

Enabling Continual Learning with Differentiable Hebbian Plasticity

Vithursan Thangarasa

School of Engineering, University of Guelph
Vector Institute for Artificial Intelligence
Guelph, Canada
vthangar@uoguelph.ca

Thomas Miconi

Uber AI
San Francisco, United States
tmiconi@uber.com

Graham W. Taylor

School of Engineering, University of Guelph
Vector Institute for Artificial Intelligence
Guelph, Canada
gwtaylor@uoguelph.ca

Abstract—Continual learning is the problem of sequentially learning new tasks or knowledge while protecting previously acquired knowledge. However, catastrophic forgetting poses a grand challenge for neural networks performing such learning process. Thus, neural networks that are deployed in the real world often struggle in scenarios where the data distribution is non-stationary (concept drift), imbalanced, or not always fully available, i.e., rare edge cases. We propose a Differentiable Hebbian Consolidation model which is composed of a Differentiable Hebbian Plasticity (DHP) Softmax layer that adds a rapid learning plastic component (compressed episodic memory) to the fixed (slow changing) parameters of the softmax output layer; enabling learned representations to be retained for a longer timescale. We demonstrate the flexibility of our method by integrating well-known task-specific synaptic consolidation methods to penalize changes in the slow weights that are important for each target task. We evaluate our approach on the Permuted MNIST, Split MNIST and Vision Datasets Mixture benchmarks, and introduce an imbalanced variant of Permuted MNIST — a dataset that combines the challenges of class imbalance and concept drift. Our proposed model requires no additional hyperparameters and outperforms comparable baselines by reducing forgetting.

Index Terms—neural networks, plasticity, catastrophic forgetting, continual learning, hebbian learning

I. INTRODUCTION

A key aspect of human intelligence is the *ability to continually adapt and learn* in dynamic environments, a characteristic which is challenging to embed into artificial intelligence. Recent advances in machine learning (ML) have shown tremendous improvements in various problems, by learning to solve one complex task very well, through extensive training on large datasets with millions of training examples or more. However, most of the ML models that are used during deployment in the real-world are exposed to non-stationarity where the distributions of acquired data changes over time. Therefore, after learning is complete, and these models are further trained with new data, responding to distributional changes, performance degrades with respect to the original data. This phenomenon known as *catastrophic forgetting* or *catastrophic interference* [1], [2] presents a crucial problem for deep neural networks (DNNs) that are tasked with continual learning [3], also called lifelong learning [4], [5]. In continual learning, the goal is to adapt and learn consecutive tasks without forgetting

how to perform well on previously learned tasks, enabling models that are scalable and efficient over long timescales.

In most supervised learning methods, DNN architectures require independent and identically distributed (iid) samples from a stationary training distribution. However, for ML systems in real-world applications that require continual learning, the iid assumption is easily violated when: (1) There is concept drift in the training data distribution. (2) There are imbalanced class distributions and concept drift occurring simultaneously. (3) Data representing all scenarios in which the learner is expected to perform are not initially available. In such situations, learning systems face the “stability-plasticity dilemma” which is a well-known problem for artificial and biological neural networks [6], [7]. This presents a continual learning challenge for an ML system where the model needs to provide a balance between its plasticity (to integrate new knowledge) and stability (to preserve existing knowledge).

In biological neural networks, synaptic plasticity has been argued to play an important role in learning and memory [8]–[10] and two major theories have been proposed to explain a human’s ability to perform continual learning. The first theory is inspired by synaptic consolidation in the mammalian neocortex [11] where a subset of synapses are rendered less plastic and therefore preserved for a longer timescale. The general idea for this approach is to consolidate and preserve synaptic parameters that are considered important for the previously learned tasks. This is normally achieved through task-specific updates of synaptic weights in a neural network. The second is the complementary learning system (CLS) theory [12], [13], which suggests that humans extract high-level structural information and store it in different brain areas while retaining episodic memories.

Recent work on differentiable plasticity has shown that neural networks with “fast weights” that leverage Hebbian learning rules [14] can be trained end-to-end through back-propagation and stochastic gradient descent (SGD) to optimize the standard “slow weights”, as well as also the amount of plasticity in each synaptic connection [15], [16]. These works use slow weights to refer to the weights normally used to train vanilla neural networks, which are updated slowly and are often associated with long-term memory. The fast weights represent the weights that are superimposed on the slow weights and

Work done while Vithursan Thangarasa was an intern at Uber AI.

change quickly from one time step to the next based on input representations. These fast weights behave as a form of short-term memory that enable “reactivation” of long-term memory traces in the slow weights. [16] showed that simple plastic networks with learned plasticity outperform networks with uniform plasticity on various problems. Moreover, there have been several approaches proposed recently for overcoming the catastrophic forgetting problem in fixed-capacity models by dynamically adjusting the plasticity of each synapse based on its importance for retaining past memories [17].

Here, we extend the work on differentiable plasticity to the *task-incremental* continual learning setting [18], where tasks arrive in a batch-like fashion, and have clear boundaries. We develop a Differentiable Hebbian Consolidation model that is capable of adapting quickly to changing environments as well as consolidating previous knowledge by selectively adjusting the plasticity of synapses. We modify the traditional softmax layer and propose to augment the slow weights in the final fully-connected (FC) layer (softmax output layer) with a set of plastic weights implemented using Differentiable Hebbian Plasticity (DHP). Furthermore, we demonstrate the flexibility of our model by combining it with recent task-specific synaptic consolidation based approaches to overcoming catastrophic forgetting such as elastic weight consolidation [19], [20], synaptic intelligence [21] and memory aware synapses [22]. Our model unifies core concepts from Hebbian plasticity, synaptic consolidation and CLS theory to enable rapid adaptation to new unseen data, while consolidating synapses and leveraging compressed episodic memories in the softmax layer to remember previous knowledge and mitigate catastrophic forgetting. We test our proposed method on established benchmark problems including the Permuted MNIST [23], Split MNIST [21] and Vision Datasets Mixture [24]. We also introduce the Imbalanced Permuted MNIST problem and show that plastic networks with task-specific synaptic consolidation methods outperform networks with uniform plasticity.

