

# Learning to Remember: A Synaptic Plasticity Driven Framework for Continual Learning

Oleksiy Ostapenko<sup>§,\*</sup>, Mihai Puscas<sup>¶,\*</sup>, Tassilo Klein<sup>\*</sup>, Patrick Jähnichen<sup>§</sup>, Moin Nabi<sup>\*</sup>

<sup>§</sup>Humboldt-Universität zu Berlin, <sup>¶</sup>University of Trento, <sup>\*</sup>SAP ML Research

## Abstract

*Models trained in the context of continual learning (CL) should be able to learn from a stream of data over an undefined period of time. The main challenges herein are: 1) maintaining old knowledge while simultaneously benefiting from it when learning new tasks, and 2) guaranteeing model scalability with a growing amount of data to learn from. In order to tackle these challenges, we introduce Dynamic Generative Memory (DGM) - synaptic plasticity driven framework for continual learning. DGM relies on conditional generative adversarial networks with learnable connection plasticity realized with neural masking. Specifically, we evaluate two variants of neural masking: applied to (i) layer activations and (ii) to connection weights directly. Furthermore, we propose a dynamic network expansion mechanism that ensures sufficient model capacity to accommodate for continually incoming tasks. The amount of added capacity is determined dynamically from the learned binary mask. We evaluate DGM in the continual class-incremental setup on visual classification tasks.*

## 1. Introduction

Conventional Deep Neural Networks (DNN) fail to continually learn from a stream of data while maintaining knowledge. Specifically, reusing old knowledge in new contexts poses a severe challenge. Generally, there are several fundamental obstacles on the way to a continually trainable AI system: the problem of forgetting when learning from new data (catastrophic forgetting), lack of model scalability, i.e. the inability to scale up the model’s size with a continuously growing amount of training data, and finally inability to transfer knowledge across tasks.

Several recent approaches [9, 34, 2, 1] try to mitigate forgetting in ANNs while simulating synaptic plasticity directly in the task solving network. It is noteworthy that these methods typically tackle the task-incremental scenario, i.e. a separate classifier is trained to make predictions about each task. This further implies the availability of oracle knowledge of the task label at inference time. Such eval-

uation is often referred to as multi-head evaluation in which the task label is associated with a dedicated output head. Alternatively, other approaches rely on single-head evaluation [22, 2]. Here, the model is evaluated on all classes observed during the training, no matter which task they belong to. While single-head evaluation does not require oracle knowledge of the task label, it also does not reduce the output space of the model to the output space of the task. Thus single-head evaluation represents a harder, yet more realistic setup. Single-head evaluation is predominantly used in class-incremental setup, in which every newly introduced data batch contains examples of one to many new classes.

As opposed to the task-incremental situation, models in class-incremental setup typically require the previously learned information to be replayed when learning new categories [22, 2, 18]. The simplest way to accomplish this is by retaining and replaying real samples of previously seen categories to the task solver. However, retaining real samples has several intrinsic implications. First, it is very much against the notion of bio-inspired design, as natural brains do not feature the retrieval of information identical to originally exposed impressions [15]. Second, as pointed out by [32, 22] storing raw samples of previous data can violate data privacy and memory restrictions of real-world applications. Such restrictions are particularly relevant for the vision domain with its continuously growing dataset sizes and rigorous privacy constraints.

In this work, we address the “strict” class-incremental setup. That is, we demand a classifier to learn from a stream of data with different classes occurring at different times with no access to previously seen data, i.e. no storing of real samples is allowed. Such a scenario is solely addressed by methods relying on generative memory - a generative network is used to memorize previously seen data distributions, samples of which can be replayed to the classifier at any time. Several strategies exist to avoid catastrophic forgetting in generative networks. The most successful approaches rely on deep generative replay (DGR) [30] - repetitive retraining of the generator on a mix of synthesized samples of previous categories and real samples of new classes. In this work we propose Dynamic Generative

Memory (DGM) with learnable connection plasticity represented by parameter level attention mechanism. As opposed to DGR, DGM features a single generator that is able to incrementally learn about new tasks without the need to replay previous knowledge.

Another important factor in the continual learning setting is the ability to scale, i.e. to maintain sufficient capacity to accommodate for a continuously growing amount of information. Given invariant resource constraints, it is inevitable that with a growing number of tasks to learn, the model capacity is depleted at some point in time. This issue is again exacerbated when simulating neural plasticity with parameter level hard attention masking. In order to guarantee sufficient capacity and constant expressive power of the underlying DNN, we keep the number of "free" parameters (i.e. to which the gradient updates can be freely applied) constant by expanding the network with exactly the number of parameters that were blocked for the previous task.

Our contribution is twofold: **(a)** we introduce Deep Generative Memory (DGM) - an adversarially trainable generative network that features neural plasticity through efficient learning of a sparse attention masks for the network weights (DGMw) or layer activations (DGMa); To the best of our knowledge we are the first to introduce weight level masks that are learned simultaneously with the base network; Furthermore, we conduct it in an adversarial context of a generative model; DGM is able to incrementally learn new information during adversarial training *without the need to replay previous knowledge to its generator*. **(b)** We propose an adaptive network expansion mechanism, facilitating resource efficient continual learning. In this context, we compare the proposed method to the state-of-the-art approaches for continual learning. Finally, we demonstrate that DGMw accommodates for higher efficiency, better parameter reusability and slower network growth than DGMa.

