

Distillation Techniques for Pseudo-rehearsal Based Incremental Learning

Haseeb Shah, Khurram Javed and Faisal Shafait

School of Electrical Engineering and Computer Science,
National University of Science and Technology,
{hshah.bese15seecs, 14besekjaved, faisal.shafait}@seecs.edu.pk

Abstract. The ability to learn from incrementally arriving data is essential for any life-long learning system. However, standard deep neural networks forget the knowledge about the old tasks, a phenomenon called catastrophic forgetting, when trained on incrementally arriving data. We discuss the biases in current Generative Adversarial Networks (GAN) based approaches that learn the classifier by knowledge distillation from previously trained classifiers. These biases cause the trained classifier to perform poorly. We propose an approach to remove these biases by distilling knowledge from the classifier of AC-GAN. Experiments on MNIST and CIFAR10 show that this method is comparable to current state of the art rehearsal based approaches. The code for this paper is available at this [link](#).

1 Introduction

Standard deep neural networks can solve a variety of simple and complex problems. Most commonly nowadays, these networks are trained on data that is collected beforehand. This is not feasible in some applications, such as systems that require life-long learning. In such cases, the network has to be trained on incrementally arriving data. However, current standard neural networks tend to forget the knowledge they obtained from previous training sessions when trained incrementally. This phenomenon is known as catastrophic forgetting.

One particular way of overcoming forgetting is to retrain the network on old training examples in addition to the new training examples [16]. These examples can be obtained either by storing the exemplars from previous increments, or storing a GAN [3] that generates these examples. We empirically study the latter approach, using Auxiliary Classifier GAN (AC-GAN) [15] in particular.

We find that the samples produced by GANs on simple datasets like MNIST are accurate enough for retaining knowledge. However, we observe that distilling knowledge from the old classifier when incrementally training on new classes results in a biased classifier. We propose a way to remove this bias by instead using AC-GAN to generate images and using the auxiliary classifier of the discriminator for distillation.

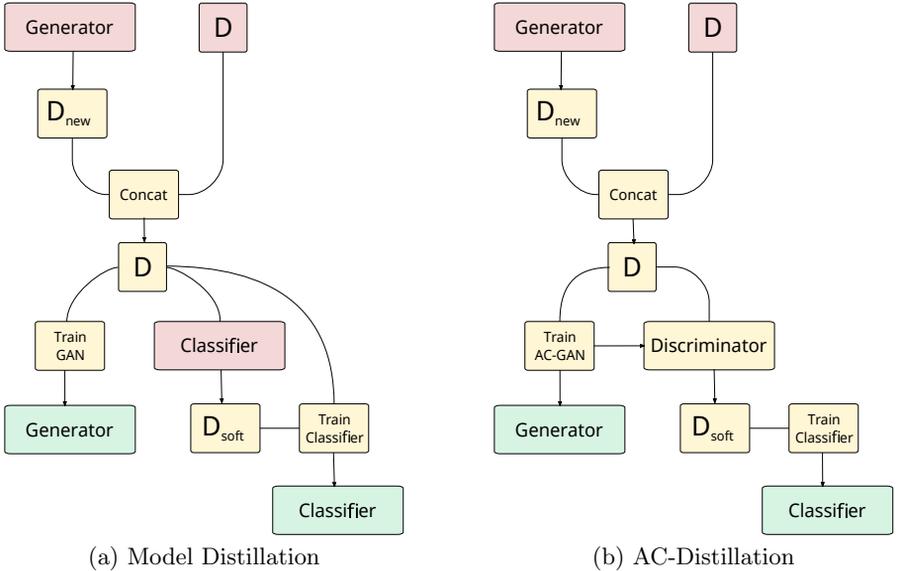


Fig. 1. (a) Architecture for the Model Distillation process as described in Section 4. (b) Architecture for AC Distillation process described in Section 6. Inputs of each learning session are highlighted as red while the outputs are highlighted green. At each increment, these inputs are fed into the algorithm. Firstly, the new dataset \mathcal{D} is augmented with the data generated belonging to old classes. Secondly, we train a new GAN using this dataset \mathcal{D} . In Model Distillation, the old classifier is used to generate soft targets which are then used in combination with real targets to train a new classifier. Whereas in AC Distillation, the discriminator of AC-GAN is used to generate the soft targets which are then used directly to train the classifier.