II. RELEVANT WORK

Neural Networks with Non-Uniform Plasticity: One of the major theories that have been proposed to explain a human’s ability to learn continually is Hebbian learning [14], which suggests that learning and memory are attributed to weight plasticity, that is, the modification of the strength of existing synapses according to variants of Hebb’s rule [25]–[27]. It is a form of activity-dependent synaptic plasticity where correlated activation of pre- and post-synaptic neurons leads to the strengthening of the connection between the two neurons. According to the Hebbian learning theory, after learning, the related synaptic strength are enhanced while the degree of plasticity decreases to protect the learned knowledge [28].

Recent approaches in the meta-learning literature have shown that we can incorporate fast weights into a neural network to perform one-shot and few-shot learning [29], [30]. [29] proposed a model that augments FC layers preceding the softmax with a matrix of fast weights to bind labels to representations. Here, the fast weights were implemented with

non-trainable Hebbian learning-based associative memory. The Hebbian Softmax layer [30] can improve learning of rare classes by interpolating between Hebbian learning and SGD updates on the output layer using a scheduling scheme.

Differentiable plasticity [16] uses SGD to optimize the plasticity of each synaptic connection, in addition to the standard fixed (slow) weights. Here, each synapse is composed of a slow weight and a plastic (fast) weight that automatically increases or decreases based on the activity over time. Although this approach served to be a powerful new method for training neural networks, it was mainly demonstrated on recurrent neural networks (RNNs) for solving pattern memorization tasks and maze exploration with reinforcement learning. Also, these approaches were only demonstrated on meta-learning problems and not the continual learning challenge of overcoming catastrophic forgetting. Our work also augments the slow weights in the FC layer with a set of plastic (fast) weights, but implements these using DHP. We only update the parameters of the softmax output layer in order to achieve fast learning and preserve knowledge over time.

Overcoming Catastrophic Forgetting: This work leverages two strategies to overcome the catastrophic forgetting problem: 1) *Task-specific Synaptic Consolidation* — Protecting previously learned knowledge by dynamically adjusting the synaptic strengths to consolidate and retain memories. 2) *CLS Theory* — A dual memory system where, the neocortex (neural network) gradually learns to extract structured representations from the data while, the hippocampus (augmented episodic memory) performs rapid learning and individuated storage to memorize new instances or experiences.

There have been several notable works inspired by task-specific synaptic consolidation for mitigating catastrophic forgetting [19], [21], [22] and they are often categorized as regularization strategies in the continual learning literature [17]. All of these regularization approaches estimate the importance of each parameter or synapse, Ω_k , where least plastic synapses can retain memories for long timescales and more plastic synapses are considered less important. The parameter importance and network parameters θ_k are updated in either an online manner or after learning task T_n . Therefore, when learning new task T_{n+1} , a regularizer is added to the original loss function $\mathcal{L}^n(\theta)$, so that we dynamically adjust the plasticity w.r.t Ω_k and prevent any changes to important parameters of previously learned tasks:

$$\tilde{\mathcal{L}}^n(\theta) = \mathcal{L}^n(\theta) + \lambda \underbrace{\sum_k \Omega_k (\theta_k^n - \theta_k^{n-1})^2}_{\text{regularizer}} \quad (1)$$

where θ_k^{n-1} are the learned network parameters after training on the previous $n - 1$ tasks and λ is a hyperparameter for the regularizer to control the amount of forgetting.

The main difference in these regularization strategies is on the method used to compute the importance of each parameter, Ω_k . Elastic Weight Consolidation (EWC) [19] used the values given by the diagonal of an approximated Fisher information

matrix for Ω_k , and this was computed offline after training on a task was completed. An online variant of EWC was proposed by [20] to improve EWC’s scalability by ensuring the computational cost of the regularization term does not grow with the number of tasks. Synaptic Intelligence (SI) [21] is an online variant for computing the parameter importance where, Ω_k is the cumulative change in individual synapses over the entire training trajectory on a particular task. Memory Aware Synapses (MAS) [22] is an online method that measures Ω_k by the sensitivity of the learned function to a perturbation in the parameters, instead of measuring the change in parameters to the loss as seen in SI and EWC.

Our work draws inspiration from CLS theory which is a powerful computational framework for representing memories with a dual memory system via the neocortex and hippocampus. There have been numerous approaches based on CLS principles involving pseudo-rehearsal [31]–[33], exact or episodic replay [34], [35] and generative replay [36], [37]. However, in our work, we are primarily interested in neuro-plasticity techniques inspired from CLS theory for alleviating catastrophic forgetting. Earlier work from [38], [39] showed how each synaptic connection can be composed of a fixed weight where slow learning stores long-term knowledge and a fast-changing weight for temporary associative memory. This approach involving slow and fast weights is analogous to properties of CLS theory to overcome catastrophic forgetting during continual learning. Recent research in this vein has included replacing soft attention mechanism with fast weights in RNNs [40], the Hebbian Softmax layer [30], augmenting slow weights in the FC layer with a fast weights matrix [29], differentiable plasticity [15], [16] and neuromodulated differentiable plasticity [41]. However, all of these methods were focused on rapid learning on simple tasks or meta-learning over a distribution of tasks or datasets. Furthermore, they did not examine learning a large number of new tasks while, alleviating catastrophic forgetting in continual learning.