## 2. Related Work

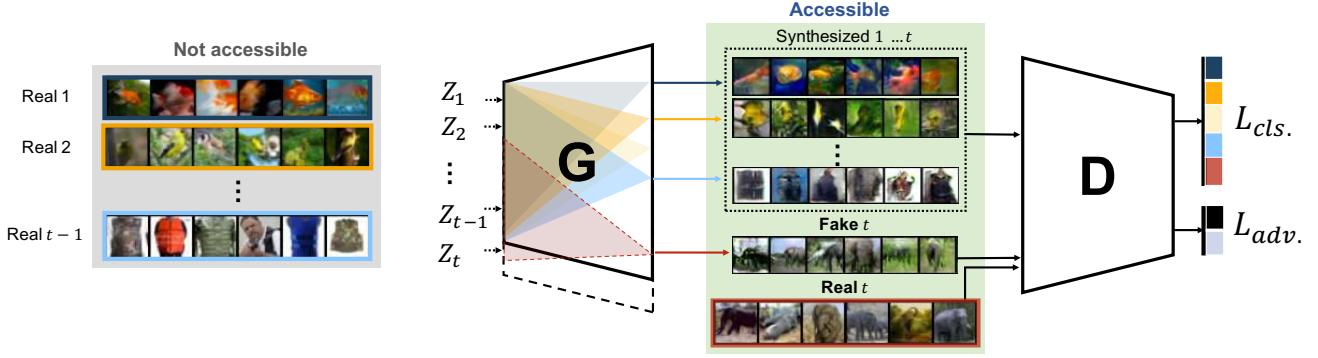
Among the first works dealing with catastrophic forgetting in the context of lifelong learning are [4, 16, 21], who tackle this problem by employing shallow neural networks, whereas our method makes use of modern deep architectures. Lately, a wealth of works dealing with catastrophic forgetting in context of DNNs have appeared in the literature, see e.g., [9, 34, 12, 29, 1, 24]. Thus, EWC [9] and RWalk [2] rely on Fisher's information to identify parameters that carry most of the information about previously learned tasks, and apply structural regularization to "discourage" change of these parameters. [34] and [1] identify important parameter segments based on the sensitivity of the loss or the learned prediction function to changes in the parameter space. Instead of relying on "soft" regularization techniques, [29] and [26] propose to dedicate separate parameter subspaces to separate tasks. Serrà et al. [29] pro-

pose a hard attention to the task (HAT) mechanism. HAT finds dedicated parameter subspaces for all tasks in a single network while allowing them to mutually overlap. The optimal solution is then found in the corresponding parameter subspace of each task. All of these methods have been proposed for a "task-incremental learning" setup. In our work we specifically propose a method to overcome catastrophic forgetting within the "class-incremental" setup. Notably, a method designed for class-incremental learning can be generally applied in a task-incremental setup.

Several continuous learning approaches [22, 18, 8], address catastrophic forgetting in the class-incremental setting, i.e. by storing raw samples of previously seen data and making use of them during the training on subsequent tasks. Thus, iCarl [22] proposes to find  $m$  most representative samples of each class whose mean feature space most closely approximates the entire feature space of the class. The final classification task is done by the means of the nearest mean-of-exemplars classifier.

Recently, there has been a growing interest in employing deep generative models for memorizing previously seen data distributions instead of storing old samples. [30, 31] rely on the idea of generative replay, which requires retraining the generator at each time step on a mixture of synthesized images of previous classes and real samples from currently available data. However, apart from being inefficient in training, these approaches are severely prone to "semantic drifting". Namely, the quality of images generated during every memory replay highly depends on the images generated during previous replays, which can result in loss of quality and forgetting over time. In contrast, we propose to utilize a single generator that is able to incrementally learn new information during the normal adversarial training without the need to replay previous knowledge. This is achieved through efficiently learning a sparse mask for the learnable units of the generator network.

Similar to our method, [28] proposed to avoid retraining the generator at every time-step on previous classes by applying EWC [9] in the generative network. We pursue a similar goal with the key difference of utilizing a hard attention mechanism similar to the one described by [29, 13, 14]. All three approaches make use of the techniques originally proposed in the context of binary-valued networks [3]. Herein, binary weights are specifically learned from a real-valued embedding matrix that is passed through a binarization function. To this end, [13, 14] learn to mask a pre-trained network without changing the weights of the base networks, whereas [29] (HAT) features binary mask-learning for the layer activations simultaneously to the training of the base network. While DGMa features HAT-like layer activation masking, DGMw accomplishes binary mask learning directly on the weights of the generator. Other works propose to use non-binary filters to define a



**Figure 1:** Dynamic Generative Memory: auxiliary output of  $D$  is trained on the real samples of the current task  $t$  and synthesized sample of previously seen tasks  $1 \dots t-1$ . Adversarial training is accomplished with real and fake samples of the current task. Connection plasticity simulated with binary mask applied to the generators weights or activations is learned simultaneously to the adversarial training.

new task solving network in terms of a linear combination of the parameters of a fixed base network [24, 23].

Similarly to [33], we propose to expand the capacity of the employed base network, in our case the samples generator. The expansion is performed dynamically with an increasing amount of attained knowledge. However, [33] propose to keep track of the *semantic drift* in every neuron, and then expand the network by duplicating neurons that are subject to sharp changes. In contrast, we compute weights importance concurrently during the course of network training by modeling the neuron behavior using learnable binary masks. As a result, our method explicitly does not require any further network retraining after expansion.

