

Continual Learning by Asymmetric Loss Approximation with Single-Side Overestimation

Dongmin Park^{1,2}Seokil Hong¹Bohyung Han¹Kyoung Mu Lee¹¹ECE & ASRI, Seoul National University, Korea²Samsung Electronics, Korea

dmpark04@gmail.com, {hongceo96, bhhan, kyoungmu}@snu.ac.kr

Abstract

Catastrophic forgetting is a critical challenge in training deep neural networks. Although continual learning has been investigated as a countermeasure to the problem, it often suffers from the requirements of additional network components and the limited scalability to a large number of tasks. We propose a novel approach to continual learning by approximating a true loss function using an asymmetric quadratic function with one of its sides overestimated. Our algorithm is motivated by the empirical observation that the network parameter updates affect the target loss functions asymmetrically. In the proposed continual learning framework, we estimate an asymmetric loss function for the tasks considered in the past through a proper overestimation of its unobserved sides in training new tasks, while deriving the accurate model parameter for the observable sides. In contrast to existing approaches, our method is free from the side effects and achieves the state-of-the-art accuracy that is even close to the upper-bound performance on several challenging benchmark datasets.

1. Introduction

It is common to learn machine learning models versatile for multiple tasks in an incremental manner when new tasks are given one by one, not in a batch. Continual learning is a concept to learn a model for a large number of tasks sequentially without forgetting knowledge obtained from the preceding tasks, where the data in the old tasks are not available any more during training new ones.

Catastrophic forgetting is a critical issue in realizing continual learning using deep neural networks. With a naïve stochastic gradient descent (SGD) method, deep neural networks easily forget the knowledge obtained from the earlier tasks while adapting to the new information quickly from the incoming tasks [26, 8]. This is mainly because, without

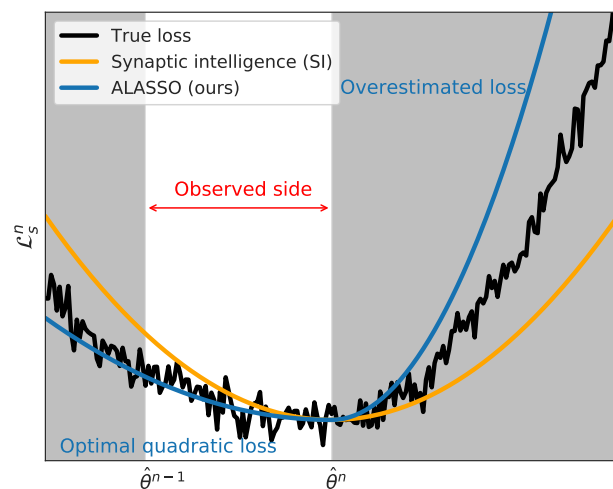


Figure 1. Conceptual diagram to illustrate why our loss approximation is required. The symmetric loss approximations such as the quadratic approximation used in SI [37] (yellow) is prone to underestimate the unobserved sides of potentially asymmetric real surrogate loss functions (black). We claim that properly introduced asymmetry in the loss approximation (blue) prevents this problem. Note that the other side are observed during the optimization of the network in training the previous tasks and we can derive a more accurate parameter for the quadratic optimization of the surrogate function.

any countermeasure, the optimization based on the losses of new tasks may not be desirable to retain the knowledge about the tasks learned in the past. The catastrophic forgetting problem limits the capability and potential of deep neural networks to be applied to dynamic real-world problems that require continuous adaptation to new environments. The continual learning or life-long learning is a well-known framework to handle the catastrophic forgetting problem. It can be categorized into three categories: architectural approach, functional approach, and structural regularization approach. The architectural and functional approaches typically need additional network components

and/or batch processing. The structural regularization methods work well for a limited number of tasks, but often have scalability issues to many tasks.

This paper presents a novel continual learning framework based on asymmetric loss approximation with single-side overestimation (ALASSO), which effectively adapts to a large number of tasks. ALASSO approximates the true loss functions corresponding to the previously considered tasks asymmetrically by overestimating their unobserved sides in the parameter space while deriving the accurate quadratic approximation on the observed sides. Figure 1 illustrates the main concept of our approach; it computes the optimal parameter through a quadratic approximation in the observed side (left) while using a steep surrogate quadratic function in the unobservable side (right). The proposed algorithm also decouples the hyperparameters for the current surrogate loss approximation and the surrogate loss change of the previous tasks. This approach is motivated by our observation that updating the model parameters of deep neural networks affects target losses asymmetrically and that using the overestimated loss functions is relatively safe for the optimization without the training data of the previous tasks.

In contrast to the existing approaches, the proposed technique is free from the additional memory requirement to store the information about the previous tasks and the overhead of batch processing. ALASSO achieves the state-of-the-art performance and is even close to the accuracy upperbounds in several challenging benchmark datasets, including the permuted MNIST, the split CIFAR-10/CIFAR-100 and the split Tiny ImageNet. In particular, we demonstrate promising results in terms of scalability and robustness to a large number of tasks.

The contribution of our work is summarized as follows:

- We propose a novel continual learning framework by overestimating the unobserved side of a loss function in the current task and approximating the loss using an asymmetric quadratic function. This strategy facilitates a reliable loss estimation even without the training data of the previous tasks.
- Our algorithm provides an accurate solution of the loss approximation for the previous tasks, which allows to derive the best approximation of a quadratic surrogate loss function on the observed side.
- The proposed technique achieves the outstanding performance on several challenging benchmark datasets by large margins.

The rest of the paper is organized as follows. We first discuss the related works in Section 2, and present the technical details of our algorithm, ALASSO, in Section 4. Section 5 demonstrates the experimental results with analysis and Section 6 concludes this paper.