III. CONTINUAL LEARNING FORMULATION

In the continual learning setup, we train a neural network model on a sequence of tasks $T_{1:n_{\max}}$, where n_{\max} is the maximum number of tasks the model is to learn. Unlike the standard supervised learning setup, continual learning trains a model on data that is fetched in sequential chunks enumerated by tasks. Therefore, in a continual learning sequence, the model receives a sequence of tasks $T_{1:n_{\max}}$ that is to be learned, each with its associated training data $(\mathcal{X}_n, \mathcal{Y}_n)$, where \mathcal{X}_n is the input data and the corresponding label data denoted by \mathcal{Y}_n . Each task T_n has its own task-specific loss \mathcal{L}^n , that will be combined with a regularizer loss term (see Eq. 1) to prevent catastrophic forgetting. After training is complete, the model will have learned an approximated mapping f to the true underlying function \tilde{f} . The learned f maps a new input \mathcal{X} to the target outputs $\mathcal{Y}_{1:n}$ for all $T_{1:n}$ tasks the network has learned so far.

IV. DIFFERENTIABLE HEBBIAN CONSOLIDATION

In our model, each synaptic connection in the softmax layer has two weights: 1) The slow weights, $\theta \in \mathbb{R}^{m \times d}$, where m is the number of units in the final hidden layer and d is the number of outputs of the last layer. 2) A Hebbian plastic component of the same cardinality as the slow weights, composed of the plasticity coefficient, α , and the Hebbian trace, Hebb. The α is a scaling parameter for adjusting the magnitude of the Hebb. The Hebbian traces accumulate the mean hidden activations of the final hidden layer h for each target label in the mini-batch $\{y_{1:B}\}$ of size B which are denoted by $\tilde{h} \in \mathbb{R}^{1 \times m}$ (refer to Algorithm 1). Given the pre-synaptic activations of neurons i in h , we can formally compute the post-synaptic activations of neurons j using Eq. 2 and obtain the unnormalized log probabilities (softmax pre-activations) z . The softmax function is then applied on z to obtain the desired predicted probabilities \hat{y} thus, $\hat{y} = \text{softmax}(z)$. The η parameter in Eq. 3 is a scalar value that dynamically learns how quickly to acquire new experiences into the plastic component, and thus behaves as the “learning rate” for the plastic connections. The η parameter also acts as a decay term for the Hebb to prevent instability caused by a positive feedback loop in the Hebbian traces.

$$z_j = \sum_{i=1}^m \underbrace{(\theta_{i,j})}_{\text{slow}} + \underbrace{(\alpha_{i,j} \text{Hebb}_{i,j})}_{\text{plastic (fast)}} h_i \quad (2)$$

$$\text{Hebb}_{i,j} \leftarrow (1 - \eta) \text{Hebb}_{i,j} + \eta \tilde{h}_{i,j} \quad (3)$$

The network parameters $\alpha_{i,j}$, η and $\theta_{i,j}$ are optimized by gradient descent as the model is trained sequentially on different tasks in the continual learning setup. In standard neural networks the weight connection has only fixed (slow) weights, which is equivalent to setting $\alpha = 0$ in Eq. 2.

Algorithm 1 Batch update Hebbian traces.

- 1: **Input:** $h_{1:B}$ (hidden activations of penultimate layer),
 $y_{1:B}$ (target labels),
Hebb (Hebbian trace)
 - 2: **Output:** $z_{1:B}$ (softmax pre-activations)
 - 3: **for** each target label $c \in \{y_{1:B}\}$ **do**
 - 4: $s \leftarrow \sum_{b=1}^B [y_b = c]$ /*Count total occurrences of $c \in y$.*/
 - 5: **if** $s > 0$ **then**
 - 6: $\tilde{h} \leftarrow \frac{1}{s} \sum_{b=1}^B h[y_b = c]$ /*Update Hebb for class c .*/
 - 7: $\text{Hebb}_{:,c} \leftarrow (1 - \eta) \text{Hebb}_{:,c} + \eta \tilde{h}$
 - 8: **end if**
 - 9: **end for**
 - 10: $z \leftarrow (\theta + \alpha \text{Hebb})h$ /*Compute softmax pre-activations.*/
-

A. Hebbian Update Rule

The Hebbian traces are initialized to zero only at the start of learning the first task T_1 and during training, the Hebb

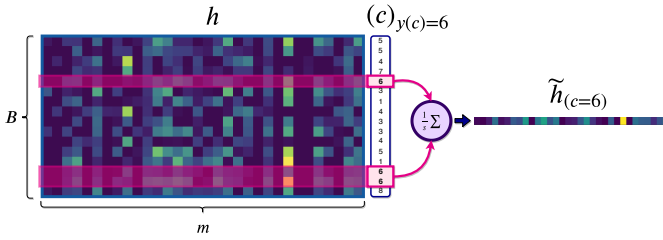


Fig. 1: An example of a Hebbian update for the class, $c = 6 \in y_{1:B}$. Here, we are given the hidden activations of the final hidden layer, h . Multiple hidden activations corresponding to class $c = 6$ (represented by the pink boxes) are averaged into one vector denoted by $\tilde{h} \in \mathbb{R}^{1 \times m}$. This Hebbian update visualization reflects Lines 4-6 in Algorithm 1 and is repeated for each unique class in the target vector $y_{1:B}$.