Other approaches like [8, 7, 27] try to explicitly model short and long term memory with separate networks. In contrast to these methods, our approach does not explicitly keep two separate memory locations, but rather incorporates it implicitly in a single memory network. Thus, the memory transfer occurs during the binary mask learning from non-binary (short term) to completely binary (long term) values.

### 3. Dynamic Generative Memory

Adopting the notation of [2], let  $S_t = \{(x_i^t, y_i^t)\}_{i=1}^{n_t}$  denote a collection of data belonging to the task  $t \in T$ , where  $x_i^t \in \mathcal{X}$  is the input data and  $y_i^t \in \mathcal{Y}^t$  are the ground truth labels. While in the non-incremental setup the entire dataset  $S = \cup_{t=0}^{|T|} S_t$  is available at once, in an incremental setup it becomes available to the model in chunks  $S_t$  specifically only during the learning of task  $t$ . Thereby,  $S_t$  can be composed of a collection of items from different classes, or even from a single class only. Furthermore, at the test time the output space covers all the labels observed so far featuring the single head evaluation:  $\mathcal{Y}^t = \cup_{j=1}^t \mathcal{Y}^j$ .

We consider a continual learning setup, in which a task solving model  $D$  has to learn its parameters  $\theta^D$  from the data  $S_t$  being available at the learning time of task  $t$ . Task

solver  $D$  should be able to maintain good performance on all classes  $\mathcal{Y}^t$  seen so far during the training. A conventional DNN, while being trained on  $S_t$ , would adapt its parameters in a way that exhibits good performance solely on the labels of the current task  $\mathcal{Y}^t$ , the previous tasks would be forgotten. To overcome this, we introduce a Generative Memory component  $G$ , who's task is to memorize previously seen data distributions. As visualized in Fig. 1, samples of the previously seen classes are synthesized by  $G$  and replayed to the task solver  $D$  at each step of continual learning to maintain good performance on the entire  $\mathcal{Y}^t$ . We train a generative adversarial network (GAN)[5] and a sparse mask for the weights of its generator simultaneously. The learned masks model connection plasticity of neurons, thus avoiding overwriting of important units by restricting SGD updates to the parameter segments of  $G$  that exhibit free capacity.

#### 3.1. Learning Binary Masks

We consider a generator network  $G_{\theta^G}$  consisting of  $L$  layers, and a discriminator network  $D_{\theta^D}$ . In our approach,  $D_{\theta^D}$  serves as both: a discriminator for generated fake samples of the currently learned task ( $L_{adv.}$ ) and as a classifier for the actual learning problem ( $L_{cls.}$ ) following the AC-GAN [19] architecture. The system has to continually learn  $T$  tasks. During the SGD based training of task  $t$ , we learn a set of binary masks  $M^t = [m_1^t, \dots, m_L^t]$  for the weights of each layer. Output of a fully connected layer  $l$  is obtained by combining the binary mask  $m_l^t$  with the layer weights:

$$y_l^t = \sigma_{act}[(m_l^t \circ W_l)^T x], \quad W_l \in \mathbb{R}^{n \times p}, \quad (1)$$

for  $\sigma_{act}$  being some activation function.  $W_l$  is the weight matrix applied between layer  $l$  and  $l-1$ ,  $\cdot \circ \cdot$  corresponds to the Hadamard product. In DGMw  $m_l^t$  is shaped identically to  $W_l$ , whereas in case of DGMa the mask  $m_l^t$  is shaped as  $1 \times p$  and should be expanded to the size of  $W_l$ . Extension to more complex models such as e.g. CNNs is straightforward.

A single binary mask for a layer  $l$  and task  $t$  is given by:

$$m_l^t = \sigma(s \cdot e_l^t), \quad (2)$$

where  $e_l^t$  is a real-valued mask embeddings matrix,  $s$  is a positive scaling parameter  $s \in \mathbb{R}_+$ , and  $\sigma$  a thresholding function  $\sigma : \mathbb{R} \rightarrow [0, 1]$ . Similarly to [29] we use the sigmoid function as a pseudo step-function to ensure gradient flow to the embeddings  $e$ . In training of DGMw, we anneal the scaling parameter  $s$  incrementally during epoch  $i$  from  $1/s_{max}^i$  to  $s_{max}^i$  (local annealing).  $s_{max}^i$  is similarly adjusted over the course of  $I$  epochs from  $1/s_{max}$  to  $s_{max}$  (global annealing with  $s_{max}$  being a fixed meta-parameter). The annealing the scheme is largely adopted from [29]:

$$s_{max}^i = \frac{1}{s_{max}} + (s_{max} - \frac{1}{s_{max}}) \frac{i-1}{I-1} \quad (3)$$

$$s = \frac{1}{s_{max}^i} + (s_{max}^i - \frac{1}{s_{max}^i}) \frac{b-1}{B-1}. \quad (4)$$

Here  $b \in \{1, \dots, B\}$  is the batch index and  $B$  the number of batches in each epoch of SGD training. DGMa only features global annealing of  $s$ , as it showed better performance.

In order to prevent the overwriting of the knowledge related to previous classes in the generator network, gradients  $g_l$  w.r.t. the weights of each layer  $l$  are multiplied by the reverse of the cumulated mask  $m_l^{\leq t}$ :

$$g_l' = [1 - m_l^{\leq t}]g_l, \quad m_l^{\leq t} = \max(m_l^t, m_l^{t-1}), \quad (5)$$

where  $g_l'$  corresponds to the new gradient matrix and  $m_l^{\leq t}$  is the cumulated mask.

