vol}(E_b) = \sqrt{(4k - b^2(4k-1)) b^{2k-2}} \cdot \text{vol}(B_k(0, 1))$$

Let $f(b) = (4k - b^2(4k - 1))b^{2k-2}$. Calculate its derivative and let it to be zero, so we get $\hat{b}^2 = 1 - \frac{3}{4k-1} < 1$.

$$f(\hat{b}) = 4 \left(1 - \frac{3}{4k-1}\right)^{k-1} = 4 \left(\left(1 - \frac{3}{4k-1}\right)^{\frac{4k-4}{3}} \right)^{\frac{3}{4}} \geq 4 \left(\frac{1}{e}\right)^{\frac{3}{4}}$$

Hence we know

$$\text{vol}(E_{\hat{b}}) \geq 2 \left(\frac{1}{e}\right)^{\frac{3}{8}} \text{vol}(B_k(0, 1)) \geq \frac{13}{10} \text{vol}(B_k(0, 1))$$

Finally for any symmetric ellipsoid $E(K) := \mathbf{A}^{1/2}B_k(0, 1)$ and any $\mathbf{u} = \mathbf{A}^{1/2}\mathbf{u}_0$ on the boundary of $E(K)$, where \mathbf{u}_0 is the unit vector corresponding to \mathbf{u} . There exists an orthogonal matrix, say \mathbf{Q} , that rotates \mathbf{u}_0 to $(1, 0, \dots, 0)^\top$. That is $\mathbf{Q}\mathbf{u}_0 = (1, 0, \dots, 0)^\top$. Define an affine bijective transformation $T := \mathbf{A}^{1/2}\mathbf{Q}^\top \mathbf{x}$, with $T^{-1}(\mathbf{x}) = \mathbf{Q}\mathbf{A}^{-1/2}\mathbf{x}$. Then

$$T\left((1, 0, \dots, 0)^\top\right) = \mathbf{A}^{1/2}\mathbf{Q}^\top \mathbf{Q}\mathbf{u}_0 = \mathbf{A}^{1/2}\mathbf{u}_0 = \mathbf{u}$$

$$\begin{aligned} T(B_k(0, 1)) &= \{T(\mathbf{y}) \mid \mathbf{y}^\top \mathbf{y} \leq 1\} \\ &= \{\mathbf{x} \mid T^{-1}(\mathbf{x})^\top T^{-1}(\mathbf{x}) \leq 1\} \\ &= \{\mathbf{x} \mid \mathbf{x}^\top \mathbf{A}^{-1/2}\mathbf{Q}^\top \mathbf{Q}\mathbf{A}^{-1/2}\mathbf{x} \leq 1\} \\ &= \{\mathbf{x} \mid \mathbf{x}^\top \mathbf{A}^{-1}\mathbf{x} \leq 1\} \\ &= E(K) \end{aligned}$$

So we get $T(E_{\hat{b}}) \subset E(K'')$. Since the ratio of volumes is invariant under affine transformation, we have

$$\frac{\text{vol}(E(K''))}{\text{vol}(E(K))} \geq \frac{\text{vol}(T(E_{\hat{b}}))}{\text{vol}(T(B_k(0, 1)))} = \frac{\text{vol}(E_{\hat{b}})}{\text{vol}(B_k(0, 1))} \geq \frac{13}{10}.$$

□

The next lemma gives us a stopping condition.

Lemma 5. *Let $P = \text{conv}(\mathbf{x}_1, \dots, \mathbf{x}_m)$ be a polytope in \mathbb{R}^k with each \mathbf{x}_i of unit Euclidean length. Then, the maximum volume ellipsoid contained in P satisfies*

$$\text{vol}(E(P)) \leq 2\sqrt{2}e \left(\sqrt{\frac{2 \log 2m}{k}} \right)^k \text{vol}(B_k(0, 1)).$$

Proof. Recall the polar of a convex body P is the convex body defined as

$$P^* = \{\mathbf{x} : \langle \mathbf{x}, \mathbf{y} \rangle \leq 1 \text{ for all } \mathbf{y} \in P\}.$$

By the Blaschke-Santaló inequality, we have

$$\text{vol}(P) \text{vol}(P^*) \leq \text{vol}(B_k(0, 1))^2.$$

Next we lower bound the volume of P^* . Note that P^* is the intersection of exactly m halfspaces, each tangent to the unit ball. Consider the ball $B_k(0, r)$ with $r = \sqrt{\frac{k-1}{2 \log(2m)}}$. By Lemma 6, each

halfspace cuts off a cap of this ball, of volume at most $e^{-(k-1)/2r^2} = \frac{1}{2m}$ of the volume of $B_k(0, r)$. Therefore, the volume that is in the intersection of all m halfspaces is at least $\text{vol}(B_k(0, r))/2$ and hence, this is a lower bound on the volume of P^* . Using this, we have,

$$\text{vol}(P) \leq \frac{\text{vol}(B_k(0, 1))^2}{\text{vol}(B_k(0, r))/2} \leq 2r^{-k} \text{vol}(B_k(0, 1))$$

Furthermore, since $1 + x \leq e^x$, we can derive the following and complete the proof.

$$2r^{-k} = 2\left(\sqrt{\frac{2\log(2m)}{k-1}}\right)^k = 2\left(\sqrt{\frac{2\log(2m)}{k}}\right)^k \left(1 + \frac{1}{k-1}\right)^{\frac{k-1}{2}} \sqrt{1 + \frac{1}{k-1}} \leq 2\sqrt{2e} \left(\sqrt{\frac{2\log(2m)}{k}}\right)^k$$

□

Lemma 6. (Lemma 4.1 from [LV07]) For any $\frac{1}{\sqrt{k}} < t < 1$ and halfspace H at distance tr from the origin,

$$\text{vol}(B_k(0, r) \cap H) \leq e^{-t^2 k/2} \text{vol}(B_k(0, r)).$$

We can now prove Theorem 4.

Proof. Let's first prove the linear case. Since the ground truth feature space lies in a k -dimensional subspace $\mathbf{V}^* = \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_m) \subseteq \mathbb{R}^k$, the span of the true features is \mathbb{R}^k . Along training, when we deal with the i_k -th task, let $\hat{\mathbf{w}}_{i_k}$ be the new feature we learn to ensure that the i_k -th task has error no more than $\epsilon_{acc} = \frac{\epsilon}{2\sqrt{k}c''(c_1+1/c')}$, where c', c'' are universal constants defined in Lemma 1 and c_1 is the approximation constant of Optimization Problem 2.1. WLOG we assume $\hat{\mathbf{w}}_{i_k}$ to be a unit vector. By Lemma 1,

$$d(\mathbf{a}_{i_k}, \hat{\mathbf{w}}_{i_k}) \leq \frac{\epsilon_{acc}}{c'} = \frac{\epsilon}{2\sqrt{k}c''(c_1c' + 1)}$$

Denote \mathbf{V}_{i_k} be the feature subspace after fine-tuning (optimization). Since there exists a k -dimensional subspace \mathbf{V}^* satisfying all constraints and the algorithm outputs a constant-factor approximation of Optimization Problem 2.1, we can get a c_1 -approximation solution with dimension c_2k for constants c_1, c_2 . So we know in the end the dimension of the feature subspace we get is $O(k)$.

Let $B_k(0, 1)$ be the unit ball on the subspace. Denote the set of all possible solutions of the optimization as $Y_{i_k} := \{\mathbf{V} | d(\tilde{\mathbf{w}}_{i_j}, \mathbf{V}) \leq c_1\epsilon_{acc}, \forall j \leq \hat{k}\}$. Let X_{i_k} be all the vectors in the unit ball that is within distance $(c_1 + \frac{1}{c'})\epsilon_{acc} = \frac{\epsilon}{2\sqrt{k}c''}$ to all of the subspaces in Y_{i_k} , that is $X_{i_k} := \{\mathbf{x} \in B_k(0, 1) | d(\mathbf{x}, \mathbf{V}) \leq \frac{\epsilon}{2\sqrt{k}c''}, \forall \mathbf{V} \in Y_{i_k}\}$. By Lemma 2 and Corollary 1, we know X_{i_k} is a convex set containing $\{\pm \mathbf{a}_{i_1}, \dots, \pm \mathbf{a}_{i_{\hat{k}}}\}$. We will show next that after learning $\tilde{k} = O(k \log(\log(k)/\epsilon))$ new tasks, $X_{i_{\tilde{k}}}$ contains the ball $B_k(0, 1/2\sqrt{k})$.