is automatically updated in the forward pass using Algorithm 1. Specifically, the Hebbian update for a corresponding class c in $y_{1:B}$ is computed on line 6. This Hebbian update $\frac{1}{s} \sum_{b=1}^B h[y_b = c]$ is analogous to another formulaic description of the Hebbian learning update rule $w_{i,j} = \frac{1}{N} \sum_{k=1}^N a_i^k a_j^k$ [14], where $w_{i,j}$ is the change in weight at connection i, j and a_i^k, a_j^k denote the activation levels of neurons i and j , respectively, for the k^{th} input. Therefore, in our model, $w = \tilde{h}$ the Hebbian weight update, $a_i = h$ the hidden activations of the last hidden layer, $a_j = y$ the corresponding target class in $y_{1:B}$ and $N = s$ the number of inputs for the corresponding class in $y_{1:B}$ (see Algorithm 1 and Figure 1 for an example). Across the model’s lifetime, we only update the Hebbian traces during training as it learns tasks in a continual manner. Therefore, during test time, we maintain and use the most recent Hebb traces to make predictions.

Our model explores an optimization scheme where hidden activations are accumulated directly into the softmax output layer weights when a class has been seen by the network. This results in better initial representations and can also retain these learned deep representations for a much longer timescale. This is because memorized activations for one class are not competing for space with activations from other classes. Fast learning, enabled by a highly plastic weight component, improves test accuracy for a given task. Between tasks this plastic component decays to prevent interference, but selective consolidation into a stable component protects old memories, effectively enabling the model to *learn to remember* by modelling plasticity over a range of timescales to form a learned neural memory (see Section V-A ablation study). In comparison to an external memory, the advantage of DHP Softmax is that it is simple to implement, requiring no additional space or computation. This allows it to scale easily with increasing number of tasks.

The plastic component learns rapidly and performs sparse parameter updates to quickly store memory traces for each recent experience without interference from other similar recent experiences. Furthermore, the hidden activations corresponding to the same class, c , are accumulated into one vector

\tilde{h} , thus forming a compressed episodic memory in the Hebbian traces to reflect individual episodic memory traces (similar to the hippocampus in biological neural networks [42], [43]). As a result, this method improves learning of rare classes and speeds up binding of class labels to deep representations of the data without introducing any additional hyperparameters.

B. Hebbian Synaptic Consolidation

Following the existing regularization strategies such as EWC [19], Online EWC [20], SI [21] and MAS [22], we regularize the loss $\mathcal{L}(\theta)$ as in Eq. 1 and update the synaptic importance parameters of the network in an online manner. We rewrite Eq. 1 to obtain the updated quadratic loss for Hebbian Synaptic Consolidation in Eq. 4 and show that the network parameters $\theta_{i,j}$ are the weights of the connections between pre- and post-synaptic activity, as seen in Eq. 2.

$$\tilde{\mathcal{L}}^n(\theta, \alpha, \eta) = \mathcal{L}^n(\theta, \alpha, \eta) + \lambda \sum_{i,j} \Omega_{i,j} (\theta_{i,j}^n - \theta_{i,j}^{n-1})^2 \quad (4)$$

We adapt the existing task-specific consolidation approaches to our model and do not compute the synaptic importance parameters on the plastic component of the network, hence we only regularize the slow weights of the network. Furthermore, when training the first task $T_{n=1}$, the synaptic importance parameter, $\Omega_{i,j}$ in Eq. 4, was set to 0 for all of the task-specific consolidation methods that we tested on except for SI. This is because SI is the only method we evaluated that estimates $\Omega_{i,j}$ while training, whereas Online EWC and MAS compute $\Omega_{i,j}$ after learning a task. The plastic component of the softmax layer in our model can alleviate catastrophic forgetting of consolidated classes by allowing gradient descent to optimize how plastic the connections should be (i.e. less plastic to preserve old information or more plastic to quickly learn new information).

V. EXPERIMENTS

In our experiments, we compare our approach to vanilla neural networks with Online EWC, SI and MAS. Since our approach increases the capacity of the DNN due to the addition of plastic weights, we add an extra set of slow weights to the softmax output layer of the standard neural network to match the capacity. We do this to show that it is not the increased model capacity from the plastic weights that is helping mitigate the forgetting when performing sequential task learning, thus ensuring a fair evaluation. We tested our model on the Permuted MNIST, Split MNIST and Vision Datasets Mixture benchmarks, and also introduce the Imbalanced Permuted MNIST problem.

For all of the benchmarks, we evaluated the model based on the average classification accuracy on all previously learned tasks as a function of n , the number of tasks trained so far. To determine memory retention and flexibility of the model, we are particularly interested in the test performance on the first task and the most recent one. We also measure forgetting using the backward transfer metric, BWT = $\frac{1}{T-1} \sum_{i=1}^{T-1} R_{T,i} - R_{i,i}$ [34], which indicates how much

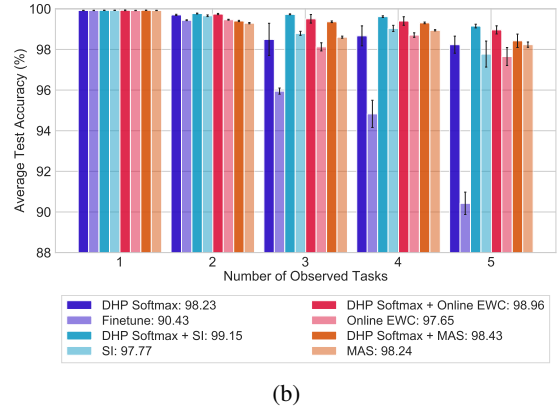
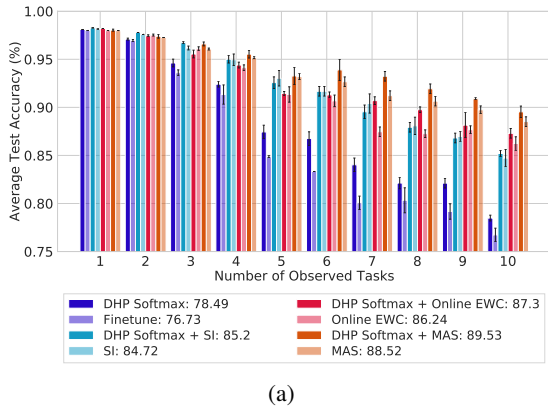