Analogously to [29], we promote sparsity of the binary mask by adding a regularization term  $R^t$  to the loss function of the AC-GAN[19] generator:

$$R^t(M^t, M^{t-1}) = \frac{\sum_{l=1}^{L-1} \sum_{i=1}^{N_l} m_{l,i}^t (1 - m_{l,i}^{\leq t})}{\sum_{l=1}^{L-1} \sum_{i=1}^{N_l} 1 - m_{l,i}^{\leq t}}, \quad (6)$$

where  $N_l$  is the number of parameters of layer  $l$ . Here, parameters that were reserved previously are not penalized, promoting reuse of units over reserving new ones.

### 3.2. Dynamic Network Expansion

As discussed by [33], significant domain shift between tasks leads to rapid network capacity exhaustion, manifesting in decreasing expressive power of the underlying network and ultimately in catastrophic forgetting. In case of DGM this effect will be caused by decreasing number of “free” parameters over the course of training due to parameter reservation. To avoid this effect, we take measures to ensure constant number of free parameters for each task.

*DGMa.* Consider a network layer  $l$  with an input vector of size  $n$ , an output vector of size  $p$ , and the mask

$m_l^1 \in [0, 1]^{1 \times p}$  initialized with mask elements  $m_l^1$  of all neurons of the layer set to 0.5 (real-valued embeddings  $e_l^1$  are initialized with 0). After the initial training cycle on task 1, the number of free output neurons in layer  $l$  will decrease to  $p - \delta_1$ , where  $\delta_t$  is the number of neurons reserved for a generation task  $t$ , here  $t = 1$ . After the training cycle, the number of output neurons  $p$  of the layer  $l$  will be expanded by  $\delta_1$ . This guarantees that the free capacity of the layer is kept constant at  $p$  neurons for each learning cycle.

*DGMw.* In case of DGMw, after the initial training cycle the number of free weights will decrease to  $np - \delta_1$ , with  $\delta_1$  corresponding to the number of weights reserved for the generation task 1. The number of output neurons  $p$  is expanded by  $\delta_1/n$ . The number of free weights of the layer is kept constant, which can be verified by the following equation:  $(p + \delta_t/n)n - \delta_t = np$ . In practice we extend the number of output neurons by  $\lceil \delta_t/n \rceil$ . The number of free weight parameters in layer  $l$  is thus either  $np$ , if  $\delta_t/n \in \mathbb{Z}$ , or  $np + p$ , otherwise.

### 3.3. Training of DGM

The proposed system combines the joint learning of three tasks: a generative, a discriminative and finally, a classification task in the strictly class-incremental setup.

Using task labels as conditions, the generator network must learn from a training set  $X_t = \{X_1^t, \dots, X_N^t\}$  to generate images for task  $t$ . To this end, AC-GAN’s conditional generator synthesizes images  $x_t = G_{\theta^G}(t, z, M_t)$ , where  $\theta^G$  represents the parameters of the generator network,  $z$  denotes a random noise vector. The parameters corresponding to each task are optimized in an alternating fashion. As such, the generator optimization problem can be seen as minimizing  $\mathcal{L}_G = \mathcal{L}_s^t - \mathcal{L}_c^t + \lambda_{RU} R^t$ , with  $\mathcal{L}_c$  a cross entropy classification loss calculated on the the auxiliary output,  $\mathcal{L}_s$  a discriminative loss function used on the adversarial output layer of the network (implemented to be compliant with architectural requirements of WGAN[6]), and  $R^t$  the regularizer term expanded upon in equation 6. To promote efficient parameter utilization, taking into consideration the proportion of the network already in use, the regularization weight  $\lambda_{RU}$  is multiplied by the ratio  $\alpha = \frac{S_t}{S_{free}}$ , where  $S_t$  is the size of the network before training on task  $t$ , and  $S_{free}$  is the number of free neurons. This ensures that less parameters are reused during early stages of training, and more during the later stages when the model already has gained a certain level of maturity.

The discriminator is optimized similarly through minimizing  $\mathcal{L}_D = \mathcal{L}_c^t + \mathcal{L}_s^t + \lambda_{GP} \mathcal{L}_{gp}^t$ , where  $\mathcal{L}_{gp}^t$  represents a gradient penalty term implemented as in [6] to ensure a more stable training process.

		MNIST (%)		SVHN(%)		CIFAR10(%)		ImageNet-50(%)	
Method		$A_5$	$A_{10}$	$A_5$	$A_{10}$	$A_5$	$A_{10}$	$A_{30}$	$A_{50}$
JT		97.66	98.10	85.30	84.82	82.20	64.20	57.35	49.88
Episodic memory	iCarl-S [22]	-	55.8	-	-	-	-	29.38	<b>28.98</b>
	EWC-S[9]	-	79.7	-	-	-	-	-	-
	RWalk-S[2]	-	82.5	-	-	-	-	-	-
	PI-S [34]	-	78.7	-	-	-	-	-	-
Generat. memory	EWC-M [28]	70.62	77.03	39.84	33.02	-	-	-	-
	DGR [30]	90.39	85.40	61.29	47.28	-	-	-	-
	MeRGAN [31]	98.19	97.00	80.90	66.78	-	-	-	-
	DGMw (ours)	98.75	96.46	<b>83.93</b>	<b>74.38</b>	<b>64.94</b>	<b>51.70</b>	<b>32.14</b>	17.82
	DGMa (ours)	<b>99.17</b>	<b>97.92</b>	81.07	66.89	62.50	50.80	25.93	15.16

**Table 1:** Comparison to the benchmark presented by [2] (episodic memory with real samples) and [31] (generative memory) of approaches evaluated in class-incremental setup. Joint training (JT) represents the upper bound. Both variants of our method are evaluated.