2. Related Work

Continual learning algorithms are categorized into three groups [37]: architectural, functional, and structural regularization approaches. This section discusses the existing methods in individual categories and their characteristics briefly.

2.1. Architectural Approaches

The approaches in this category realize continual learning by freezing model parameters learned from the previous tasks and/or providing limited architectural variations for learning the new tasks [30, 22, 3, 20]. This framework often requires additional network components for the new tasks and the size of the network gradually increases in principle. This drawback limits the applicability to large-scale problems and is not appropriate for the configurations with limited resources such as embedded systems. Existing methods [24, 34, 15, 35, 13] often keep track of the data in the previous tasks by allocating additional memory space called episodic memory or by generating training examples using generative adversarial networks [11], but they suffer from substantial memory demands to store the information of the previous tasks. Some approaches in this category aim to improve performance using different nonlinearities such as ReLU, MaxOut, and local winner-take-all [36, 10].

An interesting work among the architectural approaches, instead of adding new network components, employs a network compression technique to identify unused or rarely used parts in the target network and allows the free space to store the information of new tasks [25, 12]. It alleviates the drawbacks of the standard approaches discussed earlier, but requires the substantial overhead of batch processing for network compression.

2.2. Functional Approaches

The functional approaches [23, 14, 32, 28] often incorporate knowledge distillation to establish continual learning. The previously learned networks are fixed and used for feature computation in training new tasks. The new networks are encouraged to learn the representations that are coherent to the features computed in the previous networks. However, the feature coherency of new examples with respect to the old ones does not always ensure the output similarity of the examples in the past and current tasks. Existing functional approaches are required to store and evaluate the previous networks for training the model for the new task, which incurs additional computational overhead.

2.3. Structural Regularization Approaches

The structural regularization approaches typically augment a penalty term to the original loss functions and discourage the updates of parameters critical to the previously

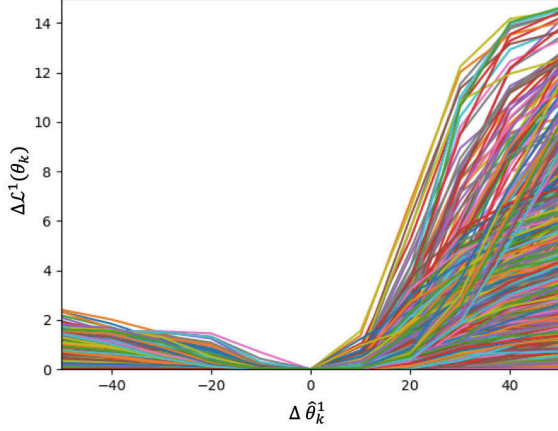


Figure 2. Illustration of asymmetric characteristics of a loss function in terms of each model parameter $\hat{\theta}_k$, where k is the index of model parameter. Each graph shows how much loss \mathcal{L}^1 changes by updating each parameter $\hat{\theta}_k^1$ after learning the first task in the permuted MNIST dataset.

learned tasks. Elastic weight consolidation (EWC) [17] and synaptic intelligence (SI) [37] employ a surrogate quadratic loss as an approximation of the real loss functions of the previous tasks. Although this is simple and effective for a small number of tasks, its performance drops drastically when there are a larger number of tasks. Memory aware synapses (MAS) [2] estimates the importance of the weights in an unsupervised manner. Incremental moment matching (IMM) [21] additionally performs a separate model-merging step after learning a new task. Variational continual learning (VCL) [27] combines the online variational inference [9, 31, 5] with Monte Carlo sampling [4] using neural networks, but it requires a relatively large amount of computational cost to infer an approximate posterior distribution. On the other hand, [33] proposes a task-specific hard attention mechanism, but, as in [30, 7, 23], the requirement of multi-head outputs—a separate output per task—limits the number of target domains considered concurrently.

Our framework falls in the structural regularization category. We design a new loss function appropriate for continual learning while it does not incur extra computational overhead such as additional network components and the need for occasional batch processing.

3. Synaptic Intelligence (SI)

The proposed method, referred to as ALASSO, is closely related to synaptic intelligence (SI) [37]. We first discuss the main idea of SI and then point out its critical problems before presenting our novel idea.

3.1. Quadratic Approximation of Loss

SI is categorized as a structural regularization method, which employs a static network architecture and does not

use additional memory throughout continual learning process. To prevent the catastrophic forgetting problem and maintain the performance with respect to the previous tasks while adapting to a new task n , this technique introduces a surrogate loss function \mathcal{L}_s^{n-1} , which approximates the loss of the previous tasks and plays a role as a regularizer. Assuming that the surrogate loss is a quadratic function, the total loss function to learn the n^{th} task, $\tilde{\mathcal{L}}^n$, in terms of the model parameter θ_k is approximated by

$$\begin{aligned}\tilde{\mathcal{L}}^n &= \mathcal{L}^n + c \mathcal{L}_s^{n-1} \\ &= \mathcal{L}^n + c \underbrace{\sum_k \hat{\Omega}_k^{n-1} (\theta_k - \hat{\theta}_k^{n-1})^2}_{\text{surrogate loss}},\end{aligned}\quad (1)$$

where \mathcal{L}^n denotes a loss for the current task n , $\hat{\theta}_k^{n-1}$ is the weight in the k^{th} dimension of the estimated parameter for the previous tasks until $n-1$, $\hat{\Omega}_k^{n-1}$ is a coefficient for the corresponding model parameter, and c is a hyperparameter for the surrogate loss.