Fig. 2: (a) The average test accuracy on a sequence of $T_{n=1:10}$ Permutated MNIST tasks and (b) $T_{n=1:5}$ binary classification tasks from the MNIST dataset. The average test accuracy over all learned tasks is provided in the legend. The addition of DHP in all cases improves the model’s ability to reduce forgetting. Error bars represent the SEM across 10 trials.

learning new tasks has influenced the performance on previous tasks. $R_{T,i}$ is the test classification accuracy on task i after sequentially finishing learning the T^{th} task. While $\text{BWT} < 0$ directly reports catastrophic forgetting, $\text{BWT} > 0$ indicates that learning new tasks has helped with the preceding tasks. To establish a baseline for comparison of well-known task-specific consolidation methods, we trained neural networks with Online EWC, SI and MAS, respectively, on all tasks in a sequential manner. The hyperparameters of the consolidation methods (i.e. EWC, SI and MAS) remain the same with and without DHP Softmax, and the plastic components are not regularized. To find the best hyperparameter combination for each of these synaptic consolidation methods, we performed a grid search using a task sequence determined by a single seed. Descriptions of the hyperparameters and other details for all benchmarks can be found in Appendix A. All experiments were run on a Nvidia RTX 2080 Ti.

A. Permutated MNIST

In this benchmark, all of the MNIST pixels are permuted differently for each task with a fixed random permutation. Although the output domain is constant, the input distribution changes between tasks and is mostly independent of each other, thus, there exists a concept drift. In the Permutated MNIST and Imbalanced Permutated MNIST benchmarks we use a multi-layered perceptron (MLP) network with two hidden layers consisting of 400 LeakyReLU nonlinearities, and a cross-entropy loss. The η of the plastic component was set to be a value of 0.001 and we emphasize that we spent little to no effort on tuning the initial value of this parameter. We swept through a range of values $\eta \in \{0.1, 0.01, 0.001, 0.0005, 0.0001\}$ and found that setting η to low values led to the best performance in terms of being able to alleviate catastrophic forgetting (refer to Figure 4).

We first compare the performance between our network with DHP Softmax and a fine-tuned vanilla MLP network we refer to as *Finetune* in Figure 2a and no task-specific consolidation methods involved. The network with DHP Softmax alone

showed improvement in its ability to alleviate catastrophic forgetting across all tasks compared to the baseline network. Then we compared the performance with and without DHP Softmax using the same task-specific consolidation methods. Figure 2a shows the average test accuracy as new tasks are learned for the best hyperparameter combination for each task-specific consolidation method. We find our DHP Softmax with consolidation maintains a higher test accuracy throughout sequential training of tasks than without DHP Softmax.




Fig. 3: **(left)** Hebbian learning rate and decay value η , **(middle)** Frobenius Norm of the Hebbian memory traces $\|\text{Hebb}\|_F$, **(right)** Frobenius Norm of the plasticity coefficients $\|\alpha\|_F$ while training each task $T_{1:10}$.

Ablation Study: We further examine the structural parameters of the network and Hebb traces to provide further interpretability into the behaviour of our proposed model. The left plot in Figure 3 shows the behaviour of η during training as 10 tasks in the Permutated MNIST benchmark are learned continually. Initially, in task T_1 , η increases very quickly from 0.001 to 0.024 suggesting that the synaptic connections become more plastic to quickly acquire new information. Eventually, η decays after the 3rd task to reduce the degree of plasticity to prevent interference between the learned representations. We also observe that within each task from T_4 to T_{10} , η initially increases then decays. The Frobenius Norm of the Hebb trace (middle plot in Figure 3) suggests that Hebb grows *without* runaway positive feedback every time a new task is learned, maintaining a memory of which synapses contributed to recent activity. The Frobenius Norm of α (right plot in Figure 3) indicates that the plasticity coefficients grow




Fig. 4: A sensitivity analysis on the Hebb decay term η in Eq. 3. We show the average test accuracy for different initial values of η after learning all tasks on the (left) Permutated MNIST, (center) Imbalanced Permutated MNIST and (right) Split MNIST problems. The shaded regions correspond to the standard error of mean (SEM) across 5 trials.

within each task, thus the network continuously leverages the structure in the plastic component. It is important to note that SGD and backpropagation are used as *meta-learning* to tune the structural parameters in the plastic component.

B. Imbalanced Permutated MNIST

We introduce the Imbalanced Permutated MNIST problem which is identical to the Permutated MNIST benchmark but, now each task is an imbalanced distribution where training samples in each class were artificially removed based on some random probability (see Appendix A). This benchmark was motivated by the fact that class imbalance and concept drift can hinder predictive performance, and the problem becomes particularly challenging when they occur simultaneously. We see that DHP Softmax achieves 80.85% after learning 10 tasks with imbalanced class distributions in a sequential manner, thus providing significant 4.41% improvement over the standard neural network baseline of 76.44%. The significance of the compressed episodic memory mechanism in the Hebbian traces is more apparent in this benchmark because the plastic component allows rare classes that are encountered infrequently to be remembered for a longer period of time. We find that DHP Softmax with MAS achieves a 0.04 decrease in BWT, resulting in an average test accuracy of 88.80% and a 1.48% improvement over MAS alone.