The rest of this paper is organized as follows. In Section 2, we introduce GANs, knowledge distillation and NCM classifier. In Section 4, we discuss model distillation. In Section 5, we discuss the biases introduced in the classifier when using model distillation. In Section 6, we discuss the use of AC-Distillation to remove these biases. In Section 3, we discuss the related work on this topic and the baselines. In Section 7, we perform empirical analysis of the approach, analyze the bias and discuss limitations. Finally in Section 8, we conclude the paper.

2 Preliminaries

2.1 Generative Adversarial Networks (GANs)

Generative Adversarial Networks, introduced by Goodfellow *et al* [3], is a framework used to mimic a target distribution. GANs consist of a generator, which generates the samples that mimic the target distribution and a discriminator, that gives the probability distribution over sources. Both these networks are

trained in turns, with discriminator acting as a critic of increasing difficulty for the generator.

Conditional GANs (cGANs) are an extension of GANs [14] which additionally utilise conditional information in order to learn the target distribution. These are used for class conditional image generation [14], text to image generation [17] and image manipulation [5]. One particular method of utilising the conditional information, known as Auxiliary Classifier Generative Adversarial Networks (AC-GAN) [15], incorporates a classifier in discriminator. This classifier allows AC-GAN’s discriminator to output the class probabilities in addition to probabilities of the source of input.

2.2 Knowledge Distillation

Knowledge distillation is a technique that can be used to transfer knowledge between neural networks. This transfer is achieved by training the target network using output probabilities of source network as target probabilities. These targets have higher entropy and thus contain similarity information that can be utilised while training. On simple datasets like MNIST, the probabilities of non-target class are often very small. Caruana *et al* [1] minimized the squared distance between logits to deal with this problem. Hinton *et al* [4] propose a more general solution by raising the temperature of softmax to produce softer targets and use them directly.

2.3 Nearest Class Mean (NCM) Classifier

Nearest Class Mean classifier [13] is a classification method where instead of using the trained classifier to perform the classification directly, the class of an image is picked as the one whose embeddings have minimum euclidean distance from the mean embeddings of all images in that class. More formally,

$$c^* = \arg \min_{c=1\dots t} \|\phi(x) - \mu(c)\| \quad (1)$$

Where c is the class, ϕ is the trained classifier and μ is the true class mean calculated using all the examples present in c . Rebuffi *et al* [16] proposed *Nearest Mean of Exemplars* rule where instead of calculating the true class mean, the mean μ is calculated on a smaller set of exemplars acting as a representative of class c .

3 Related Work

In this section, we discuss the work related to our proposed approach in the literature.

McCloskey & Cohen [12] explored the use of rehearsal in order to prevent forgetting. During rehearsal, a random set of training examples from old training sessions are picked and mixed with the ones in new training session. Robins [18] proposed pseudorehearsal, in which forgetting is prevented by generating random examples and then passing them through the network to obtain their labels. These examples along with their labels are then used during rehearsal.

Li and Hoiem [10] introduce a method named Learning without Forgetting (LwF) wherein the distillation loss is used to prevent catastrophic forgetting. This method preserves the privacy as previous class data from previous training sessions is not stored. Rebuffi *et al* [16] introduced iCarl, which involves storing a representative set of exemplars for classes encountered in old training sessions using a process called herding. At each increment, they train the classifier using distillation and the classification losses combined. During classification, they use Nearest Mean of Exemplars rule.

Venkatesan *et al* [21] demonstrated that Generative Adversarial Networks can be used to as a privacy preserving alternative to iCarl. They propose to store a separate generator at each increment, which can consume a lot of storage. Our approach stores only one generator to generate representatives of all the previous increments. Furthermore, they do not deal with the bias introduced during distillation.

Wu *et al* [22] also empirically studied the use of Generative Adversarial Networks for incrementally learning the classifier. They proposed a way to remove the bias caused by imbalance in the number of samples in the dataset by introducing a scalar multiplier whose value is estimated using a validation set. Their approach, however, is still biased as we will discuss in Section 5.