In the initial step, Y_0 is the set of all subspaces. We naturally have $B(0, \frac{\epsilon}{2\sqrt{k}c''}) \subseteq X_0$ since for any $\mathbf{x} \in B(0, \epsilon)$, $d(\mathbf{x}, \mathbf{V}) \leq d(\mathbf{O}, \mathbf{V}) + d(\mathbf{x}, \mathbf{O}) \leq \frac{\epsilon}{2\sqrt{k}c''}$. Here \mathbf{O} is the origin. So we know $\text{vol}(X_0) \geq (\frac{\epsilon}{2\sqrt{k}c''})^k V_0$, where V_0 is the volume of the unit ball in \mathbb{R}^k . Encountering the i_k -th task, the current feature subspaces $\mathbf{V}_{i_{k-1}}$ cannot ensure an ϵ error. By Lemma 1, $d(\mathbf{a}_{i_k}, \mathbf{V}_{i_{k-1}}) \geq \frac{\epsilon}{c''}$. Hence $d(\frac{\mathbf{a}_{i_k}}{2\sqrt{k}}, \mathbf{V}_{i_{k-1}}) \geq \frac{\epsilon}{2\sqrt{k}c''}$, which means $\pm \frac{\mathbf{a}_{i_k}}{2\sqrt{k}} \notin X_{i_{k-1}}$. Consequently, the vector $\mathbf{u} = u'\mathbf{a}_{i_k} \in \text{bd}(E(X_{i_{k-1}}))$ satisfies $\|\mathbf{u}\| < \frac{1}{2\sqrt{k}}$. According to lemma 4, we know that

$$\text{vol}\left(\text{conv}\left(X_{i_{k-1}}, \pm \mathbf{a}_{i_k}\right)\right) \geq \text{vol}\left(\text{conv}\left(X_{i_{k-1}}, \pm 2\sqrt{k}\mathbf{u}\right)\right) \geq \frac{13}{10} \text{vol}\left(X_{i_{k-1}}\right)$$

Also because of the convexity of $X_{i_{\tilde{k}}}$, we have $X_{i_{\tilde{k}}} \supseteq \text{conv}(X_{i_{\tilde{k}-1}}, \pm \mathbf{a}_{i_{\tilde{k}}})$. Therefore,

$$\frac{\text{vol}(X_{i_{\tilde{k}}})}{\text{vol}(X_{i_{\tilde{k}-1}})} \geq \frac{13}{10}$$

The algorithm will terminate when $X_{i_{\tilde{k}}} \supseteq E(X_{i_{\tilde{k}}}) \supseteq B_k(0, \frac{1}{2\sqrt{k}})$. So for any unit vector $\mathbf{a} \in B_k(0, 1)$, $d(\mathbf{a}, \mathbf{V}_{i_{\tilde{k}}}) \leq \frac{\epsilon}{c'}$. This means that after learning \tilde{k} new features, for any new tasks with weights lie in the same feature subspace \mathbf{V}^* , the current features can achieve error less than ϵ . By Lemma 5, we know the volume of $E(X_{i_{\tilde{k}}})$ is upper bounded by $2\sqrt{2}e \left(\sqrt{\frac{2 \log 2\tilde{k}}{k}} \right)^k \text{vol}(B_k(0, 1))$. It grows by a constant factor $\frac{13}{10}$ whenever we learn a new feature. So the number of tasks we learn with error ϵ_{acc} in the algorithm \tilde{k} satisfies:

$$\left(\frac{\epsilon}{2c'\sqrt{k}} \right)^k \cdot \left(\frac{13}{10} \right)^{\tilde{k}} \leq 2\sqrt{2}e \left(\sqrt{\frac{2 \log 2\tilde{k}}{k}} \right)^k$$

Simplify and take the log to both sides, so we will have $\tilde{k} - \frac{k}{2} \log(2 \log(2\tilde{k})) \leq \text{const} + k \log(\frac{1}{\epsilon})$. This will lead to $\tilde{k} \leq O(k \log(\log(k)/\epsilon))$.

Moreover, the sample complexity [BL13] of learning one task with input dimension d up to ϵ error is $O(d \log(1/\epsilon)/\epsilon)$. So the sample complexity of our algorithm is

$$\begin{aligned} & O\left(\frac{dk}{\epsilon_{acc}} \log(1/\epsilon_{acc}) \log(\log(k)/\epsilon)\right) + O\left(\frac{km}{\epsilon} \log(1/\epsilon)\right) \\ &= O\left(\frac{dk\sqrt{k}}{\epsilon} \log(k/\epsilon) \log(\log(k)/\epsilon) + \frac{km}{\epsilon} \log(1/\epsilon)\right) \\ &= \tilde{O}\left(\frac{dk\sqrt{k}}{\epsilon} + \frac{km}{\epsilon}\right). \end{aligned}$$

Finally, for the nonlinear case, we consider the kernel of the features. These features live in a potentially infinite-dimensional space. If we assume there is an oracle to get a constant approximation for the optimization problem 2.1, the dimension of features will be $O(k)$ in the end. Other bounds follow precisely the same as the linear case. \square

It is noteworthy that the convex set X_i (of feature vectors in \mathbb{R}^k) we keep in the proof is defined to be close to any possible subspaces that are close to the subspace spanned by new features $\tilde{\mathbf{w}}_i$. So our proof is quite general for any lifelong learning algorithm that dynamically expands the architecture, *e.g.* the basic LLL algorithm. Consequently, we can prove Theorem 3 as follows.

Proof of Theorem 3. The proof exactly follows the proof of the Theorem 4. Without the refinement of the feature subspace, the subspace we get is still in the set $Y_{i_{\tilde{k}}}$. Since $X_{i_{\tilde{k}}}$ will eventually cover the $B_k(0, \frac{1}{2\sqrt{k}})$ after learning $O(k \log(\log(k)/\epsilon))$ features, the dimension of the feature subspace is at most $O(k \log(\log(k)/\epsilon))$ for the linear features and $O(k^2 \log(\log(k)/\epsilon))$ for the nonlinear features. So the total sample complexity is $\tilde{O}(dk^{1.5}/\epsilon + km/\epsilon)$ for linear features and $\tilde{O}(dk^{2.5}/\epsilon + k^2m/\epsilon)$ for nonlinear features. \square

3.2 Proof of the Approximate Optimization for the Linear Setting

Proof of Theorem 5. Let (\mathbf{X}^*, t^*) be a solution of the SDP (2.2) with singular value decomposition (SVD) $\mathbf{X}^* = \sum_{i=1}^d \lambda_i \mathbf{u}_i \mathbf{u}_i^\top$, where $0 \leq \lambda_1 \leq \dots \leq \lambda_d \leq 1$, and $\sum_{i=1}^d \lambda_i = d - k$. Let \mathbf{V}' be the span of $\{\mathbf{u}_1, \dots, \mathbf{u}_{k'}\}$. Then the squared distance of any vector \mathbf{a} to \mathbf{V}' is $\sum_{i=k'+1}^d (\mathbf{a}^\top \mathbf{u}_i)^2$. In the meantime, the SDP assigns a value of $\mathbf{a}^\top \mathbf{X}^* \mathbf{a} = \sum_{i=1}^d \lambda_i (\mathbf{a}^\top \mathbf{u}_i)^2$. Thus, the multiplicative increase in squared distance is at most

$$\frac{\sum_{i=k'+1}^d (\mathbf{a}^\top \mathbf{u}_i)^2}{\sum_{i=1}^d \lambda_i (\mathbf{a}^\top \mathbf{u}_i)^2} \leq \frac{1}{\lambda_{k'+1}}$$

Now since the sum of all eigenvalues is $d - k$ and each one is at most 1, for $k' \geq k$, we must have

$$\lambda_{k'+1} \geq \frac{(d - k) - (d - k' - 1)}{k' + 1} = \frac{k' - k + 1}{k' + 1}$$

Choose $k' = ck - 1$, and we would have

$$\frac{\sum_{i=ck}^d (\mathbf{a}^\top \mathbf{u}_i)^2}{\sum_{i=1}^d \lambda_i (\mathbf{a}^\top \mathbf{u}_i)^2} \leq \frac{1}{\lambda_{ck}} \leq 1 + \frac{1}{c - 1}$$

Since there exists a subspace \mathbf{V}^* of k dimension with $d(\tilde{\mathbf{w}}_i, \mathbf{V}^*) \leq \epsilon_{acc}, \forall i \in [\hat{k} + 1]$, we have $\sum_{i=1}^d \lambda_i (\tilde{\mathbf{w}}_i^\top \mathbf{u}_i)^2 \leq \epsilon_{acc}^2$. Consequently, $\text{span}(\mathbf{u}_1, \dots, \mathbf{u}_{ck-1})$ is a $(ck - 1)$ -dimensional approximation of \mathbf{V}^* with the approximation factor $\sqrt{1 + \frac{1}{c-1}}$ in maximum distance.

Specifically, for $c = 2$, we have $d^2(\tilde{\mathbf{w}}_i, \mathbf{V}') = \sum_{i=2k}^d (\tilde{\mathbf{w}}_i^\top \mathbf{u}_i)^2 \leq 2\epsilon_{acc}^2$. \square

Theorem 5 guarantees that in the linear case, our refinement step will always output an $O(k)$ -dimensional feature subspace of error bounded by $O(\epsilon_{acc})$ without using any additional samples.

4 A Lower Bound for General Lifelong Learning Algorithms

In this section, we show that our sample complexity bound for *general* lifelong learning algorithms in the linear setting is asymptotically the best possible, assuming black-box access to a learner for a single linear target.

As a warm-up, we first show that the analysis of *our* lifelong algorithm is *tight*.