## 4. Experimental Results

We perform experiments measuring the classification accuracy of our system in a strictly class-incremental setup on the following benchmark datasets: MNIST [11], SVHN [17], CIFAR-10 [10], and ImageNet-50 [25]. Similarly to [22, 31, 2] we report an average accuracy ( $A_t$ ) over the held-out test sets of classes  $0 \dots t$  seen so far during the training.

**Datasets.** The MNIST and SVHN datasets are composed of 60000 and 99289 images respectively, containing digits. The main difference is in the complexity and variance of the data used. SVHN’s images are cropped photos containing house numbers and as such present varying viewpoints, illuminations, etc. CIFAR10 contains 60000 labeled images, split into 10 classes, roughly 6k images per class. Finally, we use a subset of the iILSVRC-2012 dataset containing 50 classes with 1300 images per category. All images are further resized to 32 x 32 before use.

**Implementation details.** We make use of the same architecture for the MNIST and SVHN experiments, a 3-layer DCGAN [20], with the generator’s number of parameters modified to be proportionally smaller than in [31] (10% of MeRGAN’s size for DGMw, and 44% for DGMa on MNIST and SVHN). The projection and reshape operation is performed with a convolutional layer instead of a fully connected one. For the CIFAR-10 experiments, we use the ResNet architecture proposed by [20]. For the ImageNet-50 benchmark, the discriminator features a ResNet-18 architecture. All are modified to function as an AC-GAN[19].

All datasets are used to train a classification network in an incremental way. The performance of our method is evaluated quantitatively through comparison with benchmark methods. Note that we compare our method mainly to the approaches that rely on the idea of generative memory re-





**Figure 2:** Top-5 performance of DGMw together with upper and lower performance bounds measured for ImageNet-50 benchmark. DGM+real denotes variation with different ratios of real samples added to the replay loop (25%-75% of samples being real)

play, e.g. replaying generator synthesized samples of previous classes to the task solver without storing real samples of old data. For the sake of fairness, we only consider benchmarks evaluated in class-incremental single-head evaluation setup. Hereby, to best of our knowledge [31] represent the state-of-the-art benchmark followed by [30] and [28]. Next, we relax the strict incremental setup and allow partial storage of real samples of previous classes. Here we compare to the iCarl [22], which is the state-of-the-art method for continual learning with storing real samples.

**Results.** A quantitative comparison of both variants of the proposed approach with other methods is listed in Tab. 1. We use joint training (JT) as an upper performance bound, where the task solver  $D$  is trained in a non-incremental fashion on all real samples without adversarial training being involved. The first set of methods evaluated by [2] do not adhere to the strictly incremental setup, and thus make use of stored samples, which is often referred to as ”episodic memory”. The second set of methods we com-





**Figure 3:** Masks learning dynamics. Fig. (a) illustrates the ratio of newly blocked and reused neurons over the total number of used neurons for a task  $t$ . Fig. (b) illustrates trajectories of mask value change for DGMa for a selected layer of G (bold line - layer occupation).

Method	Top-1(%)		Top-5(%)	
	$A_{30}$	$A_{50}$	$A_{30}$	$A_{50}$
JT	57.35	49.88	84.70	78.24
iCarl (K=1000)	29.38	28.98	69.98	59.49
iCarl (K=2000)	39.38	29.96	70.57	60.07
DGMw (K=1000)	36.87	18.84	69.13	43.12
DGMw (K=2000)	41.93	22.56	69.20	51.84
DGMw ( $r=0.75$ )	50.80	38.22	78.27	64.64
DGMw ( $r=0.5$ )	48.80	40.72	79.40	71.72
DGMw ( $r=0.25$ )	46.93	35.04	75.80	65.60
DGMw ( $r=0.1$ )	41.67	30.48	71.80	61.04
DGMw ( $r=0$ )	32.14	17.82	62.53	40.76

**Table 2:** Performance comparison of DGM and iCarl for different values of  $r$  and memory size  $K$ .

pare with do not store any real data samples. Our method outperforms the state of the art [28, 30] on the MNIST and SVHN benchmarks through the integration of the memory learning mechanism directly into the generator, and the expansion of said network as it saturates to accommodate new information. We yield an increase in performance over [31], a method that is based on a replay strategy for the generator and does not provide dynamic expansion mechanism of the memory network, leading to increased training time and sensitivity to semantic drift. As it can be observed for both, our method and [31], the accuracy reported between 5 and 10-tasks of the MNIST benchmark has changed a little, suggesting that for this dataset and evaluation methodology both approaches have largely curbed the effects of catastrophic forgetting. Interestingly, *DGM* and *MeRGAN* outperform JT on MNIST ( $A_5$ ) using the same architecture. This suggests that the strictly incremental training methodology forced the network to learn better generalizations compared to what it would learn given all the data.


Given the high accuracy reached on the MNIST dataset largely gives rise to questions concerning saturation, we opted to perform a further evaluation on the more visually

diverse SVHN dataset. In this context, increased data diversity translates to more difficult generation and susceptibility to catastrophic forgetting. In fact, as can be seen in Tab. 1, the difference between 5- and 10-task accuracies is significantly larger in all methods than what can be observed in the MNIST experiments. DGM strongly outperforms all other methods on the SVHN benchmark. This can be attributed primarily to the efficient network expansion that allows for more redundancy in reserving representative neurons for each task, and a less destructive joint use of neurons between tasks. Additionally, replay based methods (like [30, 31]) can be potentially prone to generation of samples that represent class mixtures, especially for classes that semantically interfere with each other. DGM is immune to this problematic, since no generative replay is involved in the generator’s training. Thus, DGM becomes more stable in the face of catastrophic forgetting.