Another direction of research commonly found in literature trying to solve the problem of catastrophic forgetting involves freezing of certain weights in neural network. Xiao *et al* [23] propose to learn a model structured like a tree that expands as it encounters more classes. Rusu *et al* [19] propose progressive neural networks, wherein they expand each layer horizontally. Kirkpatrick *et al* [7] propose Elastic Weight Consolidation (EWC), a regularization scheme to slow down learning on weights that are important to previous tasks. EWC builds a Fisher matrix using training data to determine which weights it should regularize.

Kemker *et al* [6] propose FearNet, a brain inspired model for incremental learning which does not require storing old examples. It consists of a dual memory system, where one is for short-term memories and the other for long-term memory. However, instead of learning the representation, they use pretrained ResNet embeddings to obtain features to be fed into their model.

Lopez-Paz & Ranzato propose Gradient Episodic Memory [11] where they solve forgetting problem by storing data from previous training sessions and performing gradient update such that the error on previously learned tasks do not increase. PathNet, proposed by Fernando *et al* [2] uses a genetic algorithm to find the optimal path through the network for each training session. For each training

Algorithm 1: Incremental Knowledge Transfer Using Model Distillation

Input: Dataset \mathcal{D} , Classifier ϕ , Generator \mathcal{G} and Temperature T
Output: Classifier ϕ and Generator \mathcal{G}

- 1 **if** *old Classifier ϕ exists* **then**
- 2 | Use generator \mathcal{G} to generate (x_i, y_i) and append them into \mathcal{D}
- 3 **end**
- 4 Train a new generator \mathcal{G} and discriminator ψ on $(x_i, y_i) \sim \mathcal{D}$
- 5 **if** *old Classifier ϕ exists* **then**
- 6 | $\mathcal{D}_{soft} = (x, y^{\phi(T)}) \forall x \in \mathcal{D}$
- 7 | Train ϕ by minimizing
 $\mathbb{E}_{(x,y) \sim \mathcal{D}, (\cdot, y') \sim \mathcal{D}_{soft}} [\alpha * \mathcal{H}(y, \phi(x)) + (1 - \alpha) * \mathcal{H}(y', \phi(x))]$
- 8 **else**
- 9 | Train ϕ by minimizing $\mathbb{E}_{(x,y) \sim \mathcal{D}} [\mathcal{H}(y, \phi(x))]$
- 10 **end**

session, only optimal path through the network will be trainable, preventing any new training sessions from causing the model to lose knowledge about previous training data.

4 Model Distillation

In this section we formally describe and discuss the approach for incremental learning by distilling knowledge from an old classifier as is common in the literature [21] [22]. This approach is applicable to any GAN, i.e. we are not limited to GANs which have auxiliary classifier in them. We assume that we have a source from which we receive real datasets i.e. $\mathcal{S} = \{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_i\}$ Each dataset represents one increment of learning and contains tuples of training samples x_i and their true labels y_i i.e. $\mathcal{D}_i = \{(x_i, y_i)\}$. We only have a limited amount of observations in datasets obtained from \mathcal{S} . In contrast, we can generate a very large amount of training data from the generator at almost no cost.

The proposed setup consists of a classifier ϕ to which the knowledge has to be transferred, a generator \mathcal{G} able to generate training samples and a discriminator ψ . Throughout the paper, we will use the notation $y_i^{M(t)}$ to denote the label of x_i generated using model M with temperature t , y_i to denote the true label of x_i and \mathcal{H} to denote the Cross Entropy. Whenever we receive a new real dataset from \mathcal{S} , it is fed along with other prerequisites to Algorithm 1, which consists of three phases as discussed below:

Generate and add samples to dataset using the latest generator (Steps 1 - 3)

The main goal of this step is to generate samples belonging to each class present in the previous datasets. Therefore, this step is skipped if there are no previous datasets i.e. the classifier has not been trained yet. We generate a fixed amount k samples for each class. In case the generator \mathcal{G} is not conditional, we generate

$k * m$ samples where m is the number of classes the classifier ϕ has been trained on. These samples are then appended into dataset \mathcal{D} . We discard \mathcal{G} after this step since there is no further use of it.