Theorem 6 (Tight Example). *Using the same condition and algorithm as in Theorem 4, the total sample complexity is $\Omega(dk^{1.5}/\epsilon + km/\epsilon)$.*

Proof. Assume the task vectors $\mathbf{a}_i = \mathbf{e}_i, 1 \leq i \leq k - 1$, where $\mathbf{e}_i \in \mathbb{R}^k$ has 1 in i -th coordinate and 0 otherwise. Assume that our algorithm accurately learns these $k - 1$ tasks and returns $\tilde{\mathbf{w}}_i = \mathbf{a}_i + \epsilon_{acc} \mathbf{a}_k$. Then for a new task's weight $\frac{1}{\sqrt{k}} \sum_{i=1}^{k-1} \alpha_i \mathbf{a}_i$, $\alpha_i \in \{1, -1\}$. Based on the current features, the error we make on this task is $O(\sqrt{k} \epsilon_{acc})$. If we assume that $\epsilon_{acc} = \omega(\epsilon/\sqrt{k})$, then we need to learn these 2^{k-2} tasks accurately as well given learning any of them will not help with others (except its negative). Then the total complexity will be exponential with respect to k . So $\epsilon_{acc} = O(\epsilon/\sqrt{k})$, and thus we have the sample complexity $\Omega(dk^{1.5}/\epsilon + km/\epsilon)$. \square

Theorem 6 shows that the analysis of our algorithm's sample complexity is tight for the linear setting. The main result of this section is a lower bound for general lifelong learning algorithms.

Theorem 2 (Lower bound). *Suppose that a lifelong learner has black-box access to a single task learner that takes an error parameter ϵ as input and is allowed to return any vector that is within distance ϵ of the true target unit vector, using $\Theta(d/\epsilon)$ samples in \mathbb{R}^d . Then, there exists a distribution of m tasks, $m = 2^{\Theta(k)}$ such that for any lifelong learning algorithm, WHP, the total number of samples required to learn all m tasks up to error ϵ is $\Omega(dk^{1.5}/\epsilon + km/\epsilon)$.*




Figure 4.1: Geometric illustration of the lower bound examples when $k = 2, d = 3$. The errors that the algorithm makes concentrate on the third coordinate, and thus lead to the large angle between the learned feature subspace V and the underlying one U . Any new task, e.g., \mathbf{a}_3 that lies on the span of $\{\mathbf{a}_1, \mathbf{a}_2\}$ cannot help improve the representation V .

Proof idea. Consider a sequence of tasks $\mathbf{a}_i = \begin{cases} \mathbf{e}_i, & 1 \leq i \leq k \\ \sum_{j \in S} x_j \mathbf{e}_j, & i > k \end{cases}$, where \mathbf{e}_j is the standard unit vector, $x_j \stackrel{\text{i.i.d}}{\sim} \text{Bernoulli}(1/2)$, $S \subset [k]$ is a subset of indices. The proof is mainly by constructing an adversarial output of the algorithm where the errors that it makes concentrate on one coordinate, say $k + 1$. (See Figure 4.1.) We will show by the following steps:

- (1) After learning the first k tasks, each with error ϵ_i , the angle between the learned feature subspace and the underlying one is at least $\Omega(\sqrt{\sum_{i=1}^k \epsilon_i^2})$.
- (2) For each new task followed, the angle of the new task to the learned subspace is at least $\Omega(\sqrt{\sum_{i \in S} \epsilon_i^2})$ with high probability.
- (3) Learning such a new task does not improve the representation.
- (4) To solve all tasks up to error ϵ , we will need each $\epsilon_i = O(\epsilon/\sqrt{k})$, and it leads to the sample complexity.

We will show them accordingly with the following lemmas.

Lemma 7. *For k orthonormal tasks $\mathbf{a}_i = \mathbf{e}_i, 1 \leq i \leq k$, for any algorithm that learns task i within error ϵ_i , i.e., $d_D(\mathbf{a}_i, \hat{\mathbf{a}}_i) \leq \epsilon_i$. Let $U = \text{span}(\mathbf{a}_1, \dots, \mathbf{a}_k)$ be original feature subspace, and $V = \text{span}(\hat{\mathbf{a}}_1, \dots, \hat{\mathbf{a}}_k)$ be the learned subspace. Then there exist feasible outputs $\hat{\mathbf{a}}_i$ such that $\theta(U, V) = \Omega(\sqrt{\sum_{i=1}^k \epsilon_i^2})$.*

Proof. Since the algorithm learns task i within error ϵ_i , with Lemma 1, there exists a constant c such that $\|\mathbf{a}_i - \hat{\mathbf{a}}_i\| \leq \epsilon_i/c$. Because we does not consider constant factor, we assume WLOG

that $\|\mathbf{a}_i - \hat{\mathbf{a}}_i\| \leq \epsilon_i$. Consider the case where all errors made by the algorithm concentrate on the $k+1$ coordinate. Then the features learned by the algorithm are $\hat{\mathbf{a}}_i = \begin{pmatrix} \mathbf{e}_i \\ \epsilon_i \end{pmatrix}, 1 \leq i \leq k$. Denote $\mathbf{A} = (\mathbf{a}_1, \dots, \mathbf{a}_k) = \begin{pmatrix} \mathbf{I}_k \\ \mathbf{0}^\top \end{pmatrix}$, columns of which are the task vectors. Denote $\mathbf{s} = (\epsilon_1, \dots, \epsilon_k)^\top$. $\hat{\mathbf{A}} = (\hat{\mathbf{a}}_1, \dots, \hat{\mathbf{a}}_k) = \begin{pmatrix} \mathbf{I}_k \\ \mathbf{s}^\top \end{pmatrix}$. By the definition of the angles between subspace, we have $\theta(U, V) = \max_{\mathbf{x} \in U} \theta(\mathbf{x}, \mathbf{P}_V \mathbf{x})$, where \mathbf{P}_V is the projection matrix of the subspace V .

$$\begin{aligned} \mathbf{P}_V &= \hat{\mathbf{A}}(\hat{\mathbf{A}}^\top \hat{\mathbf{A}})^{-1} \hat{\mathbf{A}}^\top \\ &= \begin{pmatrix} \mathbf{I}_k \\ \mathbf{s}^\top \end{pmatrix} (\mathbf{I}_k + \mathbf{s} \mathbf{s}^\top)^{-1} (\mathbf{I}_k \quad \mathbf{s}) \\ &= \begin{pmatrix} \mathbf{I}_k \\ \mathbf{s}^\top \end{pmatrix} (\mathbf{I}_k - \mathbf{s} \mathbf{s}^\top + o((\mathbf{s} \mathbf{s}^\top)^3)) (\mathbf{I}_k \quad \mathbf{s}) \\ &\approx \begin{pmatrix} \mathbf{I}_k - \mathbf{s} \mathbf{s}^\top & \mathbf{s}(1 - \mathbf{s}^\top \mathbf{s}) \\ \mathbf{s}^\top(1 - \mathbf{s}^\top \mathbf{s}) & \mathbf{s}^\top \mathbf{s} - (\mathbf{s}^\top \mathbf{s})^2 \end{pmatrix} \end{aligned}$$

Let $\mathbf{x} = \mathbf{s} = (\epsilon_1, \dots, \epsilon_k)^\top \in U$. Since

$$\mathbf{P}_V \mathbf{x} = \left(\left(1 - \sum_{i=1}^k \epsilon_i^2\right) \epsilon_1, \dots, \left(1 - \sum_{i=1}^k \epsilon_i^2\right) \epsilon_k, \left(1 - \sum_{i=1}^k \epsilon_i^2\right) \sum_{i=1}^k \epsilon_i^2 \right)^\top,$$

we have

$$\tan \theta(\mathbf{x}, \mathbf{P}_V \mathbf{x}) = \frac{\left(1 - \sum_{i=1}^k \epsilon_i^2\right) \sum_{i=1}^k \epsilon_i^2}{\left(1 - \sum_{i=1}^k \epsilon_i^2\right) \sqrt{\sum_{i=1}^k \epsilon_i^2}} = \sqrt{\sum_{i=1}^k \epsilon_i^2}$$

So we know that $\theta(U, V) \geq \theta(\mathbf{x}, \mathbf{P}_V \mathbf{x}) = \Omega(\sqrt{\sum_{i=1}^k \epsilon_i^2})$. So we prove that the angle between the learned subspace V and the underlying one U is at least $\Omega(\sqrt{\sum_{i=1}^k \epsilon_i^2})$. \square

Lemma 8. For any vector $\mathbf{x} = \sum_{i \in S} x_i \mathbf{e}_i$, where $S \subset [k]$, $x_i \stackrel{i.i.d}{\sim} \text{Bernoulli}(1/2)$. Assume that $\forall i \in S, 0 \leq \epsilon_i \leq 2\sqrt{\sum_{i \in S} \epsilon_i^2 / |S|}$. Let $V = \text{span}(\{\hat{\mathbf{a}}_i\}, i \in S)$, where $\hat{\mathbf{a}}_i = \begin{pmatrix} \mathbf{e}_i \\ \epsilon_i \end{pmatrix}$. Then with high probability, $\theta(\mathbf{x}, V) = \Omega(\sqrt{\sum_{i \in S} \epsilon_i^2})$.