The quality of the generated images after 10 stages of incremental training for MNIST and SVHN can be observed in Fig. 5. The generative network is able to provide an informative and diverse set of samples for all previously seen classes without catastrophic forgetting.

Finally, in the ImageNet-50 benchmark, we incrementally add 50 classes with 10 classes per step and evaluate the classification performance of DGM using single-head evaluation. The dynamics of the top-5 classification accuracy of our system is provided in Fig. 2. Looking at the qualitative results shown in Fig. 5, it can be observed that generated samples clearly feature class discriminative features which are not forgotten after incremental training on 5 tasks of the benchmark. Nevertheless, for each newly learned task the discriminator network’s classification layer is extended with 10 new outputs, making the complexity of the classification problem to grow constantly (from 10-way classification to 50-way classification). With the more complex ImageNet samples also the generation task becomes much harder than in datasets like MNIST and SVHN. These factors negatively impact the classification performance of the

task solver presented in Fig. 2, where DGMw performs significantly worse than the JT upper bound.



**Figure 4:** Network growth (numb. of neurons) in the incremental MNIST setup. Comparison of best performing DGMw and comparable DGMa (DGMw  $A_{10}$ : 96.46%, DGMa - 96.98%).

Next, we relax the strict incremental setup and allow the DGM to partially store real samples of previous classes. We compare the performance of DGM to the state-of-the-art iCarl [22]<sup>1</sup>. Noteworthy, iCarl relies only on storing real samples of previous classes introducing a smart sample selection strategy. We define a ratio of stored real and total replayed samples  $r = n_r/N$ , where  $N$  is the total number of samples replayed per class and  $n_r$  is the number of randomly selected real samples stored per each previously seen class. To keep the number of replayed samples balanced with the number of real samples of the currently observed classes,  $N$  is set to be equal to the average number of samples per class in the currently observed data chunk  $S_t$ . Furthermore, similarly to iCarl [22] we define  $K$  to be the total number of real samples that can be stored by the algorithm at any point of time. We compare DGMw with iCarl for different values of  $K$  allowing the storage of  $K/|\mathcal{Y}^t|$  samples per class.

From Tab. 2 we observe that DGM is outperformed by iCarl when no real samples are replayed (i.e.  $r = 0$ ) after 50 classes in top-1 and after 30 and 50 classes in top-5 accuracy. DGMw with  $r = 0$  outperforms iCarl with  $K = 1000$  in top-1 accuracy after 30 classes. Furthermore, we observe that adding real samples to the replay loop boosts DGM’s classification accuracy beyond the iCarl’s one. Thus, already for  $r = 0.1$  the performance of our system can be improved significantly. We now consider DGM and iCarl with the same memory size  $K$  (we test for  $K = 1000$  and  $K = 2000$ ). Here DGM outperforms iCarl in top-1 accuracy after 30 classes, and almost reaches it in Top-5 accuracy. This is largely attributed to the advantage of DGM using generated samples additionally to the stored real once. Yet, a significant performance drop is observed after learning 5 tasks ( $A_{50}$ ), where DGMw is outperformed by iCarl. This can be attributed to (a) the fact that the number of samples replayed per class decreases over time due to fixed  $K$

<sup>1</sup> We use classes of ImageNet-50 with  $32 \times 32$  resolution and the iCarl implementation under <https://github.com/srebuffi/iCaRL>

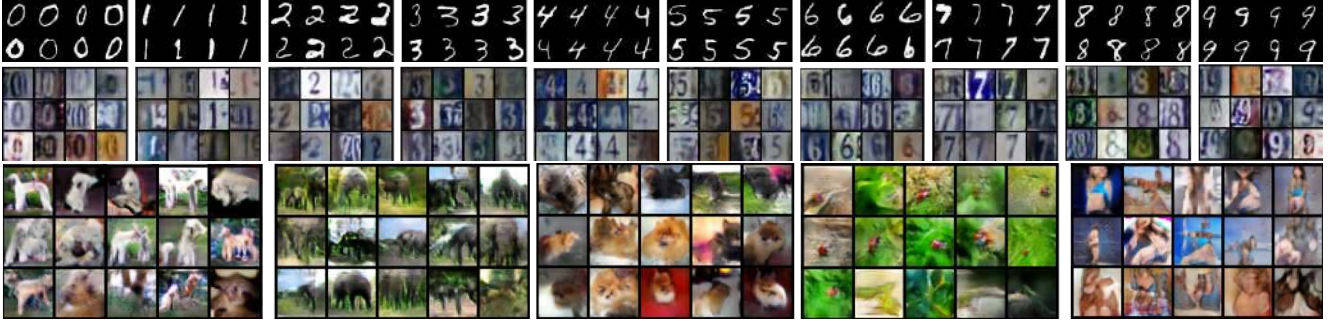
Dataset	Method	Size init.	Size final	$A_{10}(\%)$
MNIST	DGMw	5.58e+4	3.83e+5	96.46
	MeRGAN	5.18e+5	5.18e+5	97.00
SVHN	DGMw	5.58e+4	3.99e+5	74.38
	MeRGAN	5.18e+5	5.18e+5	66.78

**Table 3:** Comparison of generator’s base network’s size (number of parameters) of DGMw and MeRGAN.

and increasing number of classes (e.g. for  $K = 2000$ , 66 samples are played per class after seeing 30 classes, and 40 samples after 50 classes), as well as (b) iCarl’s smart samples selection strategy that favors samples that better approximate the mean of all training samples per class. Such samples selection strategy appears to work better in a situation where the number of real samples available per class decreases over time. It is noteworthy that iCarl’s samples selection strategy can also be applied to DGM.