We minimize the total loss, $\tilde{\mathcal{L}}^n$, with respect to θ_k , and obtain the optimized parameters, $\hat{\theta}_k^n$. Assuming that the surrogate loss function for each parameter up to the n^{th} task is defined by a quadratic function as

$$\mathcal{L}_s^n(\theta_k) = \hat{\Omega}_k^n (\theta_k - \hat{\theta}_k^n)^2, \quad (2)$$

$\hat{\Omega}_k^n$ is derived approximately by

$$\hat{\Omega}_k^n \approx \frac{\omega_k^n}{(\hat{\theta}_k^n - \hat{\theta}_k^{n-1})^2} + \hat{\Omega}_k^{n-1}, \quad (3)$$

where ω_k^n denotes the difference between the losses before and after training a new task n , *i.e.*, $\omega_k^n = \mathcal{L}^n(\hat{\theta}_k^{n-1}) - \mathcal{L}^n(\hat{\theta}_k^n)$. Refer to [37] for more details about SI.

3.2. Underestimation of Loss

SI employs a symmetric loss function for approximation with respect to the previous tasks while it can observe only a single side of the symmetric function along the trajectory of parameter update from $\hat{\theta}_k^{n-1}$ to $\hat{\theta}_k^n$ during the optimization process for the n^{th} task. In other words, the surrogate loss function is assumed to be symmetric as shown in Eq. (1) and may not be able to model its unobserved side accurately unless the true loss function is symmetric.

However, it turns out that the true loss functions are typically asymmetric and the symmetric loss functions adopted by SI are prone to underestimate the true losses in practice. Figure 2 illustrates the variation of the losses with respect to the changes of each model parameter; x -axis denotes the offset from the optimal model parameter and y -axis is the difference of the cross-entropy losses, which are obtained

from the first task of the permuted MNIST dataset. Note that, to visualize the asymmetry more clearly, the steeper halves of the individual graphs are located on the right hand side. Based on this observation, we consider an asymmetric loss function formulation for a more accurate and stable estimation of the surrogate loss function parametrized by $\hat{\Omega}_k^n$.

4. Proposed Algorithm

We propose a novel structure regularizer, which mitigates the limitations of SI. The main contributions of our algorithm, ALASSO, include the introduction of the asymmetric loss function with single-side overestimation and the more accurate quadratic approximation of the loss function, which lead to remarkable performance improvement. This section presents the details about ALASSO, especially, in comparison to the existing approach, SI.

4.1. Overview

Our algorithm overestimates the unobserved sides of the approximate loss function and allows the models to learn under a harsher condition. It also derives the accurate parameter estimation of the approximate quadratic loss functions on their observed sides. To accelerate the optimization procedure and handle the conflicts between the loss computation of the current task and the loss approximation of the previous tasks, we introduce a hyperparameter decoupling technique although the values of the decoupled hyperparameters should be identical conceptually. The proposed algorithm inherits the merits of the standard structural regularization approaches such as online learning and local updates while providing the capability to maintain the crucial knowledge about the prior tasks by making the models less adaptive without the observation of the loss function. We now discuss the technical contributions and characteristics of our continual learning framework, ALASSO.

4.2. Asymmetric Loss Approximation

One possible option for a better structural regularizer in continual learning is asymmetric loss approximation of the previous tasks. We believe that the symmetric regularizer as in SI overly simplifies the true loss functions and have a critical limitation in maintaining the knowledge obtained from the previous tasks. Figure 1 illustrates why asymmetric loss approximation is effective. The approximate quadratic loss functions may be sufficiently accurate on the sides, where the true losses are observable along the model parameter updates during training. However, they may incur substantial error on their unobserved sides, so it is dangerous to assume that the true loss functions are symmetric, which is supported by Figure 2.

Based on this motivation, we now propose a simple but effective approximation approach of the true loss functions.

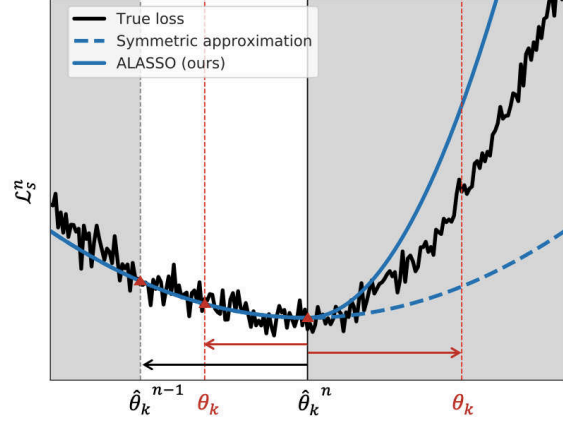


Figure 3. Visualization of quadratic surrogate loss functions based on the two piece-wise approximation of the true loss. When learning the n^{th} task, θ_k moves by a gradient descent method and the loss approximation is performed differently depending on the location of θ_k —whether it is on the observed side or not. Note that the area between $\hat{\theta}_k^{n-1}$ and $\hat{\theta}_k^n$ is observed.

We believe that the coefficient, $\hat{\Omega}_k^n$, is reliable if the learnable parameter θ_k and the fixed parameter $\hat{\theta}_k^{n-1}$ are on the same side from $\hat{\theta}_k^n$ because θ_k can observe the true loss function during optimization process. Conversely, if θ_k and $\hat{\theta}_k^{n-1}$ are on the opposite sides with respect to $\hat{\theta}_k^n$, $\hat{\Omega}_k^n$ is unreliable. Figure 3 visualizes the relations between the parameters for the both cases.