C. Split MNIST

We split the original MNIST dataset [44] into a sequence of 5 binary classification tasks: $T_1 = \{0/1\}$, $T_2 = \{2/3\}$, $T_3 = \{4/5\}$, $T_4 = \{6/7\}$ and $T_5 = \{8/9\}$. The output spaces are disjoint between tasks, unlike the previous two benchmarks. Similar to the network used by [21], we use an MLP network with two hidden layers of 256 LeakyReLU nonlinearities each, and a cross-entropy loss. The initial η value was set to 0.001 as seen in previous benchmark experiments. We observed that DHP Softmax alone achieves 98.23% thus, provides a 7.80% improvement on test performance compared to a finetuned MLP network (Figure 2b). Also, combining DHP Softmax with task-specific consolidation consistently decreases BWT, leading to a higher average test accuracy across all tasks, especially the most recent one, T_5 .

D. Vision Datasets Mixture

Following previous works [24], [45], we perform continual learning on a sequence of 5 vision datasets: MNIST, notMNIST¹, FashionMNIST [46], SVHN [47] and CIFAR-10 [48]. The MNIST, notMNIST and FashionMNIST datasets are zero-padded to be of size 32×32 and are replicated 3 times to create grayscale images with 3 channels, thus matching the resolution of the SVHN and CIFAR-10 images.

Here, we use a CNN architecture that is similar to the one used in [24], [45], which consists of 2 convolutional layers with 20 and 50 channels respectively, and a kernel size of 5. Each convolution layer is followed by LeakyReLU nonlinearities (negative threshold of 0.3) and 2×2 max-pooling operations with stride 2. The two convolutional layers are followed by an FC layer of size 500 before the final softmax output layer. The initial η parameter value was set to 0.0001. We train the network with mini-batches of size 32 and optimized using plain SGD with a fixed learning rate of 0.01 for 50 epochs per task.

We found that DHP Softmax plus MAS decreases BWT by 0.04 resulting in a 2.14% improvement in average test accuracy over MAS on its own (see Table I). Also, SI with DHP Softmax outperforms other competitive methods with an average test performance of 81.75% and BWT of -0.04 after learning all five tasks. In Table I, we present a summary of the final average test performance after learning all tasks in the respective continual learning problems. Here, we summarize the average test accuracy and BWT across ten trials for each of the benchmarks.

VI. DISCUSSION AND CONCLUSION

We have shown that the problem of catastrophic forgetting in continual learning environments can be alleviated by adding compressed episodic memory in the softmax layer through DHP and performing task-specific updates on synaptic parameters based on their individual importance for solving previously learned tasks. The compressed episodic memory

¹Originally published at <http://yaroslavvb.blogspot.com/2011/09/notmnist-dataset.html> and downloaded from <https://github.com/davidflanagan/notMNIST-to-MNIST>.

TABLE I: The average test accuracy (% , higher is better) and backward transfer (BWT, lower is better) after learning all tasks on each benchmark, respectively. The results are averaged over 10 trials.

Method	Permuted-MNIST	Imbalanced Permuted-MNIST	SplitMNIST	5-Vision Mixture
Finetune	76.73 / -0.19	76.44 / -0.20	90.43 / -0.13	60.02 / -0.33
DHP Softmax	78.49 / -0.16	80.85 / -0.14	98.23 / -0.02	62.94 / -0.26
SI	84.72 / -0.13	85.92 / -0.06	97.77 / -0.04	81.26 / -0.06
DHP Softmax + SI	85.20 / -0.09	85.39 / -0.06	99.15 / 0.00	81.75 / -0.04
Online EWC	86.24 / -0.11	87.18 / -0.09	97.65 / -0.03	78.61 / -0.07
DHP Softmax + Online EWC	87.30 / -0.09	87.43 / -0.08	98.96 / -0.01	79.10 / -0.04
MAS	88.52 / -0.08	87.32 / -0.09	98.24 / -0.02	78.51 / -0.05
DHP Softmax + MAS	89.53 / -0.06	88.80 / -0.05	98.43 / -0.01	80.66 / -0.01

allows new information to be learned in individual traces without overlapping representations, thus avoiding interference when added to the structured knowledge in the slow changing weights and allowing the model to generalize across experiences. The α parameter in the plastic component automatically learns to scale the magnitude of the plastic connections in the Hebbian traces, effectively choosing when to be less plastic (protect old knowledge) or more plastic (acquire new information quickly). The neural network with DHP Softmax showed noticeable improvement across all benchmarks when compared to a neural network with a traditional softmax layer. The DHP Softmax does not introduce any additional hyperparameters since all of the structural parameters of the plastic part α and η are learned, and setting the initial η value required very little tuning effort.