Proof. Denote $\mathbf{y} = \mathbf{P}_V(\mathbf{x})$, then $y_j = \begin{cases} x_j - \epsilon_j \sum_{i \in S} \epsilon_i x_i & \text{if } j \in S \\ (1 - \sum_{i=1}^k \epsilon_i^2) \sum_{i \in S} \epsilon_i x_i & \text{if } j = k+1 \end{cases}$. Then we can calculate the angle between \mathbf{x} and the subspace V as follows.

$$(\tan(\theta(\mathbf{x}, \mathbf{P}_V(\mathbf{x}))))^2 = \frac{(1 - \sum_{i \in S} \epsilon_i^2)^2 (\sum_{i \in S} \epsilon_i x_i)^2}{\sum_{i \in S} x_i^2 + (\sum_{i \in S} \epsilon_i^2) (\sum_{i \in S} \epsilon_i x_i)^2 - 2 (\sum_{i \in S} \epsilon_i x_i)^2} \geq \frac{\frac{1}{4} (\sum_{i \in S} \epsilon_i x_i)^2}{\sum_{i \in S} x_i^2}$$

The last inequality is because we learn each task well, so we can assume $\sum_{i \in S} \epsilon_i^2 \leq 1/2$. Then the probability that $\theta(\mathbf{x}, \mathbf{P}_V(\mathbf{x}))$ is greater than $O(\sqrt{\sum_{i \in S} \epsilon_i^2})$ is as follows.

$$\begin{aligned}
\mathbb{P} \left(\theta(\mathbf{x}, \mathbf{P}_V(\mathbf{x})) \geq \frac{1}{16} \sqrt{\sum_{i \in S} \epsilon_i^2} \right) &= \mathbb{P} \left(\tan^2 \theta(\mathbf{x}, \mathbf{P}_V(\mathbf{x})) \geq \frac{1}{256} \sum_{i \in S} \epsilon_i^2 \right) \\
&\geq \mathbb{P} \left(\frac{\frac{1}{4} \left(\sum_{i \in S} \epsilon_i x_i \right)^2}{\sum_{i \in S} x_i^2} \geq \frac{1}{256} \sum_{i \in S} \epsilon_i^2 \right) \\
&= \mathbb{P} \left(\left(\sum_{i \in S} \epsilon_i x_i \right)^2 \geq \frac{1}{64} \sum_{i \in S} \epsilon_i^2 \sum_{i \in S} x_i^2 \right) \\
&\geq \mathbb{P} \left(\left(\sum_{i \in S} \epsilon_i x_i \right)^2 \geq \frac{1}{64} |S| \sum_{i \in S} \epsilon_i^2 \right)
\end{aligned}$$

By Chernoff bound, for any $t > 0$,

$$\mathbb{P} \left(\sum_{i \in S} \epsilon_i x_i - \frac{1}{2} \sum_{i \in S} \epsilon_i \geq -t \sqrt{\sum_{i \in S} \epsilon_i^2} \right) \geq 1 - e^{-t^2/2}$$

Choose $t = \sqrt{|S|}/8$, we have

$$\mathbb{P} \left(\sum_{i \in S} \epsilon_i x_i \geq \frac{1}{2} \sum_{i \in S} \epsilon_i - \frac{1}{8} \sqrt{|S| \sum_{i \in S} \epsilon_i^2} \right) \geq 1 - e^{-|S|/128}$$

Note that with $\epsilon = \sqrt{\sum_{i \in S} \epsilon_i^2}$ and $0 \leq \epsilon_i \leq 2\epsilon/\sqrt{|S|}$ for all $i \in S$, we have

$$\sum_{i \in S} \epsilon_i \geq \frac{\sqrt{|S|}}{2} \epsilon.$$

Consequently, we have

$$\begin{aligned}
\sqrt{\frac{1}{64} |S| \sum_{i \in S} \epsilon_i^2} &\leq \frac{1}{2} \sum_{i \in S} \epsilon_i - \frac{1}{8} \sqrt{|S| \sum_{i \in S} \epsilon_i^2} \\
\mathbb{P} \left(\theta(\mathbf{x}, V) \geq \frac{1}{16} \sqrt{\sum_{i \in S} \epsilon_i^2} \right) &\geq \mathbb{P} \left(\left(\sum_{i \in S} \epsilon_i b_i \right)^2 \geq \frac{1}{64} |S| \sum_{i \in S} \epsilon_i^2 \right) \\
&\geq \mathbb{P} \left(\sum_{i \in S} \epsilon_i b_i \geq \frac{1}{2} \sum_{i \in S} \epsilon_i - \frac{1}{8} \sqrt{|S| \sum_{i \in S} \epsilon_i^2} \right) \\
&\geq 1 - e^{-|S|/128}
\end{aligned}$$

□

Lemma 9. For $b_1, \dots, b_k \geq 0$, with $\bar{b} = \sqrt{\sum_{i=1}^k b_i^2/k}$, let $b_i \geq \bar{b}/C$ for some constant $C > 1$. Then there exists a subset $S \subset [k]$ with $|S| \geq k(1-p)$ s.t. for all $i \in S$, we have

$$b_i \leq \sqrt{\frac{1}{p} \ln \frac{C^2}{1-p}} \sqrt{\frac{\sum_{i \in S} b_i^2}{|S|}}.$$

Proof. Let $S_1 = \{1, \dots, k\}$. Choose a constant $\gamma > 1$. We repeat the following procedure. For the j -th step, we are given a set $\{x_i, i \in S_j\}$, Let $S_{j+1} = \{i \in S, b_i \leq \gamma \sqrt{\sum_{i \in S_j} b_i^2/|S_j|}\}$. The algorithm terminates when $S_j = S_{j+1}$. Denote $p_j = 1 - |S_{j+1}|/|S_j|$. For the j -th step, we have

$$\sum_{i \in S_j} b_i^2 = \sum_{i \in S_{j+1}} b_i^2 + \sum_{i \in S_j \setminus S_{j+1}} b_i^2 \geq \sum_{i \in S_{j+1}} b_i^2 + p_j \gamma^2 \sum_{i \in S_j} b_i^2$$

This derives that

$$\sum_{i \in S_{j+1}} b_i^2 \leq (1 - p_j \gamma^2) \sum_{i \in S_j} b_i^2 \leq e^{-p_j \gamma^2} \sum_{i \in S_j} b_i^2$$

Accumulating all J steps, we have

$$\sum_{i \in S_J} b_i^2 \leq e^{-\gamma^2 \sum_{j=1}^{J-1} p_j} \sum_{i \in S_1} b_i^2 = e^{-\gamma^2 \sum_{j=1}^{J-1} p_j} k \bar{b}^2$$

From the condition that $b_i \geq \bar{b}/C$, we have

$$|S_J| \frac{\bar{b}^2}{C^2} \leq \sum_{i \in S_J} b_i^2 \leq e^{-\gamma^2 \sum_{j=1}^{J-1} p_j} k \bar{b}^2$$

From the definition of p_j , we know that

$$|S_J| = k \prod_{j=1}^{J-1} (1 - p_j) \geq k \left(1 - \sum_{j=1}^{J-1} p_j\right)$$

Denote $p = \sum_{j=1}^{J-1} p_j$, then we have

$$1 - p \leq e^{-\gamma^2 p C^2}$$

Choose $\gamma = \sqrt{\frac{1}{p} \ln \frac{C^2}{1-p}}$, then we get a set S_J with $|S_J| \geq k(1-p)$ satisfying for all $i \in S_J$,

$$b_i \leq \sqrt{\frac{1}{p} \ln \frac{C^2}{1-p}} \sqrt{\frac{\sum_{i \in S_J} b_i^2}{|S_J|}}$$

□

We can now prove Theorem 2 as follows.

Proof. Denote the underlying feature subspace as U . Consider a sequence of tasks with first k tasks the basis of the feature subspace, i.e., $\mathbf{a}_i = \mathbf{e}_i, 1 \leq i \leq k$. The lifelong learning algorithm learns task i up to error ϵ_i , and get k features $\hat{\mathbf{a}}_1, \dots, \hat{\mathbf{a}}_k$. If the number of tasks for which $\epsilon_i < \sqrt{2 \sum_{i=1}^k \epsilon_i^2 / 3k}$

is more than $k/4$, then we have that the total sample complexity is already $\Omega(dk^{1.5}/\sqrt{\sum_{i=1}^k \epsilon_i^2})$ and the theorem follows. So we assume that for at least $3k/4$ tasks, we have

$$\epsilon_i \geq \sqrt{2 \sum_{i=1}^k \epsilon_i^2 / 3k}.$$

Calling this subset S_1 , it follows that for each $i \in S_1$, we have $\epsilon_i \geq \sqrt{\sum_{i \in S_1} \epsilon_i^2 / 2|S_1|}$. Next, applying Lemma 9 to the set S_1 with cardinality at least $3k/4$, using $p = 1/3$ and $C = \sqrt{2}$, we get that there exists a set $S \subseteq S_1$ with $|S| \geq k/2$ and $\epsilon_i \leq 2\sqrt{\frac{\sum_{i \in S} \epsilon_i^2}{|S|}}$ for all $i \in S$. Consider the span $V := \{\hat{\mathbf{a}}_i, i \in S\}$. By Lemma 7, we know there exists feasible $\hat{\mathbf{a}}_i$ such that $\theta(U, V) = \Omega(\sqrt{\sum_{i \in S} \epsilon_i^2})$.

