

---

# Continual Learning with Self-Organizing Maps

---

**Pouya Bashivan, Martin Schrimpf, Robert Ajemian**

McGovern Institute for Brain Research and  
Department of Brain and Cognitive Sciences  
Massachusetts Institute of Technology  
Cambridge, MA, USA  
{bashivan, mschrimpf, ajemian}@mit.edu

**Irina Rish, Matthew Riemer, Yuhai Tu**

IBM Research and MIT-IBM Watson AI Lab  
Yorktown Heights, NY, USA  
{rish, mdriemer, yuhai}@us.ibm.com

## Abstract

Despite remarkable successes achieved by modern neural networks in a wide range of applications, these networks perform best in domain-specific stationary environments where they are trained only once on large-scale controlled data repositories. When exposed to non-stationary learning environments, current neural networks tend to forget what they had previously learned, a phenomena known as *catastrophic forgetting*. Most previous approaches to this problem rely on memory replay buffers which store samples from previously learned tasks, and use them to regularize the learning on new ones. This approach suffers from the important disadvantage of not scaling well to real-life problems in which the memory requirements become enormous. We propose a memoryless method that combines standard supervised neural networks with self-organizing maps to solve the continual learning problem. The role of the self-organizing map is to adaptively cluster the inputs into appropriate task contexts – without explicit labels – and allocate network resources accordingly. Thus, it selectively routes the inputs in accord with previous experience, ensuring that past learning is maintained and does not interfere with current learning. Our method is intuitive, memoryless, and performs on par with current state-of-the-art approaches on standard benchmarks.

## 1 Introduction

Learning algorithms typically operate on the entire training data set, in an offline mode, and may require costly re-training from scratch when new data becomes available. In contrast, humans learn continuously ("online"), adapting to the environment while leveraging past experiences. In this setting of *continual learning*, an agent is presented with a stream of samples (input-output pairs) from a non-stationary data distribution, such as a sequence of different classification tasks. The agent has to learn the input-output relationships associated with different data distributions in order to be able to adapt well to new tasks without forgetting the tasks learned previously, i.e. to avoid the *catastrophic forgetting* [7] problem in continual learning. Several types of approaches to this problem have been proposed in the past that introduce structural robustness [2, 9], regularize the parameters [10, 11, 4] or utilize memory buffers [6, 8]. While the memory-based methods are among the most successful methods, these approaches generally utilize explicit task labels when storing instances of past learning and they also don't scale well to real life problems.

Herein, we propose to capture the essence of the multimodal data distributions encountered in continual learning via the unsupervised learning mechanism known as Self-Organizing Map (SOM).

The map learns simultaneously with a supervised feedforward neural net in such a way that the SOM routes each input sample to the most relevant part of the network. Unlike previous methods, SOMs do not require explicit information about the change in task, or an explicit memory of previous samples, while still achieving performance levels close to the current state-of-art on several benchmarks.

## 2 Self-Organizing Maps for Continual Learning

Self-Organizing Maps (SOMs) [5] are a type of artificial neural network which reduce the input from a high-dimensional ( $N$ ) space to a low-dimensional representation without supervision. Unlike supervised training via error-correction (backpropagation), self-organizing maps employ unsupervised competitive learning in order to map similar input vectors to physically nearby nodes in the map layer. Namely, each SOM node will be associated with an  $N$ -dimensional weight vector, and, at each iteration, the next input sample will be assigned to the most similar network node called the Best Matching Unit (BMU). Trainable weights associated with the BMU along with weights of its topographical neighbors (determined by neighborhood radius  $\sigma$ ) are tuned to reduce the error between the weights and the input, (proportional to a learning rate  $\alpha$ ) so that the neighborhood becomes more similar to the input. This allows the SOM to only adapt parts of its parameters in response to each observed input. To stabilize the learning, both  $\alpha$  and  $\sigma$  parameters are decayed during learning (with an exponential decaying rate of  $\tau$ ).