We demonstrated the flexibility of our model where, in addition to DHP Softmax, we can perform Hebbian Synaptic Consolidation by regularizing the slow weights using EWC, SI or MAS to improve a model’s ability to alleviate catastrophic forgetting after sequentially learning a large number of tasks with limited model capacity. DHP Softmax combined with SI outperforms other consolidation methods on the Split MNIST and 5-Vision Datasets Mixture. The approach where we combine DHP Softmax and MAS consistently leads to overall superior results compared to other baseline methods on the Permuted MNIST and Imbalanced Permuted MNIST benchmarks. This is interesting because the local variant of MAS does compute the synaptic importance parameters of the slow weights $\theta_{i,j}$ layer by layer based on Hebb’s rule, and therefore synaptic connections i, j that are highly correlated would be considered more important for the given task than those connections that have less correlation. Furthermore, our model consistently exhibits lower negative BWT across all benchmarks, leading to higher average test accuracy over methods without DHP. This gives a strong indication that Hebbian plasticity enables neural networks to learn continually and remember distant memories, thus reducing catastrophic forgetting when learning from sequential datasets in dynamic environments. Furthermore, continual synaptic plasticity can play a key role in learning from limited labelled data while being able to adapt and scale at long timescales. We hope that our work will open new investigations into gradient descent

optimized Hebbian consolidation for learning and memory in DNNs to enable continual learning.

REFERENCES

- [1] M. McCloskey and N. J. Cohen, “Catastrophic interference in connectionist networks: The sequential learning problem,” *The Psychology of Learning and Motivation*, 1989.
- [2] R. French, “Catastrophic forgetting in connectionist networks,” *Trends in Cognitive Sciences*, Apr. 1999.
- [3] M. B. Ring, “Continual learning in reinforcement environments,” Ph.D. dissertation, Austin, TX, USA, 1994, uMI Order No. GAX95-06083.
- [4] S. Thrun and T. M. Mitchell, “Lifelong robot learning,” *Robotics and Autonomous Systems*, 1995, the Biology and Technology of Intelligent Autonomous Agents.
- [5] S. Thrun, *Lifelong Learning Algorithms*. Boston, MA: Springer US, 1998.
- [6] G. A. Carpenter and S. Grossberg, “A massively parallel architecture for a self-organizing neural pattern recognition machine,” *Computer Vision, Graphics, and Image Processing*, Jan. 1987.
- [7] W. C. Abraham and A. Robins, “Memory retention – the synaptic stability versus plasticity dilemma,” *Trends in Neurosciences*, Feb. 2005.
- [8] J. G. Howland and Y. T. Wang, “Chapter 8 synaptic plasticity in learning and memory: Stress effects in the hippocampus,” in *Essence of Memory*, ser. Progress in Brain Research. Elsevier, 2008.
- [9] T. Takeuchi, A. J. Duzkiewicz, and R. G. M. Morris, “The synaptic plasticity and memory hypothesis: encoding, storage and persistence,” *Philosophical Transactions of the Royal Society B: Biological Sciences*, Dec. 2013.
- [10] C. H. Bailey, E. R. Kandel, and K. M. Harris, “Structural components of synaptic plasticity and memory consolidation,” *Cold Spring Harbor Perspectives in Biology*, Jul. 2015.
- [11] M. K. Benna and S. Fusi, “Computational principles of synaptic memory consolidation,” *Nature Neuroscience*, Oct. 2016.
- [12] J. L. McClelland, B. L. McNaughton, and R. C. O’Reilly, “Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory,” *Psychological Review*, 1995.
- [13] D. Kumaran, D. Hassabis, and J. L. McClelland, “What learning systems do intelligent agents need? complementary learning systems theory updated,” *Trends in Cognitive Sciences*, 2016.
- [14] D. O. Hebb, *The organization of behavior; a neuropsychological theory*. Oxford, England: Wiley, 1949.
- [15] T. Miconi, “Learning to learn with backpropagation of hebbian plasticity,” *CoRR*, 2016.
- [16] T. Miconi, K. O. Stanley, and J. Clune, “Differentiable plasticity: training plastic neural networks with backpropagation,” in *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018.
- [17] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, and S. Wermter, “Continual lifelong learning with neural networks: A review,” *Neural Networks*, May 2019.
- [18] G. M. van de Ven and A. S. Tolias, “Three scenarios for continual learning,” *CoRR*, 2019.