To model the relationship between the variables, we introduce the following function:

$$\alpha(\theta_k) \equiv (\theta_k - \hat{\theta}_k^n) (\hat{\theta}_k^{n-1} - \hat{\theta}_k^n). \quad (4)$$

In our framework, since the true loss function is not fully observable, we introduce an additional parameter, $\hat{\Omega}_k^n (\geq 0)$, for the quadratic approximation, which results in the asymmetric loss function as

$$\mathcal{L}_s^n(\theta_k, a) = \begin{cases} \hat{\Omega}_k^n (\theta_k - \hat{\theta}_k^n)^2 & \text{if } \alpha(\theta_k) > 0 \\ (a\hat{\Omega}_k^n + \epsilon)(\theta_k - \hat{\theta}_k^n)^2 & \text{if } \alpha(\theta_k) \leq 0 \end{cases}, \quad (5)$$

where $a (> 1)$ is a hyperparameter to control the degree of overestimation and ϵ is a small positive number to make the loss overestimated even when $\hat{\Omega}_k^n = 0$ ¹. From now, we omit the hyperparameter a in $\mathcal{L}_s^n(\theta_k, a)$ for notational simplicity, i.e., $\mathcal{L}_s^n(\theta_k, a) \equiv \mathcal{L}_s^n(\theta_k)$.

The asymmetric loss function in Eq. (5) is used to define the total loss $\tilde{\mathcal{L}}^n$, which is given by

$$\begin{aligned} \tilde{\mathcal{L}}^n &= \mathcal{L}^n + c\mathcal{L}_s^{n-1} \\ &= \mathcal{L}^n + c \sum_k \mathcal{L}_s^{n-1}(\theta_k). \end{aligned} \quad (6)$$

¹Since $\hat{\Omega}_k^n$ is non-negative, $(a\hat{\Omega}_k^n + \epsilon) > \hat{\Omega}_k^n$.

Note that \mathcal{L}_s^{n-1} can also be interpreted as a regularizer for the n^{th} task. One remaining concern is how to compute $\hat{\Omega}_k^n$ in Eq. (5), which is discussed next.

4.3. Accurate Quadratic Approximation

In addition to the overestimation of the true loss in the unobserved sides for a new task n , our algorithm estimates the optimal coefficient $\hat{\Omega}_k^n$ for the quadratic approximation of the loss function in its observed sides. Note that, as presented in Eq. (5), $\hat{\Omega}_k^n$ affects the approximation of the loss function in both sides; it is critical to derive $\hat{\Omega}_k^n$ accurately for performance improvement.

The quadratic surrogate loss function in SI is determined by approximating its parameters $\hat{\Omega}_k^n$ as in Eq. (3), which is derived from Eq. (2). Instead of using the equation, we present a new derivation that leads to the exact quadratic approximation, which is given by

$$\begin{aligned} \hat{\Omega}_k^n &= \frac{\mathcal{L}_s^n(\theta_k)}{(\theta_k - \hat{\theta}_k^n)^2} = \frac{\mathcal{L}_s^n(\hat{\theta}_k^{n-1})}{(\hat{\theta}_k^{n-1} - \hat{\theta}_k^n)^2} \\ &= \frac{\mathcal{L}_s^n(\hat{\theta}_k^{n-1}) - \mathcal{L}_s^n(\hat{\theta}_k^n)}{(\hat{\theta}_k^{n-1} - \hat{\theta}_k^n)^2} \\ &= \frac{\omega_k^n + \omega_k^{1:(n-1)}}{(\hat{\theta}_k^n - \hat{\theta}_k^{n-1})^2}. \end{aligned} \quad (7)$$

Note that $\hat{\Omega}_k^n$ means the change of the surrogate loss \mathcal{L}_s^n when the model parameter changes from $\hat{\theta}_k^{n-1}$ to $\hat{\theta}_k^n$ during the learning process of the n^{th} task. The following properties and definitions are required to derive Eq. (7):

$$\mathcal{L}_s^n(\hat{\theta}_k^n) \propto (\hat{\theta}_k^n - \hat{\theta}_k^n)^2 = 0 \quad (8)$$

$$\mathcal{L}_s^n(\hat{\theta}_k^n) = \mathcal{L}^n(\hat{\theta}_k^n) + c\mathcal{L}_s^{n-1}(\hat{\theta}_k^n) \quad (9)$$

and

$$\omega_k^n \equiv \mathcal{L}^n(\hat{\theta}_k^{n-1}) - \mathcal{L}^n(\hat{\theta}_k^n) \quad (10)$$

$$\begin{aligned} \omega_k^{1:(n-1)} &\equiv c\left(\mathcal{L}_s^{n-1}(\hat{\theta}_k^{n-1}) - \mathcal{L}_s^{n-1}(\hat{\theta}_k^n)\right) \\ &= -c\mathcal{L}_s^{n-1}(\hat{\theta}_k^n). \end{aligned} \quad (11)$$

The second equality of Eq. (7) is simply obtained from plugging $\hat{\theta}_k^{n-1}$ into θ_k in both the numerator and denominator. The numerator in Eq. (7) has two terms corresponding to Eq. (10) and (11), which implies that $\hat{\Omega}_k^n$ is given by the sum of the two parameter importances—the loss differences with respect to the parameter changes—of the new and previous tasks. Eq. (7) is actually different from Eq. (3) only