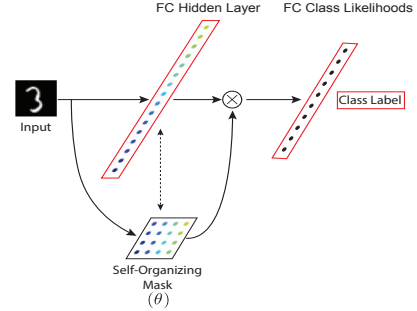


Figure 1: **SOMLP setup.** Red borders indicate supervised training; SOM is trained in an unsupervised manner.

SOMs has been previously used as a potential solution to the catastrophic forgetting problem [1] but it has been shown that the proposed algorithms often fail on small-scale benchmarks like MNIST-permutations and rotations [3]. Here we introduce the Self-Organized Multi-Layer Perceptron (SOMLP) where a SOM layer is used in parallel with a fully-connected layer, as illustrated in Figure 1. The SOM receives the same input as the fully-connected layer and is trained without supervision on the network inputs (Algorithm 2.1). During the training it learns the input distribution for each task and, most importantly, a 2D map of relationships between the tasks. For every input, an output mask  $\Gamma$  is computed based on the Euclidean distances between the input and SOM weight vectors, as shown in Algorithm 2.1 (tuned with hyperparameter  $\epsilon$ ), and is multiplied by the output of the fully connected layer (Figure 1). This allows the network to allocate resources (i.e. nodes) to learning input-output associations for each task, and mask irrelevant nodes. In addition, the SOM shares nodes between similar tasks while using separate nodes for dissimilar tasks. The overall training procedure is shown in Algorithm 2.2, which includes the option of pretraining SOM weights on unlabeled data from different tasks (if available).

## 3 Experiments

We evaluated SOMLP against three other baseline methods: a naive Multi-Layer Perceptron (MLP), Elastic Weight Consolidation (EWC) [4] and Gradient Episodic Memory (GEM) [6]. We evaluated all methods on two standard benchmark datasets, *MNIST-permutations* and *MNIST-rotations* [6]. *MNIST-permutations* consists of 20 random pixel permutations of the standard MNIST dataset while *MNIST-rotations* contains 20 rotations of MNIST digits between 0 and 180 degrees.

All experiments were performed on a fully-connected network with a single hidden layer with  $n_{h1}$  units, varying across the experiments, and an output layer with  $n_{h2} = 10$  units. For SOMLP, the hidden layer is augmented with a SOM of the same size and the output of  $n_{h1}$  is multiplied by the SOM output before being fed to the next layer. For a fair comparison with respect to the numbers of parameters, we used  $n_{h1} = 3200$  for MLP, EWC and GEM and  $n_{h1} = 1600$  for SOMLP<sup>1</sup>.

All networks were trained for 1 epoch on the training set, learning 20 tasks in sequence. For each task, the network is presented with 60,000 samples. The hyperparameter settings used in the experiments are summarized in supplementary Table S1.

<sup>1</sup>The total number of parameters for MLP, EWC and GEM is  $2,540,800 = 3200 \times 784 + 3200 \times 10$ ; and for SOMLP  $2,524,800 = 1600 \times 2 \times 784 + 1600 \times 10$

---

**Algorithm 2.1** SOM update

---

**Given:** input  $x$  and SOM weights  $\theta$  at step  $t$ ,  
nodes coordinate matrix  $L$ , batch size  $N_{bs}$ ,  
number of nodes  $N_n$ , and hyperparameters  
 $\alpha, \sigma, \epsilon, \tau$ .