Next we consider the following tasks as $\mathbf{a}_j = \sum_{i \in S} x_{ji} \mathbf{e}_i$, $j > k$, where $x_{ji} \stackrel{\text{i.i.d.}}{\sim} \text{Bernoulli}(1/2)$. There are $2^{k/2}$ such tasks. By Lemma 8, we know that with high probability each new task is far from the learned subspace, *i.e.*,

$$\mathbb{P} \left(\theta(\mathbf{a}_j, V) \geq \frac{1}{16} \sqrt{\sum_{i \in S} \epsilon_i^2} \right) \geq 1 - e^{-k/256}$$

Assume that the single-task learning algorithm applied to this new task also induce error in the $(k+1)$ 'st coordinate. Say the new learned feature $\hat{\mathbf{a}}_j = (\mathbf{a}_j^\top, \sum_{i \in S} b_{ji} \epsilon_i)^\top$. Then the new learned feature is in the learned subspace V , which means that learning new tasks does not improve the learned subspace.

Since each new task is generated randomly, there are exponentially many new tasks that are far from the learned subspace but make no improvement by learning them. Therefore to ensure each task is learned with error ϵ , the only way is to let $\theta(U, V) \leq \epsilon$. This implies that $\sum_{i \in S} \epsilon_i^2 \leq c\epsilon^2$. By the generalized Hölder inequality [Fin92], we have

$$\sum_{i \in S} \frac{1}{\epsilon_i} \sum_{i \in S} \frac{1}{\epsilon_i} \sum_{i \in S} \epsilon_i^2 \geq \left(\sum_{i \in S} 1 \right)^3 \geq \left(\frac{k}{2} \right)^3$$

So $\sum_{i \in k} \frac{1}{\epsilon_i} \geq c'k^{1.5}/\epsilon$. The number of the samples needed to learn the tasks in set S is $\sum_{i \in S} \frac{d}{\epsilon_i} = \Omega(dk^{1.5}/\epsilon)$. So the overall sample complexity for the sequence of tasks are $\Omega(dk^{1.5}/\epsilon + km/\epsilon)$. \square

5 Simulations and Empirical Results

In this section, we describe our experimental studies. In Section 5.1, we run the basic LLL and LLL-RR algorithms in a task-incremental binary classification setting. Then we conduct class-incremental experiments on real dataset using our H-LLL algorithm in Section 5.2. The performance shows the benefits of our algorithm compared to existing continual learning algorithms.

5.1 Linear Features

Here we consider task-incremental lifelong learning in the setting of binary classification where $y = \text{sign}(\langle \mathbf{c}_i^*, \mathbf{W}^* \mathbf{x} \rangle)$. We choose the input dimension, $d = 100$, the number of tasks, $m = 100$, the

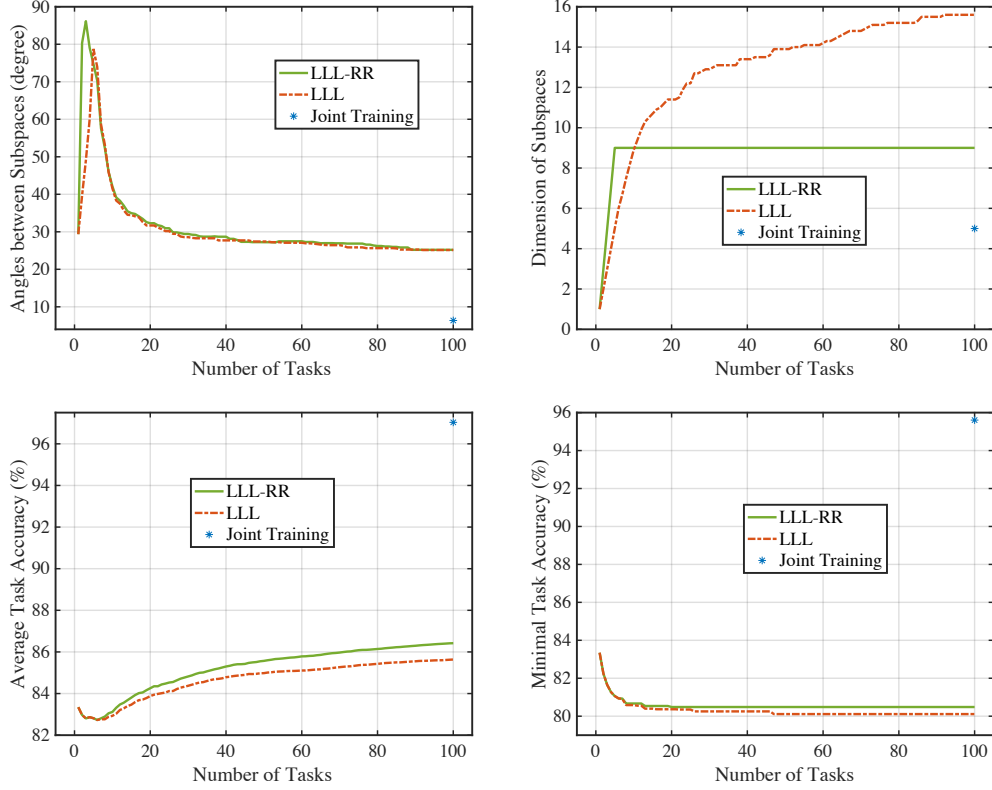


Figure 5.1: Simulation on linear features. $N=200, k=5, m=100, d=100$, averaged on 10 trials.

number of examples per task, $N = 200$, the dimension of feature subspace, $k = 5$. The parameters $c_{ij}^*, W_{ij}^* \sim \mathcal{N}(0, 1)$. The input data $X_i \sim \mathcal{N}(0, 1)$. We set the error threshold to be $\epsilon = 0.1$. We compare three methods: LLL (Basic lifelong learning algorithm), LLL-RR (lifelong learning algorithm with representation refinement), and Joint Training (offline training with all data jointly).

The average task accuracy and minimal task accuracy are computed for tasks encountered so far based on the current model. The angle between feature subspaces is calculated as their maximal principal angle. Formally, for two subspaces F and G , let \mathbf{P}, \mathbf{Q} to be the orthogonal matrices whose columns form an orthonormal basis of F and G . For the singular value decomposition $\mathbf{P}^\top \mathbf{Q} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^\top$, we define the principal angles between F and G as $\theta_i = \arccos(\Sigma_{ii})$, $\frac{\pi}{2} \geq \theta_1 \geq \dots \geq \theta_k \geq 0$. We calculate the angle between two subspaces F and G as the maximal principal angle, *i.e.*, $\arccos(\Sigma_{11})$.

As we can see in Figure 5.1, lifelong learning can continually learn better features while learning more tasks. Moreover, lifelong learning with refinement improves average accuracy, min accuracy, model size and convergence to the underlying feature subspace.

5.2 Image Classification

Experimental settings. We generally follow the experimental settings and evaluation protocol in [RKSL17]. In our experiments, we evaluate our H-LLL algorithm on CIFAR-100. We train all 100 classes in 10 splits and each split contains 10 classes. There is no class overlap between different splits. Each training data split can be viewed as a task and is fed to the neural network incrementally. Similar to [RKSL17], we use a fixed memory size of 2,000 exemplars. The final result are curves of the classification accuracies after each batch of classes. We use ResNet-18 [HZRS16] for all the

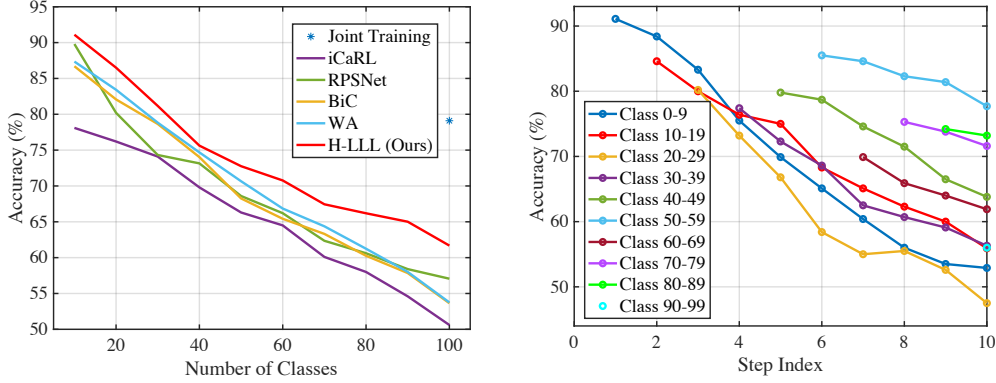


Figure 5.2: Classification accuracy on CIFAR-100.

encoders $f_j, \forall j$ and SGD with weight decay 0.0005. All the ResNet encoders are trained from scratch.

Accuracy vs. number of classes. In Fig. 5.2, we first show the comparison of incremental accuracies to some of the state-of-the-art methods including iCaRL [RKSL17], RPSNet [RHK⁺19], BiC [HPL⁺19] and WA [ZXG⁺20]. One can observe that our H-LLL algorithm significantly outperforms the other methods and yields an average incremental accuracy [RKSL17] of 73.8%, while the second best approach (WA) only achieves 69.8% accuracy.

Accuracy for different classes. In order to gain deeper understanding of the H-LLL algorithm, we examine the accuracy of different class splits in each step. From Fig. 5.2, we can see that the incremental accuracy for different class groups decreases in a slow and smooth way. This indicates that H-LLL is able to preserve knowledge of class concepts and effectively avoid catastrophic forgetting.