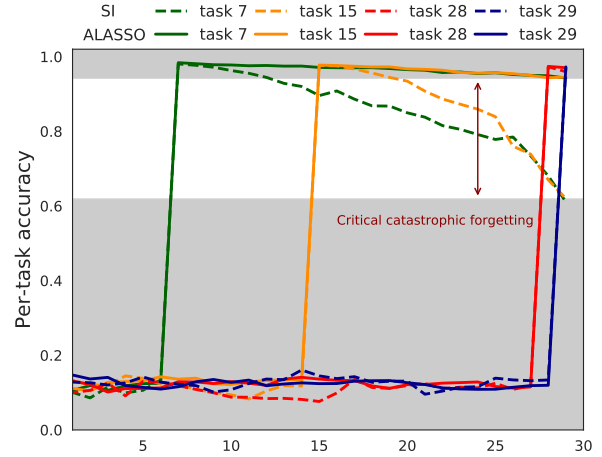


Figure 4. Analysis of results from ALASSO in comparison to SI, which is one of the state-of-the-art methods. This figure presents per-task test accuracy of the selected tasks on the permuted MNIST dataset over time. In this graph, x -axis denotes the index of most recently trained task, and y -axis is accuracy. All the visualized tasks achieve very high accuracy by the two algorithms when they are trained initially, but their accuracies are degraded as newer tasks come in. This forgetting problem is more catastrophic in SI than ALASSO by a large margin; SI is not effective to maintain accuracy of tasks learned in the past as the newer tasks are considered for training.

in the importance term of the previous tasks. In principle, the parameter importance of the previous tasks should be updated depending on the identified local minimum ($\hat{\theta}_k^n$) of the quadratic function. Note that our formulation presented in Eq. (7) realizes this exactly while Eq. (3) uses the fixed importance, $\hat{\Omega}_k^{n-1}$, resulting in poor approximation.

4.4. Parameter Decoupling

The overestimated loss approximation is effective to reduce errors that happen inevitably in the previous tasks, but the optimization with this technique may suffer from slow convergence due to its inherent limitation. Specifically, some hyperparameters, such as a and c , affect the objective function in one way or another depending on where they occur in our formulation. This is because we approximate the real loss for all the previous tasks using a single asymmetric function and it is almost impossible to consider all the possible combinations of the approximate functions estimated in the past.

In our formulation, there are two different occurrences of hyperparameters; one is related to the gradient computation in SGD and the other is in calculating $\hat{\Omega}_k^n$ for loss approximation. For example, increasing c in Eq. (6) makes the model consider the previous tasks more while increasing c in Eq. (11) results in more weight on the current task by underestimating $\hat{\Omega}_k^n$.

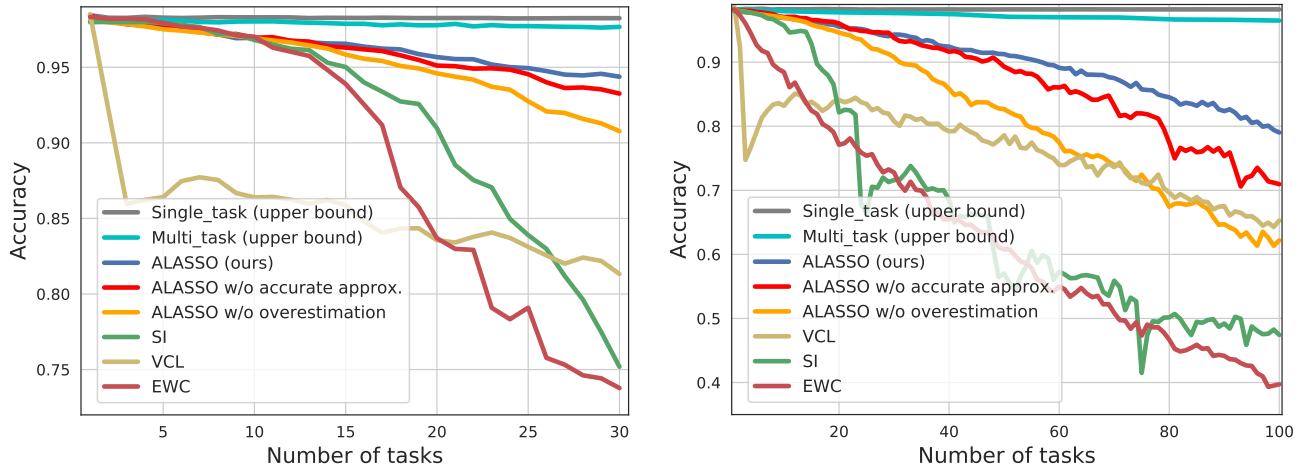


Figure 5. Test accuracy on the permuted MNIST dataset with 30 (left) and 100 (right) tasks. The results from several continual learning techniques including ours (ALASSO) are presented together with two upper-bound methods, *single_task* and *multi_task*. Note that x -axis denotes the index of a task given by the order of training. ALASSO achieves the state-of-the-art performance for both cases. Each method uses the hyper-parameters optimized for itself. The accuracies of SGD, SGD+dropout, IMM and MAS are much worse and we do not include them in this graph for better visualization.

To handle the inconsistent impact of identical hyper-parameters on the optimization process, we decouple the parameters into two sets; a set of parameters used for computation of $\omega_k^{1:(n-1)}$ and the other set of parameters used for surrogate loss estimation.

Then, when we compute $\omega_k^{1:(n-1)}$ in Eq. (11), the hyper-parameters a and c are decoupled from the other equations as

$$\omega_k^{1:(n-1)} = -c' \mathcal{L}_s^{n-1} (\hat{\theta}_k^n, a'), \quad (12)$$

where a and c in Eq. (11) are replaced by a' and c' , respectively.

4.5. Discussion

Figure 4 illustrates the promising results of the proposed algorithm in comparison to SI, which is one of the state-of-the-art methods. The figure presents how the accuracy of each task that is learned earlier changes over time as new tasks are added one by one. We notice that the amount of degradation in ALASSO (ours) is much smaller than SI and the accuracy differences of the two algorithms at the same number of tasks are getting larger as time goes by. These results clearly show the potential of our algorithm.