- 1:  $\bar{x} = \frac{1}{N_{bs}} \sum_{i \in \text{batch}} x_i$   $\triangleright$  average over batch  
 $\triangleright$  Find best matching unit (BMU) for  $\bar{x}$
- 2:  $i^* = \operatorname{argmin}_i (\|\bar{x} - \theta_i\|)$
- 3: **for**  $i = 1, 2, \dots, N_n$  **do**  $\triangleright$  compute
- 4:  $D_i = \|L_i - L_{i^*}\|$   $\triangleright$  distances to BMU
- 5: **end**
- 6:  $\phi = e^{D/\sigma^2}$   $\triangleright$  neighborhood mask
- 7:  $\Gamma = e^{-\|x - \theta\|/\epsilon}$   $\triangleright$  output mask
- 8:  $\alpha \leftarrow \alpha e^{-\tau t / N_{steps}}$   $\triangleright$  Adjust  $\sigma$  and  $\alpha$
- 9:  $\sigma \leftarrow \sigma e^{-\tau t / N_{steps}}$
- 10:  $\theta_i \leftarrow \theta_i + \alpha \phi(x - \theta_i)$   $\triangleright$  Update weights

---



---

**Algorithm 2.2** Training procedure

---

**Given:** Training datasets  $(X_t, Y_t)$  for each  
task  $t$  out of  $N_T$  tasks, and data subsets  
 $(X_t^{pr}, Y_t^{pr})$  available for SOM pretraining.

**Initialize:** SOM and MLP weights randomly.  
 $\triangleright$  Pretrain sequentially w/ unlabeled data

- 1: **for**  $t = 1, 2, \dots, N_T$  **do**
- 2: **for** batch in  $X_t^{pr}$  **do**
- 3: Update SOM weights using Alg. 2.1
- 4: **end**
- 5: **end**  
 $\triangleright$  Train sequentially using labeled data
- 6: **for**  $t = 1, 2, \dots, N_T$  **do**
- 7: **for** batch in  $(X_t, Y_t)$  **do**
- 8: Update SOM weights using Alg. 2.1
- 9: Update MLP weights using SGD
- 10: **end**
- 11: **end**

---

In addition, EWC and GEM methods require additional memory slots. EWC uses a memory buffer to save samples before computing the Fisher matrix before moving to the next task. On the other hand, GEM uses a memory buffer per task to retain samples for experience replay. A scalar "memory strength" additionally dictates how much to weigh previous samples. In our experiments, we varied the number of memory slots between 8 to 5120 for EWC with a memory strength of 3, and 256 to 5120 for GEM with a memory strength of 0.5. An advantage of SOMLP is that it does not require a sample memory buffer.

In our approach, we used two methods for pretraining SOM weights. The first one pretrains the weights on the unlabeled MNIST training set, with large neighborhood values (SOMLP<sub>s</sub>), while for the second one, the weights are pretrained on 10% of the training set from all the tasks in the benchmark (SOMLP<sub>m</sub>). We only use the input images for pretraining.

## 4 Results

Table 1 shows the networks' performances on each benchmark dataset after learning a battery of 20 tasks. On MNIST-permutations, SOMLP<sub>m</sub> performed on par with EWC but lower than GEM

Table 1: Comparison of average performances on benchmark datasets. Number of parameters in each network were matched by changing the network size.

Network	Memory Size	#Parameters (M)	Performance (%) (MNIST-Permutations)	Performance (%) (MNIST-Rotations)
MLP	0	2.5	29±0.6	64.4±0.4
EWC-Low Mem	8	2.5	84.8	59.1
EWC-High Mem	5120	2.5	83.6	49.8
GEM-Low Mem	320	2.5	68.3	71.5
GEM-High Mem	5120	2.5	93.3	92.6
SOMLP <sub>s</sub>	0	2.5	65.9±5.2	59.1±6.6
SOMLP <sub>m</sub>	0	2.5	85.4±1.8	79.1±0.8

with large memory buffers, whereas it outperformed EWC on MNIST-rotations but still was lower than GEM. On the other hand, in the low-memory case that the size of available memory is reduced, SOMLP<sub>m</sub> performs better than all baseline methods. In contrast, SOMLP<sub>s</sub> only performed better than EWC on MNIST-rotations. Naive MLP worked fairly well on MNIST-rotations but lower than all methods on MNIST-permutations.