**Growth Pattern Analysis.** One of the primary strengths of DGM is an efficient generator network expansion component, removing which would lead to the inability of the generator to accommodate for memorizing new task. Performance of DGM is directly related to how the network parameters are reserved during incremental learning, which ultimately depends on the generator’s ability to generalize from previously learned tasks. Fig. 4 reports network growth against the number of tasks learned. We find that learning masks directly for the layer weights (DGMw) significantly slows down the network growth. Furthermore, one can observe the high efficiency of DGM’s sub-linear growth pattern as compared to the worst-case linear growth scenario. Interestingly, as shown in Tab.3, after incrementally learning 10 classes the final number of generator’s base network’s parameters is lower than the one of the benchmarked MeRGAN [31]. More specifically, we observe the final network’s size reduction of 26% on MNIST, and 23% on SVHN as compared to MeRGAN’s fixed generator. In general, growth pattern of DGM depends on various factors: e.g. initialization size, similarity and order of classes etc.. A rather low saturation tendency of DGM’s growth pattern observed in Fig. 4 can be attributed to the fact that with growing amount of information stored in the network, selecting relevant knowledge becomes increasingly hard.

**Plasticity Evolution Analysis.** We analyze how learning is accomplished within a given task  $t$ , and how this further affects the wider algorithm. For a given task  $t$ , its binary mask  $M_t$  is initialized with the scaling parameter  $s = 1$ . Fig. 3(b) shows the learning trajectories of the mask values over the learning time of task  $t$ . Here, at task initialization of DGMa the mask is completely non-binary (all mask values are 0.5). As training progresses, the scaling parameter  $s$  is annealed, the network is encouraged to search



**Figure 5:** Images generated by DGM for MNIST(top), SVHN(middle) after training on 10 tasks, and ImageNet(bottom) after 5 tasks.

for the most efficient parameter constellation (epoch 2-10). But with most mask values near 0 (most of the units are not used, high efficiency is reached), the network’s capacity to learn is greatly curtailed. The optimization process pushes the mask to become less sparse, the number of non-zero mask values is steadily increasing until the optimal mask constellation is found, a trend observed in the segment between the epoch 10 and 55. This behaviour can be seen as a *short-term memory* formation - if learning was stopped at e.g. epoch 40 only a relatively small fraction of learnable units would be masked in a binary way, the units with non-binary mask values would be still partially overwritten during the subsequent learning resulting in forgetting. A transition from short to the *long-term memory* occurs largely within the epochs 45-65. Here the most representative units are selected and reserved by the network, parameters that have not made this transition are essentially left as unused for the learning task  $t$ . Finally, the optimal neuron constellation is optimized for the given task from epoch 60 onwards.

For a given task  $t$ , masked units (neurons in DGMa, network weights in DGMw) can be broadly divided into three types: (i) units that are not used at all (U) [masked with 0], (ii) units that are newly blocked for the task ( $NB_t$ ), (iii) units that have been reused from previous tasks ( $R_t$ ). Figure 3(a) presents the evolution of the ratio of the ( $NB_t$ ) and ( $R_t$ ) types over the total number of units blocked for the task  $t$ . Of particular importance is that the ratio of reused units is increasing between tasks, while the ratio of newly blocked units is decreasing. These trends can be justified by the network learning to generalize better, leading to a more efficient capacity allocation for new tasks.

**Memory Usage Analysis.** We evaluate the viability of generative memory usage from the perspective of required disc space. Storing the generator for the ImageNet-50 benchmark (weights and masks) corresponds to the disc space requirement of  $228MB$ . Thereby storing the preprocessed training samples of ImageNet-50 results in the required disc space of  $315MB$ . In this particular case storing the generator 27.5% more memory efficient than storing the training samples. Naturally, this effect will become more

pronounced for larger datasets.

As discussed in Sec. 4, DGMw features a more efficient network growth pattern as compared to DGMa. Yet, DGMw’s attention masks are shaped identically to the weight matrices and thus require more memory. Tab. 4 gives an overview of the required disc space for different components of DGMa and DGMw (masks are stored in a sparse form). Less total disc space is required to store DGMw’s model as compared to DGMa, which suggests that DGMw’s model growth efficiency compensates for the higher memory required for storing its masks. During the training, DGMw still exhibits a larger memory consumption, as the real-valued mask embeddings for the currently learned task must be kept in memory in a non-sparse form.

	$Size_5(MB)$		$Size_{10}(MB)$	
	DGMa	DGMw	DGMa	DGMw
Weights	6.7	1.1	14.0	2.0
Masks	0.4	3.8	0.8	9.9
Total	7.1	<b>4.9</b>	14.8	<b>11.9</b>

**Table 4:** Disc space required to store different components of DGMw and DGMa in Megabytes (MB) after 5 and 10 tasks. Compared models exhibit comparable performance on MNIST.