We claim that the proposed algorithm is practically good because it does not involve side effects such as architectural modification, network size increase, additional memory requirements, batch processing, multiple execution of networks, and inference on multi-head networks. We only need to store several variables such as Ω_k^n , $\hat{\theta}_k^n$, and ω_k^n during training, and perform additional operations to compute the surrogate losses.

5. Experiments

This section demonstrates performance of our algorithm compared to existing approaches on the standard datasets.

5.1. Datasets and Algorithms

We employ three standard benchmark datasets to evaluate the proposed continual learning framework, which include the permuted MNIST, the split CIFAR-10/CIFAR-100 and the split Tiny ImageNet. The permuted MNIST is a synthetic dataset based on MNIST [19], where all pixels of an image in MNIST are permuted differently but coherently in each task. This dataset contains a large number of tasks and is widely used for evaluation of continual learning algorithms [36, 10, 17, 37]. The split CIFAR-10/CIFAR-100 dataset is generated from CIAFR-10 and CIFAR-100 [18] while the split Tiny ImageNet is derived from Tiny ImageNet [1]. These datasets divide their target classes into multiple subsets, which correspond to individual tasks.

We compare our algorithm with the existing state-of-the-art methods including naïve SGD [29, 16], SGD with dropout (SGD+dropout) [10], VCL [27], EWC [17], SI [37], IMM [21] and MAS [2]. To estimate the upper-bound accuracy of our approach, we present the results from the models learned for individual tasks and all the tasks in a batch, which are denoted by *single_task* and *multi_task*, respectively.

5.2. Training Details

Our models are trained based on the description in SI [37]; the models are optimized by Adam and the learning rate is 0.001 for all the tested datasets. The batch size is 256

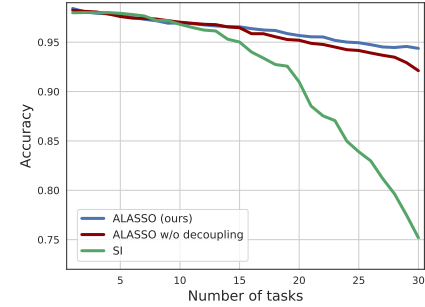
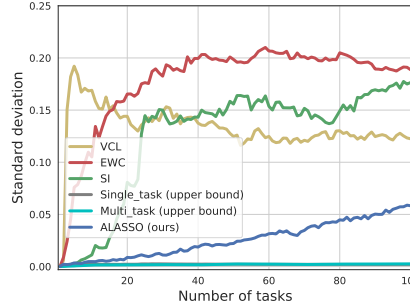
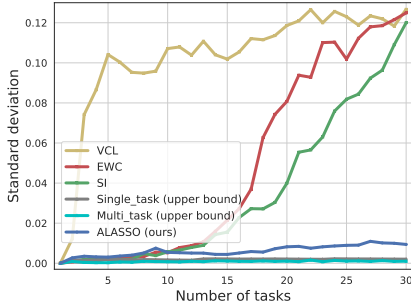


Figure 6. The standard deviation of per-task test accuracy on the permuted MNIST dataset when 30 (left) and 100 (right) tasks are learned online. Figure 7. Result with and without parameter decoupling on the permuted MNIST.

for the permuted MNIST and the split CIFAR-10/CIFAR-100, and 128 for the split Tiny ImageNet while the numbers of epochs are 20 for the permuted MNIST and 60 for the other two datasets. We plan to release our source code and raw results for better reproducibility.

5.3. Results on the Permuted MNIST

We present the results from all compared algorithms on 30 and 100 tasks of the permuted MNIST dataset². The network architectures used in this experiment are simple multi-layer perceptrons composed of two hidden layers with additional ReLU and softmax output layers. Dropout is not used except for SGD+dropout.

As Figure 5 shows, ALASSO outperforms all compared methods with large margins (about 15% point at least) for both 30 and 100 tasks. Note that all other methods undergo large performance drops as the number of tasks increases while ALASSO is more robust than others; it presents only moderate performance loss even after learning 100 tasks.

To verify the effectiveness of the asymmetric loss overestimation on the unobserved side and the accurate quadratic approximation on the observable one, we perform the ablation study to analyze their benefit. Figure 5 also illustrates the clear contribution of the components.

Figure 6 shows that the accuracy of ALASSO is stable compared to the other methods. The overall performance of ALASSO is not affected by parameter decoupling but rather improves accuracy slightly as illustrated in Figure 7. Note that it facilitates the fast convergence of the models.

5.4. Results from the Split CIFAR-10/CIFAR-100

We evaluate our algorithm on a more realistic scenario using the split CIFAR-10/CIFAR-100 dataset. In this experiment, each task is composed of 3 classes and the number of classes in a task gets larger as the index of the task increases. Our model is based on a convolutional

²For better visualization, we omit results from SGD, SGD+dropout, IMM and MAS because their accuracies are significantly lower than the presented ones.