- [19] J. Kirkpatrick, R. Pascanu, N. C. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, D. Hassabis, C. Clopath, D. Kumaran, and R. Hadsell, "Overcoming catastrophic forgetting in neural networks," *Proceedings of the National Academy of Sciences (PNAS)*, 2017.
- [20] J. Schwarz, W. Czarnecki, J. Luketina, A. Grabska-Barwinska, Y. W. Teh, R. Pascanu, and R. Hadsell, "Progress & compress: A scalable framework for continual learning," in *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018.
- [21] F. Zenke, B. Poole, and S. Ganguli, "Continual learning through synaptic intelligence," in *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017.
- [22] R. Aljundi, F. Babiloni, M. Elhoseiny, M. Rohrbach, and T. Tuytelaars, "Memory aware synapses: Learning what (not) to forget," in *The European Conference on Computer Vision (ECCV)*, September 2018.
- [23] I. J. Goodfellow, D. Warde-Farley, M. Mirza, A. Courville, and Y. Bengio, "Maxout networks," in *Proceedings of the 30th International Conference on International Conference on Machine Learning (ICML)*, 2013.
- [24] H. Ritter, A. Botev, and D. Barber, "Online structured laplace approximations for overcoming catastrophic forgetting," in *Advances in Neural Information Processing Systems*, 2018.
- [25] O. Paulsen and T. J. Sejnowski, "Natural patterns of activity and long-term synaptic plasticity," *Current Opinion in Neurobiology*, 2000.
- [26] S. Song, K. D. Miller, and L. F. Abbott, "Competitive hebbian learning through spike-timing-dependent synaptic plasticity," *Nature Neuroscience*, Sep. 2000.
- [27] E. Oja, "Oja learning rule," *Scholarpedia*, 2008.
- [28] F. Zenke, W. Gerstner, and S. Ganguli, "The temporal paradox of hebbian learning and homeostatic plasticity," *Current Opinion in Neurobiology*, 2017, neurobiology of Learning and Plasticity.
- [29] T. Munkhdalai and A. Trischler, "Metalearning with hebbian fast weights," *CoRR*, 2018.
- [30] J. W. Rae, C. Dyer, P. Dayan, and T. P. Lillicrap, "Fast parametric learning with activation memorization," in *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018.
- [31] A. Robins, "Catastrophic forgetting, rehearsal and pseudorehearsal," *Connection Science*, 1995.
- [32] B. Ans, S. Rousset, R. M. French, and S. Musca, "Self-refreshing memory in artificial neural networks: learning temporal sequences without catastrophic forgetting," *Connection Science*, Jun. 2004.
- [33] C. Atkinson, B. McCane, L. Szymanski, and A. V. Robins, "Pseudo-recursal: Solving the catastrophic forgetting problem in deep neural networks," *CoRR*, 2018.
- [34] D. Lopez-Paz and M. Ranzato, "Gradient episodic memory for continual learning," in *Advances in Neural Information Processing Systems*, 2017.
- [35] Z. Li and D. Hoiem, "Learning without forgetting," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Dec 2018.
- [36] H. Shin, J. K. Lee, J. Kim, and J. Kim, "Continual learning with deep generative replay," in *Advances in Neural Information Processing Systems*, 2017.
- [37] C. Wu, L. Herranz, X. Liu, Y. Wang, J. van de Weijer, and B. Raducanu, "Memory replay gans: Learning to generate new categories without forgetting," in *Advances in Neural Information Processing Systems*, 2018.
- [38] G. E. Hinton and D. C. Plaut, "Using fast weights to deblur old memories," in *Proceedings of the 9th Annual Conference of the Cognitive Science Society*, 1987.
- [39] A. R. Gardner-Medwin, "Doubly modifiable synapses: A model of short and long term auto-associative memory," *Proceedings of the Royal Society B: Biological Sciences*, Nov. 1989.
- [40] J. Ba, G. E. Hinton, V. Mnih, J. Z. Leibo, and C. Ionescu, "Using fast weights to attend to the recent past," in *Advances in Neural Information Processing Systems*, 2016.
- [41] T. Miconi, A. Rawal, J. Clune, and K. O. Stanley, "Backpropamine: training self-modifying neural networks with differentiable neuromodulated plasticity," in *International Conference on Learning Representations (ICLR)*, 2019.
- [42] M. J. Chadwick, D. Hassabis, N. Weiskopf, and E. A. Maguire, "Decoding individual episodic memory traces in the human hippocampus," *Current Biology*, Mar. 2010.
- [43] A. C. Schapiro, N. B. Turk-Browne, M. M. Botvinick, and K. A. Norman, "Complementary learning systems within the hippocampus: a neural network modelling approach to reconciling episodic memory with statistical learning," *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 2017.
- [44] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," in *IEEE Intelligent Signal Processing*, 2001.
- [45] C. Zeno, I. Golan, E. Hoffer, and D. Soudry, "Task agnostic continual learning using online variational bayes," *CoRR*, 2018.
- [46] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms," *CoRR*, 2017.
- [47] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng, "Reading digits in natural images with unsupervised feature learning," in *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [48] A. Krizhevsky, "Learning multiple layers of features from tiny images," Tech. Rep., 2009.

APPENDIX

A. Details on Hyperparameters and Experimental Setup

Permuted MNIST: We train the network on a sequence of tasks $T_{n=1:10}$ with mini-batches of size 64 and optimized using plain SGD with a learning rate of 0.01 for 20 epochs on each task. The regularization hyperparameter λ for each of the task-specific consolidation methods is set to $\lambda = 100$ for Online EWC [20], $\lambda = 0.1$ for SI [21] and $\lambda = 0.1$ for MAS [22]. We note that for the SI method, λ refers to the parameter c in the original work [21] but we use λ to keep the notation consistent across other task-specific consolidation methods. In SI, the damping parameter, ξ , was set to 0.1.

Imbalanced Permuted MNIST: For each task, we artificially removed training samples from each class in the original MNIST dataset [44] based on some random probability. For each class and each task, we draw a different removal probability from a standard uniform distribution $U(0, 1)$, and then remove each sample from that class with that probability. The distribution of classes in each dataset corresponding to tasks $T_{n=1:10}$. The λ for each of the task-specific consolidation methods is $\lambda = 400$ for Online EWC [20], $\lambda = 1.0$ for SI [21] and $\lambda = 0.1$ for MAS [22]. In SI, the damping parameter, ξ , was set to 0.1. Across all experiments, we maintained the same random probabilities determined by a single seed to artificially remove training samples from each class.

Split MNIST: For the Split MNIST experiments shown in Figure 2b, the regularization hyperparameter λ for each of the task-specific consolidation methods is $\lambda = 400$ for Online EWC [20], $\lambda = 1.0$ for SI [21] and $\lambda = 1.5$ for MAS [22]. In SI, the damping parameter, ξ , was set to 0.001. We train the network on a sequence of $T_{n=1:5}$ tasks with mini-batches of size 64 and optimized using plain SGD with a fixed learning rate of 0.01 for 10 epochs for each task.

Vision Datasets Mixture: For the 5-Vision Datasets Mixture experiments, the regularization hyperparameter λ for each of the task-specific consolidation methods is $\lambda = 100$ for Online EWC [20], $\lambda = 0.1$ for SI [21] and $\lambda = 1.0$ for MAS [22]. In SI, the damping parameter, ξ , was set to 0.1.