It should be noted that both EWC and GEM make use of a "sample memory buffer" as well as explicit information about when the task is switched while SOMLP requires neither a memory buffer

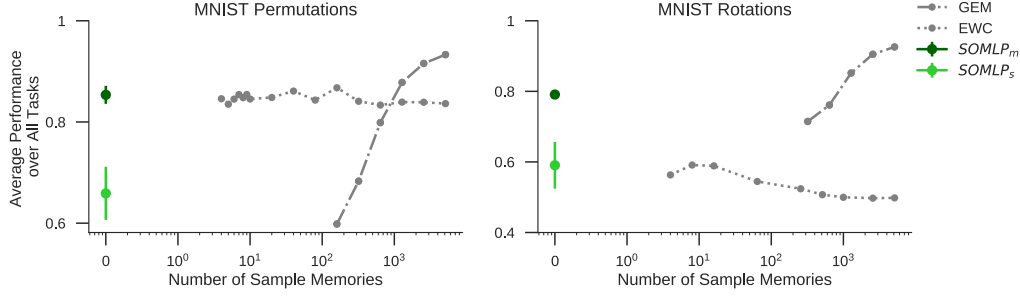


Figure 2: **Sample memory size effect on performance.** Both GEM and EWC depend on the right amount of memory. For GEM, more memory generally leads to better performance, whereas the optimal EWC memory size seems to depend on the dataset. SOMLP does not require sample memory.

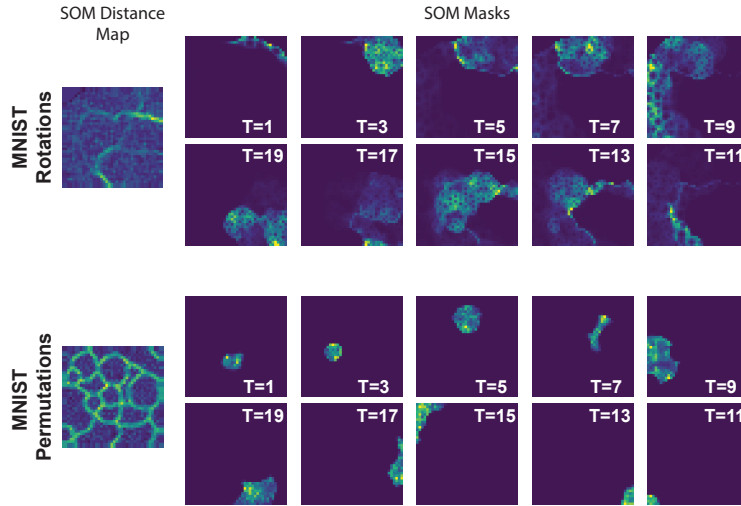


Figure 3: **Learned SOM.** (left) Euclidean distance between SOM nodes for each benchmark during a single run. (right) SOM masks found for different tasks during a single run on each benchmark.

nor knowledge of task switching. We examined the effect of the amount of sample memory on mean performance<sup>2</sup>. We found that the performance of GEM on both benchmarks is significantly reduced with smaller memory sizes. EWC performance is almost constant in the MNIST-permutations benchmark and even decreases with enlarged memory buffer in MNIST-rotations (Figure 2).

We also examined the mean performance during sequential learning of all tasks (Figure S1) on each dataset. As more tasks are learned, EWC and SOMLP<sub>s</sub> experience some degree of forgetting and therefore the mean performance decreases with more tasks. While both GEM and SOMLP<sub>m</sub> are able to maintain the same level of mean performance throughout, SOMLP<sub>m</sub> is consistently below GEM’s level of performance. This is potentially due to lower network capacity in our SOMLP approach because of (1) smaller network size and (2) per-task SOM masks that limits the resources available to learn each task.

The learned feature maps and per-task masks for each benchmark are shown in Figure 3. In MNIST-permutations, because of the random pixel permutations in each task, masks corresponding to each task are independent of each other. Conversely, in MNIST-rotations, the learned masks share nodes between tasks that are more similar (i.e. slightly rotated tasks) and use independent nodes for more dissimilar tasks (e.g. tasks that are separated by larger rotation angles).