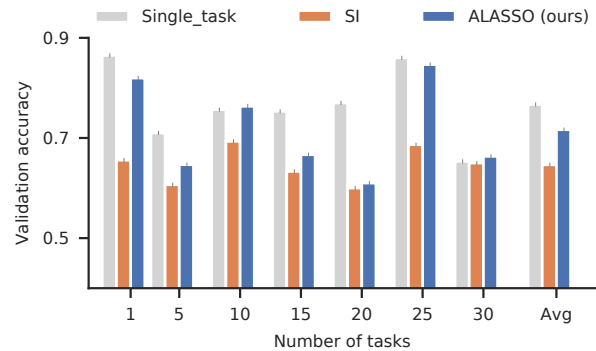


Figure 8. Per-task accuracy at every 5 task on the split CIFAR-10/CIFAR-100 when training 30 tasks in an online manner. ALASSO works better than SI and comparable to *single_task* in this more realistic dataset.

neural network with 4 convolutional layers followed by 2 fully connected layers with dropouts; ReLU and 2×2 max pooling are also employed to add nonlinearity in the model. Figure 8 demonstrates the results from *single_task* (to show the upper-bound accuracy), SI, and ALASSO on the split CIFAR-10/CIFAR-100 dataset. The proposed method outperforms SI consistently even on this more realistic dataset and achieves 5.7% point better than SI in average; the accuracy of SI drops by 13.9% point and 41% of the lost accuracy in SI is recovered by ALASSO. Note that ALASSO even shows the comparable performance to *single_task* method in some cases.

5.5. Results from the Split Tiny ImageNet

We conduct an experiment with a larger network on the split Tiny ImageNet dataset, which has more classes and consists of the tasks with 6 classes. The model has 5 convolutional layers followed by 2 fully connected layers with dropouts while 2×2 max pooling. ReLU activation function are employed as well to add nonlinearity of the network.

Figure 9 illustrates the results from *single_task*, SI, and ALASSO. Note that the overall accuracy of ALASSO is as competitive to the *single_task*, which is supposed to show a

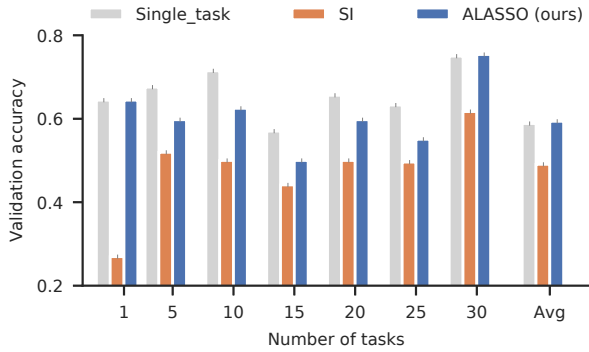


Figure 9. Per-task accuracy of every 5 task on the split Tiny ImageNet when training 30 tasks in an online manner. Even in this more realistic dataset with more classes, the accuracy ALASSO is better than SI and even comparable to *single_task*.

Table 1. Sensitivity to a on the permuted MNIST

a	0.8	1.0	2.0	3.0	4.0	5.0
30 tasks	0.91	0.92	0.94	0.94	0.94	0.94
100 tasks	0.59	0.62	0.79	0.79	0.78	0.75

practical upper-bound performance conceptually. In reality, the average validation accuracy of our method is 59.4% and even higher than *single_task* accuracy, 58.9%. ALASSO also outperforms SI by approximately 10% point in average.

5.6. More Analysis

The primary hyperparameter in our algorithm is a , which is introduced to overestimate the unobserved side of the loss function. Since a determines the factor of overestimation, it is reasonable to set its value larger than 1.0. We determine the value of a empirically in the permuted MNIST dataset; we choose the similar values for the other datasets, and they are fixed within each dataset. According to our experience, the overall performance is not sensitive to a wide range of a 's value as presented in Table 1. if a is 1, loss function is approximated by symmetric function. Note that the performance of our algorithm degrades when $a = 0.8$ (< 1). Another hyperparameter c balances between the losses in the current and the past task, and it is set to 1.0 throughout the experiments.

Although our original formulation uses no additional memory, the performance of ALASSO is compared with VCL [27], which is a method with episodic memory, in the same condition. Figure 10 illustrates that our algorithm still achieve outstanding performance.

Table 2 present accuracy (A), forgetting (F) and intransigence (I) measures of EWC, SI, and ALASSO on the permuted MNIST and the split CIFAR-10/CIFAR-100, where the subscripts besides A , F , and I denote the number of tasks when the performance is computed. Note that the accuracy measure is most comprehensive and is conceptu-

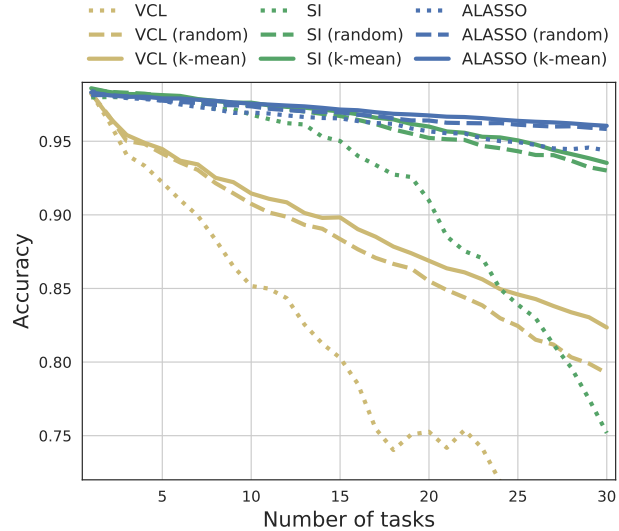


Figure 10. Results with episodic memory for 200 examples per task on the permuted MNIST.

Table 2. Algorithm comparison in terms of accuracy, forgetting and intransigence measures [6]

Measure	Permuted MNIST			Split CIFAR10/100		
	A_{30}	F_{30}	I_{30}	A_{30}	F_{30}	I_{30}
EWC	0.738	0.243	0.010	-	-	-
SI	0.752	0.232	0.008	0.632	0.110	0.137
Ours	0.944	0.027	0.015	0.697	0.026	0.123

ally correlated to the other two ones. ALASSO outperforms EWC and SI, especially in terms of accuracy and forgetting measures.