6 Discussion

We study, theoretically and empirically, the efficiency of lifelong learning when tasks share a low-dimensional feature representation. We introduce a refinement algorithm and bound its representation and sample complexity, and prove a matching lower bound for the sample complexity (for any lifelong learning algorithm). Our results show that: (1) lifelong learning provably converges for nonlinear feature representations, (2) refinement has provable benefits, and (3) lifelong learning is an efficient approach to multi-class/multi-task learning. Our work also indicates that (a) refinement can be practical and can dynamically keep the dimension of the representation bounded and (b) remembering only a small subset of previous examples suffices for efficient lifelong learning.

These results raise further questions. In the general setting of nonlinear features, how can we guarantee that the refinement is efficient in terms of time complexity? Our experiments suggest that SGD does well in practice. One complication with nonlinear features is that even solving a single offline task is nontrivial and needs further assumptions. E.g., is there an efficient lifelong learning algorithm for two-layer ReLU networks under nice input distributions?

Acknowledgements. This work was supported in part by NSF awards CCF-1909756, CCF-2007443 and CCF-2134105. Weiyang Liu is supported by a Cambridge-Tübingen Fellowship, an NVIDIA GPU grant, DeepMind and the Leverhulme Trust via CFI. We thank Le Song for helpful discussions.

References

- [Bax97] Jonathan Baxter. A bayesian/information theoretic model of learning to learn via multiple task sampling. *Machine learning*, 28(1):7–39, 1997.
- [BBV15] Maria-Florina Balcan, Avrim Blum, and Santosh Vempala. Efficient representations for lifelong learning and autoencoding. In *Conference on Learning Theory*, pages 191–210. PMLR, 2015.
- [BDS20] Mehdi Abbana Bennani, Thang Doan, and Masashi Sugiyama. Generalisation guarantees for continual learning with orthogonal gradient descent. *arXiv preprint arXiv:2006.11942*, 2020.
- [BHLM19] Peter L Bartlett, Nick Harvey, Christopher Liaw, and Abbas Mehrabian. Nearly-tight vc-dimension and pseudodimension bounds for piecewise linear neural networks. *J. Mach. Learn. Res.*, 20(63):1–17, 2019.
- [BKT19] Maria-Florina Balcan, Mikhail Khodak, and Ameet Talwalkar. Provable guarantees for gradient-based meta-learning. In *International Conference on Machine Learning*, pages 424–433. PMLR, 2019.
- [BL13] Maria-Florina Balcan and Phil Long. Active and passive learning of linear separators under log-concave distributions. In *Conference on Learning Theory*, pages 288–316. PMLR, 2013.
- [CLL21] Kurtland Chua, Qi Lei, and Jason D Lee. How fine-tuning allows for effective meta-learning. *arXiv preprint arXiv:2105.02221*, 2021.
- [CMJG⁺18] Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *ECCV*, 2018.
- [DBM⁺21] Thang Doan, Mehdi Abbana Bennani, Bogdan Mazoure, Guillaume Rabusseau, and Pierre Alquier. A theoretical analysis of catastrophic forgetting through the ntk overlap matrix. In *International Conference on Artificial Intelligence and Statistics*, pages 1072–1080. PMLR, 2021.
- [DHK⁺20] Simon S Du, Wei Hu, Sham M Kakade, Jason D Lee, and Qi Lei. Few-shot learning via learning the representation, provably. *arXiv preprint arXiv:2002.09434*, 2020.
- [DPC20] Giulia Denevi, Massimiliano Pontil, and Carlo Ciliberto. The advantage of conditional meta-learning for biased regularization and fine-tuning. *arXiv preprint arXiv:2008.10857*, 2020.
- [DPC21] Giulia Denevi, Massimiliano Pontil, and Carlo Ciliberto. Conditional meta-learning of linear representations. *arXiv preprint arXiv:2103.16277*, 2021.
- [DSCP19] Giulia Denevi, Dimitris Stamos, Carlo Ciliberto, and Massimiliano Pontil. Online-within-online meta-learning. In *ADVANCES IN NEURAL INFORMATION PROCESSING SYSTEMS 32 (NIPS 2019)*, volume 32, pages 1–11. Neural Information Processing Systems (NeurIPS 2019), 2019.

- [FAL17] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135. PMLR, 2017.
- [FAML20] Mehrdad Farajtabar, Navid Azizan, Alex Mott, and Ang Li. Orthogonal gradient descent for continual learning. In *International Conference on Artificial Intelligence and Statistics*, pages 3762–3773. PMLR, 2020.
- [Fin92] Helmut Finner. A generalization of holder’s inequality and some probability inequalities. *The Annals of probability*, pages 1893–1901, 1992.
- [HPL⁺19] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *CVPR*, 2019.
- [HZRS16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [JGH18] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *arXiv preprint arXiv:1806.07572*, 2018.
- [KBT19] Mikhail Khodak, Maria-Florina Balcan, and Ameet Talwalkar. Adaptive gradient-based meta-learning methods. *arXiv preprint arXiv:1906.02717*, 2019.
- [KPR⁺17] 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.
- [LH17] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017.
- [LLL⁺21] Weiyang Liu, Rongmei Lin, Zhen Liu, James M Rehg, Liam Paull, Li Xiong, Le Song, and Adrian Weller. Orthogonal over-parameterized training. In *CVPR*, 2021.
- [LLRS19] Weiyang Liu, Zhen Liu, James Rehg, and Le Song. Neural similarity learning. In *NeurIPS*, 2019.
- [LV07] László Lovász and Santosh Vempala. The geometry of logconcave functions and sampling algorithms. *Random Structures & Algorithms*, 30(3):307–358, 2007.
- [MC89] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989.
- [PKP⁺19] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019.
- [Rat90] Roger Ratcliff. Connectionist models of recognition memory: constraints imposed by learning and forgetting functions. *Psychological review*, 97(2):285, 1990.
- [RHK⁺19] Jathushan Rajasegaran, Munawar Hayat, Salman H Khan, Fahad Shahbaz Khan, and Ling Shao. Random path selection for continual learning. In *NeurIPS*, 2019.

- [RKSL17] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017.
- [RRD⁺16] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016.
- [TJJ20] Nilesch Tripuraneni, Chi Jin, and Michael I Jordan. Provable meta-learning of linear representations. *arXiv preprint arXiv:2002.11684*, 2020.
- [TM95] Sebastian Thrun and Tom M Mitchell. Lifelong robot learning. *Robotics and autonomous systems*, 15(1-2):25–46, 1995.
- [WDC20] Ruohan Wang, Yiannis Demiris, and Carlo Ciliberto. A structured prediction approach for conditional meta-learning. *Advances in Neural Information Processing Systems*, 2020.
- [Wel09] Max Welling. Herding dynamical weights to learn. In *ICML*, 2009.
- [XZ18] Ju Xu and Zhanxing Zhu. Reinforced continual learning. *arXiv preprint arXiv:1805.12369*, 2018.
- [YFL⁺20] Dong Yin, Mehrdad Farajtabar, Ang Li, Nir Levine, and Alex Mott. Optimization and generalization of regularization-based continual learning: a loss approximation viewpoint. *arXiv preprint arXiv:2006.10974*, 2020.
- [YXH21] Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *CVPR*, 2021.
- [YYLH17] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. *arXiv preprint arXiv:1708.01547*, 2017.
- [ZXG⁺20] Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shu-Tao Xia. Maintaining discrimination and fairness in class incremental learning. In *CVPR*, 2020.
- [ZYK⁺18] Xu Zhang, Felix Xinnan Yu, Svebor Karaman, Wei Zhang, and Shih-Fu Chang. Heated-up softmax embedding. *arXiv preprint arXiv:1809.04157*, 2018.

Appendix

A Formal Algorithm of LLL-RR

For completeness, we describe the formal algorithm of LLL-RR. Different from the basic LLL algorithm, we are memorizing a list $\tilde{\mathbf{w}}_1, \dots, \tilde{\mathbf{w}}_{\hat{k}k_0}$ all along with the algorithm. Each time when we need to learn the new features $\tilde{\mathbf{w}}_{\hat{k}k_0+1}, \dots, \tilde{\mathbf{w}}_{\hat{k}k_0+k_0}$, we add them to the list, and feed the list to Algorithm 2 to get a new feature subspace. The formal algorithm is in Algorithm 4.

Algorithm 4 Lifelong Learning Algorithm with Representation Refinement (LLL-RR)

Input: d, m, k , labeled examples of m tasks, threshold parameters ϵ_{acc}, ϵ .