**Acknowledgements:** This research was supported by the MIT-IBM Watson AI Lab.

<sup>2</sup>We note that EWC additionally stores the Fisher matrix associated with each task which we did not account for in our memory requirement estimate.

## References

- [1] Alexander Gepperth and Cem Karaoguz. A Bio-Inspired Incremental Learning Architecture for Applied Perceptual Problems. *Cognitive Computation*, 8(5):924–934, 2016.
- [2] Ian J 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.
- [3] Ronald Kemker, Marc McClure, Angelina Abitino, Tyler Hayes, Christopher Kanan, and Rochester Ny. Measuring Catastrophic Forgetting in Neural Networks. *arXiv*, 2017.
- [4] James Kirkpatrick, Pascanua Razvan, Neil Rabinowitz, Joel Venessa, Guillaume Desjardins, Andrei A. Rusua, Kieran Milana, John Quana, Tiago Ramalhoa, Agnieszka Grabska-Barwinskaa, Demis Hassabisa, Claudia Clopathb, Dharshan Kumarana, and Raia Hadsella. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Science (PNAS)*, 2017.
- [5] Teuvo Kohonen. The Self-Organizing Map. *Proceedings of the IEEE*, 78(9):1464–1480, 1990.
- [6] David Lopez-Paz and Marc’ Aurelio Ranzato. Gradient Episodic Memory for Continual Learning. *Neural Information Processing Systems*, 2017.
- [7] 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.
- [8] Pablo Sprechmann, Siddhant M. Jayakumar, and Jack W. Rae. Memory-based Parameter Adaptation. *ICLR*, pages 1–13, 2018.
- [9] Rupesh K Srivastava, Jonathan Masci, Sohrab Kazerounian, Faustino Gomez, and Jürgen Schmidhuber. Compete to compute. *Neural Information Processing Systems*, pages 2310–2318, 2013.
- [10] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? *Neural Information Processing Systems*, pages 3320–3328, 2014.
- [11] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual Learning Through Synaptic Intelligence. *arXiv*, 2017.

## Supplementary Material

Table S1: List of hyperparameter choices for SOMLP and baseline method.

Hyperparameter	SOMLP <sub>s</sub>		SOMLP <sub>m</sub>		EWC		GEM	
	Permutations	Rotations	Permutations	Rotations	Permutations	Rotations	Permutations	Rotations
$\alpha$	0.5	3	1	3	-	-	-	-
$\sigma$	4	4	2	3	-	-	-	-
$\epsilon$	0.5	0.5	0.5	0.5	-	-	-	-
$\tau$	2	3	15	3	-	-	-	-
Batch Size	10	100	10	100	10	10	10	10
Learning Rate	0.01	0.02	0.01	0.02	0.1	0.01	0.1	0.1

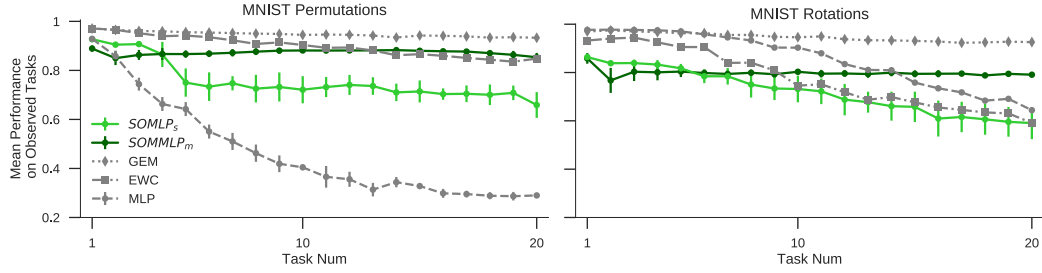


Figure S1: **Performance during task learning.** Average performance on observed tasks during sequential learning of different task-sets decreases for a standard MLP, while continual learning methods maintain performance with varying degrees of success.