6. Conclusion

We presented a novel continual learning framework based on the overestimated asymmetric loss approximation with the better parametrization for the quadratic approximation, which is a carefully designed generalized version of SI. Our algorithm alleviates the catastrophic forgetting issue, which is common in deep neural networks, and is particularly helpful for the scenario with a large number of tasks in continual learning. The proposed solution is motivated by the observation that network parameter updates do not affect target loss function symmetrically, and does not incur substantial side effects. It achieves the state-of-the-art performance on several challenging standard benchmark datasets.

Acknowledgments This work was supported by the National Research Foundation (NRF) grant funded by the Korea Government (MSIT) (NRF2017R1A2B2011862).

References

- [1] Tiny ImageNet Visual Recognition Challenge. Available at

- tiny-imagenet.herokuapp.com.
- [2] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision*, 2018.
 - [3] Rahaf Aljundi, Punarjay Chakravarty, and Tinne Tuytelaars. Expert gate: Lifelong learning with a network of experts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
 - [4] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International Conference on Machine Learning*, 2015.
 - [5] Tamara Broderick, Nicholas Boyd, Andre Wibisono, Ashia C Wilson, and Michael I Jordan. Streaming variational bayes. In *Advances in Neural Information Processing Systems*, 2013.
 - [6] Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European Conference on Computer Vision*, 2018.
 - [7] Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A Rusu, Alexander Pritzel, and Daan Wierstra. Pathnet: Evolution channels gradient descent in super neural networks. *arXiv preprint arXiv:1701.08734*, 2017.
 - [8] Robert M French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 1999.
 - [9] Zoubin Ghahramani and H Attias. Online variational bayesian learning. In *NIPS workshop on Online Learning*, 2000.
 - [10] Ian Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*, 2013.
 - [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2014.
 - [12] Xu He and Herbert Jaeger. Overcoming catastrophic interference using conceptor-aided backpropagation. In *International Conference on Learning Representations*, 2018.
 - [13] Wenpeng Hu, Zhou Lin, Bing Liu, Chongyang Tao, Zhengwei Tao, Jinwen Ma, Dongyan Zhao, and Rui Yan. Overcoming catastrophic forgetting via model adaptation. In *International Conference on Learning Representations*, 2019.
 - [14] Heechul Jung, Jeongwoo Ju, Minju Jung, and Junmo Kim. Less-forgetful learning for domain expansion in deep neural networks. In *AAAI Conference on Artificial Intelligence*, 2018.
 - [15] Ronald Kemker and Christopher Kanan. Fearnnet: Brain-inspired model for incremental learning. In *International Conference on Learning Representations*, 2018.
 - [16] Jack Kiefer, Jacob Wolfowitz, et al. Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, 1952.
 - [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*, 2017.
 - [18] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. The CIFAR-10 and CIFAR-100 datasets. Available at www.cs.toronto.edu/~kriz/cifar.html.
 - [19] Yann LeCun, Cortes Cortes, and Christopher J.C. Burges. The MNIST database of handwritten digits. Available at yann.lecun.com/exdb/mnist/.
 - [20] Jeongtae Lee, Jaehong Yun, Sungju Hwang, and Eunho Yang. Lifelong learning with dynamically expandable networks. In *International Conference on Learning Representations*, 2018.
 - [21] Sang-Woo Lee, Jin-Hwa Kim, Jaehyun Jun, Jung-Woo Ha, and Byoung-Tak Zhang. Overcoming catastrophic forgetting by incremental moment matching. In *Advances in Neural Information Processing Systems*, 2017.
 - [22] Sang-Woo Lee, Chung-Yeon Lee, Dong-Hyun Kwak, Jiwon Kim, Jeonghee Kim, and Byoung-Tak Zhang. Dual-memory deep learning architectures for lifelong learning of everyday human behaviors. In *International Joint Conference on Artificial Intelligence*, 2016.
 - [23] Zhizhong Li and Derek Hoiem. Learning without forgetting. In *Proceedings of the European Conference on Computer Vision*. Springer, 2016.
 - [24] David Lopez-Paz et al. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, 2017.
 - [25] Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
 - [26] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*. Elsevier, 1989.
 - [27] Cuong V Nguyen, Yingzhen Li, Thang D Bui, and Richard E Turner. Variational continual learning. In *International Conference on Learning Representations*, 2018.
 - [28] 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*, 2017.
 - [29] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 1951.
 - [30] 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.
 - [31] Masa-Aki Sato. Online model selection based on the variational bayes. *Neural computation*, 2001.
 - [32] Jonathan Schwarz, Jelena Luketina, Wojciech Czarnecki, Agnieszka Grabska-Barwinska, Yee Whye Teh, Razvan Pascanu, and Raia Hadsell. Progress compress: A scalable

- framework for continual learning. In *International Conference on Machine Learning*, 2018.
- [33] Joan Serra, Dídac Surís, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning*, 2018.
- [34] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems*, 2017.
- [35] Pablo Sprechmann, Siddhant M Jayakumar, Jack W Rae, Alexander Pritzel, Adrià Puigdomènech Badia, Benigno Uria, Oriol Vinyals, Demis Hassabis, Razvan Pascanu, and Charles Blundell. Memory-based parameter adaptation. In *International Conference on Learning Representations*, 2018.
- [36] Rupesh K Srivastava, Jonathan Masci, Sohrab Kazerounian, Faustino Gomez, and Jürgen Schmidhuber. Compete to compute. In *Advances in Neural Information Processing Systems*, 2013.
- [37] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, 2017.