1. Using data from the first task to learn a set of features $\tilde{\mathbf{W}}_1(\cdot) = (\tilde{\mathbf{w}}_1(\cdot), \dots, \tilde{\mathbf{w}}_{k_0}(\cdot))^\top$ and a linear function $\tilde{\mathbf{c}}_1$ such that $\mathbf{x} \rightarrow \text{sign}(\tilde{\mathbf{c}}_1^\top \tilde{\mathbf{W}}_1(\mathbf{x}))$ has error smaller than ϵ_{acc} .
/* Number of features $1 \leq k_0 \leq k$. For linear features, $k_0 = 1$. */
Let $\tilde{k} = 1$. Set the feature subspace $\mathbf{V}_1 = \tilde{\mathbf{W}}_1$, and the temporary features $\tilde{\mathbf{V}}_1 = \tilde{\mathbf{W}}_1$.
2. For the task $i = 2, \dots, m$
 - Using the data from the i task, attempt to learn the linear function $\tilde{\mathbf{c}}_i$ using the temporary features $\tilde{\mathbf{V}}_{i-1}$.
 - Check whether $\mathbf{x} \rightarrow \text{sign}(\tilde{\mathbf{c}}_i^\top \tilde{\mathbf{V}}_{i-1}(\mathbf{x}))$ has error less than ϵ .
 - (a) If yes, set $\tilde{\mathbf{V}}_i = \tilde{\mathbf{V}}_{i-1}$. // Small error with current features.
 - (b) Otherwise, learn a new set of features $\tilde{\mathbf{W}}_i(\cdot)$ and a linear function $\tilde{\mathbf{c}}_i$ such that the predictor $\mathbf{x} \rightarrow \text{sign}(\tilde{\mathbf{c}}_i^\top \tilde{\mathbf{W}}_i(\mathbf{x}))$ has error less than ϵ_{acc} .
Update the feature subspace $\mathbf{V}_i = (\mathbf{V}_i; \tilde{\mathbf{W}}_i)$, and feed into Algorithm 2. It returns the refined subspace \mathbf{V}' . Set the temporary features $\tilde{\mathbf{V}}_i = \mathbf{V}'$. Let $\hat{k} = \hat{k} + 1$.

return m predictors: $\mathbf{x} \rightarrow \text{sign}(\tilde{\mathbf{c}}_i^\top \tilde{\mathbf{V}}_i(\mathbf{x}))$, $1 \leq i \leq m$.

B Extensions to Task-Incremental Regression and Class-Incremental Learning

In our main text, we study the setting of solving m tasks of binary classification incrementally. The classification error is defined as $\text{err}(\hat{l}) = \mathbb{P}_{(\mathbf{x}, y) \sim P}[\hat{l}(\mathbf{x}) \neq y]$. By Assumption 1, we know the task error is small if and only if the parameters are close to each other. Now we would like to extend to task-incremental regression and class-incremental classification by connecting the parameter's l_2 distance to the error of the model.

Task-incremental regression. Consider the regression tasks shared with low-dimensional common features. For $i \in [m]$, $y = \langle \mathbf{c}_i^*, \boldsymbol{\sigma}^*(\mathbf{x}) \rangle + \epsilon_i$. The regression error is $\text{err}(\hat{l}) = \mathbb{E}_{(\mathbf{x}, y) \sim P}[\|\hat{l}(\mathbf{x}) - y\|_2^2]$. We can further weaken our assumption to $c_1 I \preceq \mathbb{E}[\mathbf{x}\mathbf{x}^\top] \preceq c_2 I$ for $0 < c_1 < c_2$. Then for any unit vector \mathbf{u}, \mathbf{v} , we have

$$\mathbb{E}_{(\mathbf{x}, y) \sim P} \left[\left(\mathbf{u}^\top \mathbf{x} - \mathbf{v}^\top \mathbf{x} \right)^2 \right] = \mathbb{E}_{(\mathbf{x}, y) \sim P} \left[(\mathbf{u} - \mathbf{v})^\top \mathbf{x} \mathbf{x}^\top (\mathbf{u} - \mathbf{v}) \right]$$

So we get a lemma similar to Assumption 1 that $c_1 \|\mathbf{u} - \mathbf{v}\|^2 \leq \mathbb{E}_{\mathbf{x} \in D} \|\mathbf{u}^\top \mathbf{x} - \mathbf{v}^\top \mathbf{x}\|^2 \leq c_2 \|\mathbf{u} - \mathbf{v}\|^2$, which is sufficient for analysis.

Class-incremental classification. Let $X = \mathbb{R}^d$ be the input space and $Y = \{1, 2, \dots, m\}$ be the class labels. We assume that the labels can be recovered by passing the input through a linear/nonlinear layer and then taking the maximum of m linear combinations. Formally, the label is given by

$$\ell(\mathbf{x}) = \operatorname{argmax}_{i \in [m]} \langle \mathbf{c}_i^*, \boldsymbol{\sigma}^*(\mathbf{x}) \rangle$$

The classification error is $\operatorname{err}(\hat{l}) = \mathbb{P}_{(\mathbf{x}, y) \sim P} [\hat{l}(\mathbf{x}) \neq y]$. Noticing that, when we meet a new class, the classifier should determine whether the label belongs to the current class or not. In this sense, we regard the problem as a binary classification. To get the negative samples in the current class, we also need a small proportional of data from previous classes. Practically, we propose the heuristic lifelong learning (H-LLL) algorithm in Section 2.3 to solve the class-incremental learning. Experiments in Section 5.2 complement our results.

C Another Approach – Theoretical Guarantees for LLL

Here we give a simpler analysis for the basic LLL algorithm, along the lines of [BBV15]. The result is weaker than Theorem 3, but we include the proof here for completeness, along with an extension to nonlinear features.

Theorem 7 (Basic LLL). *Let $\gamma = c\epsilon$ and ϵ_{acc} s.t. $4k \frac{\epsilon_{acc}}{\gamma} + \gamma = c'\epsilon$ for sufficiently small constants $c, c' > 0$. Assume that all targets share k common features. Then, under Assumption 1, and sequential presentation of the tasks in any order, the basic LLL algorithm will incrementally learn a representation of dimension k for linear features and k^2 for nonlinear features with error at most ϵ on all tasks. The total number of samples used by the algorithm is $O(dk^2 \log(k/\epsilon)/\epsilon^2 + km \log(1/\epsilon)/\epsilon) = \tilde{O}(dk^2/\epsilon^2 + mk/\epsilon)$ in the linear setting and a factor of k higher in the nonlinear setting.*

Before we prove the theorem, we define the γ -separated term. We use the definition γ -separated from [BBV15] that a subsequence of vectors $\mathbf{a}_{i_1}, \mathbf{a}_{i_2}, \dots$ is γ -separated if for any \mathbf{a}_{i_j} , $\theta(\mathbf{a}_{i_j}, \operatorname{span}(\mathbf{a}_{i_1}, \dots, \mathbf{a}_{i_{j-1}})) \geq \gamma$. Define the γ -effective dimension of $\mathbf{a}_1, \dots, \mathbf{a}_m$ as the size of the largest γ -separated subsequence. Note that when $\gamma = 0$, γ -effective dimension is exactly the dimension of the spanned subspace. We prove Theorem 7 by showing two facts: each target is far from the span of previous ones; we learn the new target accurately. (Lemma 11). We start with a helper lemma (Lemma 10).

Lemma 10. *Let \mathbf{w}, \mathbf{v} be two unit vectors in \mathbb{R}^d and \mathbf{U} be a subspace. Then,*

$$\sin \theta(\operatorname{span}(\mathbf{U}, \mathbf{w}), \operatorname{span}(\mathbf{U}, \mathbf{v})) \leq \frac{\sin \theta(\mathbf{w}, \mathbf{v})}{\max\{\sin \theta(\mathbf{w}, \mathbf{U}), \sin \theta(\mathbf{v}, \mathbf{U})\}}.$$

Proof. If $\mathbf{w} \in \mathbf{U}$ or $\mathbf{v} \in \mathbf{U}$, $\theta(\operatorname{span}(\mathbf{U}, \mathbf{w}), \operatorname{span}(\mathbf{U}, \mathbf{v})) = 0$. The inequality becomes trivial. Now we assume that $\mathbf{w}, \mathbf{v} \notin \mathbf{U}$. From the symmetry of \mathbf{w} and \mathbf{v} , we prove the following and then replacing \mathbf{v} with \mathbf{w} leads to the original inequality.

$$\sin \theta(\operatorname{span}(\mathbf{U}, \mathbf{w}), \operatorname{span}(\mathbf{U}, \mathbf{v})) \leq \frac{\sin \theta(\mathbf{w}, \mathbf{v})}{\sin \theta(\mathbf{v}, \mathbf{U})}$$

By definition, $\exists \mathbf{x} \in \text{span}(\mathbf{U}, \mathbf{w})$ s.t. $\theta(\mathbf{x}, \text{span}(\mathbf{U}, \mathbf{v})) = \theta(\text{span}(\mathbf{U}, \mathbf{w}), \text{span}(\mathbf{U}, \mathbf{v}))$. Here \mathbf{x} is a combination of \mathbf{u}_1 and \mathbf{w} , \mathbf{u}_1 is some vector in \mathbf{U} . Using the fact that $\theta(\mathbf{x}, \text{span}(\mathbf{U}, \mathbf{v})) \leq \theta(\mathbf{x}, \text{span}(\mathbf{u}_1, \mathbf{v})) \leq \theta(\text{span}(\mathbf{u}_1, \mathbf{w}), \text{span}(\mathbf{u}_1, \mathbf{v}))$ and $\theta(\mathbf{v}, \mathbf{U}) \leq \theta(\mathbf{v}, \mathbf{u}_1)$, it is sufficient to prove that

$$\sin \theta(\text{span}(\mathbf{u}_1, \mathbf{w}), \text{span}(\mathbf{u}_1, \mathbf{v})) \leq \frac{\sin \theta(\mathbf{w}, \mathbf{v})}{\sin \theta(\mathbf{v}, \mathbf{u}_1)}$$