## 5. Conclusion

In this work we study the continual learning problem in a single-head, strictly incremental context. We propose a Dynamic Generative Memory approach for class-incremental continual learning. Our results suggest that DGM successfully overcomes catastrophic forgetting by making use of a conditional generative adversarial model where the generator is used as a memory module endowed with neural masking. We find that neural masking works more efficient when applied directly to layers’ weights instead of activations. Future work will address the limitations of the DGM including missing backward knowledge transfer and limited saturation of the network growth pattern.



## References

- [1] R. Aljundi, F. Babiloni, M. Elhoseiny, M. Rohrbach, and T. Tuytelaars. Memory aware synapses: Learning what (not) to forget. *CoRR*, abs/1711.09601, 2017.
- [2] A. Chaudhry, P. K. Dokania, T. Ajanthan, and P. H. S. Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. *CoRR*, abs/1801.10112, 2018.
- [3] M. Courbariaux, Y. Bengio, and J.-P. B. David. Training deep neural networks with binary weights during propagations. arXiv preprint. *arXiv preprint arXiv:1511.00363*, 2015.
- [4] R. M. French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135, 1999.
- [5] I. J. Goodfellow, M. Mirza, D. Xiao, A. Courville, and Y. Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*, 2013.
- [6] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*, pages 5767–5777, 2017.
- [7] N. Kamra, U. Gupta, and Y. Liu. Deep generative dual memory network for continual learning. *arXiv preprint arXiv:1710.10368*, 2017.
- [8] R. Kemker and C. Kanan. Fearnnet: Brain-inspired model for incremental learning. *arXiv preprint arXiv:1711.10563*, 2017.
- [9] 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. *CoRR*, abs/1612.00796, 2016.
- [10] A. Krizhevsky, V. Nair, and G. Hinton. The cifar-10 dataset. online: <http://www.cs.toronto.edu/kriz/cifar.html>, 2014.
- [11] Y. LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- [12] Z. Li and D. Hoiem. Learning without forgetting. *CoRR*, abs/1606.09282, 2016.
- [13] A. Mallya and S. Lazebnik. Piggyback: Adding multiple tasks to a single, fixed network by learning to mask. *arXiv preprint arXiv:1801.06519*, 2018.
- [14] M. Mancini, E. Ricci, B. Caputo, and S. R. Bulò. Adding new tasks to a single network with weight transformations using binary masks. *arXiv preprint arXiv:1805.11119*, 2018.
- [15] M. Mayford, S. A. Siegelbaum, and E. R. Kandel. Synapses and memory storage. *Cold Spring Harbor perspectives in biology*, page a005751, 2012.
- [16] M. McCloskey and N. J. Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. volume 24 of *Psychology of Learning and Motivation*, pages 109 – 165. Academic Press, 1989.
- [17] 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*, volume 2011, page 5, 2011.
- [18] C. V. Nguyen, Y. Li, T. D. Bui, and R. E. Turner. Variational continual learning. *arXiv preprint arXiv:1710.10628*, 2017.
- [19] A. Odena, C. Olah, and J. Shlens. Conditional image synthesis with auxiliary classifier gans. *arXiv preprint arXiv:1610.09585*, 2016.
- [20] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- [21] R. Ratcliff. Connectionist models of recognition memory: Constraints imposed by learning and forgetting functions. *Psychological Review*, pages 285–308, 1990.
- [22] S. Rebuffi, A. Kolesnikov, and C. H. Lampert. icarl: Incremental classifier and representation learning. *CoRR*, abs/1611.07725, 2016.
- [23] S.-A. Rebuffi, H. Bilen, and A. Vedaldi. Efficient parametrization of multi-domain deep neural network. In *CVPR*, 2018.
- [24] A. Rosenfeld and J. K. Tsotsos. Incremental learning through deep adaptation. *TPAMI*, 2018.
- [25] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [26] A. A. Rusu, N. C. Rabinowitz, G. Desjardins, H. Soyer, J. Kirkpatrick, K. Kavukcuoglu, R. Pascanu, and R. Hadsell. Progressive neural networks. *CoRR*, abs/1606.04671, 2016.
- [27] J. Schwarz, J. Luketina, W. M. Czarnecki, A. Grabska-Barwinska, Y. W. Teh, R. Pascanu, and R. Hadsell. Progress & compress: A scalable framework for continual learning. *arXiv preprint arXiv:1805.06370*, 2018.
- [28] A. Seff, A. Beatson, D. Suo, and H. Liu. Continual learning in generative adversarial nets. *arXiv preprint arXiv:1705.08395*, 2017.
- [29] J. Serrà, D. Surís, M. Miron, and A. Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. *CoRR*, abs/1801.01423, 2018.
- [30] H. Shin, J. K. Lee, J. Kim, and J. Kim. Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems*, pages 2990–2999, 2017.
- [31] C. Wu, L. Herranz, X. Liu, Y. Wang, J. van de Weijer, and B. Raducanu. Memory Replay GANs: learning to generate images from new categories without forgetting. In *Advances In Neural Information Processing Systems*, 2018.
- [32] Y. Wu, Y. Chen, L. Wang, Y. Ye, Z. Liu, Y. Guo, Z. Zhang, and Y. Fu. Incremental classifier learning with generative adversarial networks. *CoRR*, abs/1802.00853, 2018.
- [33] J. Yoon, E. Yang, J. Lee, and S. J. Hwang. Lifelong learning with dynamically expandable networks. 2018.
- [34] F. Zenke, B. Poole, and S. Ganguli. Improved multitask learning through synaptic intelligence. *CoRR*, abs/1703.04200, 2017.