Denote $\alpha = \theta(\text{span}(\mathbf{u}_1, \mathbf{w}), \text{span}(\mathbf{u}_1, \mathbf{v}))$, $\beta = \theta(\mathbf{v}, \mathbf{u}_1)$. WLOG we assume $\mathbf{u}_1 = (1, 0, 0)$ and $\text{span}(\mathbf{u}_1, \mathbf{w})$ is the x - y plane. Then we can write $\mathbf{v} = \cos(\beta)\mathbf{u}_1 + \sin(\beta)\mathbf{v}_1$, where $\mathbf{v}_1 = (0, \cos(\alpha), \sin(\alpha))$. Since $\sin \theta(\mathbf{w}, \mathbf{v}) \geq d(\mathbf{v}, x\text{-}y \text{ plane}) = \sin(\alpha) \sin(\beta)$, we get the lemma proved. \square

Lemma 11 (Kernel Subspace). *Let $\mathbf{U}_k, \mathbf{V}_k$ be two subspaces of \mathbb{R}^d . Let $\mathbf{U}_k = \text{span}\{\mathbf{y}_1^*, \dots, \mathbf{y}_k^*\}$, $\mathbf{V}_k = \text{span}\{\mathbf{y}_1, \dots, \mathbf{y}_k\}$. Let $\epsilon, \gamma \geq 0$ and $\epsilon \leq \gamma^2/(10k)$. Assume that*

1. $\sin \theta(\mathbf{y}_i, \text{span}\{\mathbf{y}_1, \dots, \mathbf{y}_{i-1}\}) \geq \gamma$, for $i = 2, \dots, k$.
2. $\sin \theta(\mathbf{y}_i, \mathbf{y}_i^*) \leq \epsilon$, for $i = 1, \dots, k$.

Then we have $\sin \theta(\mathbf{U}_k, \mathbf{V}_k) \leq 2k\epsilon/\gamma$. In other words, for any point $\mathbf{y}^ \in \mathbf{U}_k$, there is a point $\mathbf{y} \in \mathbf{V}_k$ s.t.*

$$\sin \theta(\mathbf{y}^*, \mathbf{y}) \leq \frac{2k\epsilon}{\gamma}.$$

Proof. Here we use the strong induction on a stronger version of the conclusion where $\mathbf{U}_k = \text{span}\{\mathbf{W}, \mathbf{y}_1^*, \dots, \mathbf{y}_k^*\}$, $\mathbf{V}_k = \text{span}\{\mathbf{W}, \mathbf{y}_1, \dots, \mathbf{y}_k\}$ for some fixed subspace \mathbf{W} . The base case is $k = 1$. This follows directly from Lemma 10 with $\mathbf{U} = \mathbf{W}$, $\mathbf{w} = \mathbf{y}$, $\mathbf{v} = \mathbf{y}^*$. Now we prove the induction step on k with strong hypothesis. Let $\mathbf{U}'_k = \text{span}(\mathbf{U}_{k-1}, \mathbf{y}_k)$. By Lemma 10 and induction hypothesis, we have

$$\sin \theta(\mathbf{U}_k, \mathbf{V}_k) \leq \sin \theta(\mathbf{U}_k, \mathbf{U}'_k) + \sin \theta(\mathbf{U}'_k, \mathbf{V}_k) \leq \frac{\sin \theta(\mathbf{y}_k, \mathbf{y}_k^*)}{\sin \theta(\mathbf{y}_k, \mathbf{U}_{k-1})} + \frac{2(k-1)\epsilon}{\gamma}$$

By triangle inequality and induction hypothesis, we further have

$$\sin \theta(\mathbf{y}_k, \mathbf{U}_{k-1}) \geq \sin \theta(\mathbf{y}_k, \mathbf{V}_{k-1}) - \sin \theta(\mathbf{V}_{k-1}, \mathbf{U}_{k-1}) \geq \gamma - \frac{2\epsilon(k-1)}{\gamma}$$

Combining the two inequalities, we get

$$\sin \theta(\mathbf{U}_k, \mathbf{V}_k) \leq \frac{\epsilon}{\gamma - \frac{2\epsilon(k-1)}{\gamma}} + \frac{2(k-1)\epsilon}{\gamma} = \frac{\epsilon}{\gamma} \left(\frac{\gamma^2}{\gamma^2 - 2\epsilon(k-1)} + 2(k-1) \right) \leq \frac{2k\epsilon}{\gamma}$$

\square

Now we put them together to analyze Algorithm 1.

Proof. We consider the kernel of nonlinear features $\boldsymbol{\sigma}(\mathbf{x})$. These features live in a potentially infinite-dimensional space (or exponential in d dimensional space if, e.g., the input is from the Boolean hypercube). Let \mathbf{U} be the span of the nonlinear features (viewed as vectors) in the model used to label data. The γ -effective dimension of \mathbf{U} is at most k . Let $\mathbf{y}_i^* = \mathbf{c}_i^{*\top} \boldsymbol{\sigma}^*(\mathbf{x}) = \mathbf{a}_i^*(\mathbf{x})$, $y_i = \mathbf{c}_i^\top \boldsymbol{\sigma}(\mathbf{x}) = \mathbf{a}_i(\mathbf{x})$. WLOG let's assume $\mathbf{a}_i, \mathbf{a}_i^*$ be vectors of unit length. From the algorithm, if the current task i has already achieved ϵ error by current features, it's done. Otherwise, we learn a new set of k_0 features and

a linear combination whose error is at most ϵ_{acc} . Denote the indices of tasks that we learn new features as $i_1, i_2, \dots, i_{\hat{k}}$. Encountering the task $i_{\hat{k}}$, denote $\mathbf{V}_{\hat{k}} = \text{span}(\mathbf{a}_{i_1}, \dots, \mathbf{a}_{i_{\hat{k}}})$, $\mathbf{U}_{\hat{k}} = \text{span}(\mathbf{a}_{i_1}^*, \dots, \mathbf{a}_{i_{\hat{k}}}^*)$. We will prove by induction that for any $\hat{k} \in [\tilde{k}]$, (1) $\theta(\mathbf{a}_{i_{\hat{k}}}, \mathbf{V}_{\hat{k}-1}) \geq \gamma$; (2) $\theta(\mathbf{a}_{i_{\hat{k}}}^*, \mathbf{U}_{\hat{k}-1}) \geq \gamma$.

The base case $\hat{k} = 1$ holds immediately. For the inductive step $\hat{k} > 1$, the task $i_{\hat{k}}$ cannot achieve ϵ error with the current features $\mathbf{V}_{\hat{k}-1}$. By Assumption 1, $\theta(\mathbf{a}_{i_{\hat{k}}}^*, \mathbf{V}_{\hat{k}-1}) \geq \epsilon/c_2$. After learning a new set of features to ensure error less than ϵ_{acc} , we know there is a new linear combination $\mathbf{a}_{i_{\hat{k}}}$ such that $\theta(\mathbf{a}_{i_{\hat{k}}}^*, \mathbf{a}_{i_{\hat{k}}}) \leq \epsilon_{acc}/c_1$. So by triangle inequality,

$$\theta(\mathbf{a}_{i_{\hat{k}}}, \mathbf{V}_{\hat{k}-1}) \geq \epsilon/c_2 - \epsilon_{acc}/c_1 \geq \gamma$$

So we have shown that (1) holds for $i_{\hat{k}}$.

To prove (2), we suppose for contradiction that $\theta(\mathbf{a}_{i_{\hat{k}}}^*, \mathbf{U}_{\hat{k}-1}) < \gamma$. From induction hypothesis, for any $j \in [\hat{k}-1]$, $\sin \theta(\mathbf{a}_{i_j}, \mathbf{V}_{j-1}) \geq \gamma/2$. By construction, we also have for any $\sin \theta(\mathbf{a}_{i_j}, \mathbf{a}_{i_j}^*) \leq \epsilon_{acc}/c_1$. Apply Lemma 11, we have $\theta(\mathbf{U}_{\hat{k}-1}, \mathbf{V}_{\hat{k}-1}) \leq 8\epsilon_{acc}k/\gamma$. By triangle inequality, we further have

$$\theta(\mathbf{a}_{i_{\hat{k}}}^*, \mathbf{V}_{\hat{k}-1}) \leq \theta(\mathbf{a}_{i_{\hat{k}}}^*, \mathbf{U}_{\hat{k}-1}) + \theta(\mathbf{U}_{\hat{k}-1}, \mathbf{V}_{\hat{k}-1}) \leq \gamma + 4\epsilon_{acc}k/(c_1\gamma) \leq \epsilon/c_2$$

By Assumption 1, there exists $\mathbf{b}_{i_{\hat{k}}} \in \mathbf{V}_{\hat{k}-1}$ with error less than ϵ , and thus leads to contradiction. So (2) is also proved. Furthermore, since we have assume that the γ -effective dimension of the true targets is at most k , we have $\tilde{k} \leq k$. So the size of the internal representation is $k' = O(kk_0)$.

The sample complexity for learning one task in d -dimension up to error ϵ is $O(dk_0 \log(1/\epsilon)/\epsilon)$. Here we learn $O(k)$ such tasks. All other tasks can be learned using the features of dimension $O(kk_0)$. Therefore the total sample complexity is $O(dkk_0/\epsilon_{acc} \log(1/\epsilon_{acc}) + kk_0m \log(1/\epsilon)/\epsilon) = \tilde{O}(dk^2k_0/\epsilon^2 + kk_0m/\epsilon)$.

□