Train generator \mathcal{G} and discriminator ψ on samples obtained from \mathcal{D} (Step 4)

Several different types of GANs can be used in this phase. The goal at this phase is to train \mathcal{G} to be used as a privacy preserving alternative to storing \mathcal{D} at the cost of reduced sample quality. This approach keeps at most one \mathcal{G} and ψ stored at a time, and they are both reinitialized at every increment instead of being trained continuously because several variants of GANs suffer from mode collapse when trained for a longer periods of time.

Transfer previous knowledge to classifier ϕ using Model-Distillation (Steps 5 - 9)

If there is no old classifier to transfer the knowledge from, then we simply minimise the cross-entropy between the real and the predicted targets. Otherwise, we generate a dataset \mathcal{D}_{soft} with softer targets $y^{\phi(T)}$ using temperature T . The classifier ϕ from which we are obtaining the soft targets at this point is trained only on classes belonging to older datasets. In the next step, we proceed by training the classifier ϕ on new classes by using standard weighted distillation loss. In first part of the loss, we calculate the cross entropy between the predicted target $\phi(x)$ and real target y obtained from \mathcal{D} . In second part, we calculate cross entropy between the predicted target $\phi(x)$ and soft target $y^{\phi(T)}$ obtained from \mathcal{D}_{soft} . α is the weight constant used to determine the weight of these two terms. It is necessary to train on real targets in addition to softened targets as the classifier used to generate the soft targets had no knowledge about the new classes, so omitting it would result in inability to learn about them.

5 Model Distillation Trains a Biased Classifier

When we perform distillation using trained classifier ϕ , the resulting model can be biased in three different ways:

1. The number of samples generated by \mathcal{G} can be different than the number of samples present per unique class in \mathcal{D} i.e $k * m \gg |\mathcal{D}|$ or $k * m \ll |\mathcal{D}|$ where k is the number of generated samples per class and m is the number of unique classes in \mathcal{D} .
2. When distillation is performed on new classes the soft targets obtained will indicate that new classes have high similarity to older classes, even if they are not similar. This introduces a bias against the new classes.
3. When distillation is performed on old classes, the soft targets will indicate that the old classes have no similarity to new classes, even if they are similar. This introduces a bias against the old classes.

We illustrate these biases through a simple example. We assume that we have an ideal softmax based image classifier ϕ_{ideal} that can classify the samples it was

Algorithm 2: Incremental Knowledge Transfer Using AC-Distillation

Input: Dataset \mathcal{D} , Classifier ϕ , Generator \mathcal{G} and Temperature T

Output: Classifier ϕ and Generator \mathcal{G}

```

1 if old Classifier  $\phi$  exists then
2   | Use generator  $\mathcal{G}$  to generate  $(x_i, y_i)$  and append them into  $\mathcal{D}$ 
3 end
4 Train a new generator  $\mathcal{G}$  and discriminator  $\psi$  on  $(x_i, y_i) \sim \mathcal{D}$ 
5 if old Classifier  $\phi$  exists then
6   |  $\mathcal{D}_{soft} = (x, y^{\psi_{AC}(T)}) \forall x \in \mathcal{D}$ 
7   | Train  $\phi$  by minimizing  $\mathbb{E}_{(x,y) \sim \mathcal{D}_{soft}} [\mathcal{H}(y, \phi(x))]$ 
8 else
9   | Train  $\phi$  by minimizing  $\mathbb{E}_{(x,y) \sim \mathcal{D}} [\mathcal{H}(y, \phi(x))]$ 
10 end

```

trained on with zero error. For the second type of bias, we assume that ϕ_{ideal} is already trained on classes $\mathcal{D} = (\text{Buildings}, \text{Boats})$. If this ϕ_{ideal} has to be trained on $\mathcal{D} = (\text{Cats})$, the second term of equation in step 7 will of Algorithm 1 will indicate that the Cat class images have a high similarity to the images of Buildings and Boats, which is clearly not the case. This introduces a bias against the Cat class as their class probabilities are incorrect. This happens because the soft labels y' are obtained from a classifier that was trained only on the old classes i.e. $D = (\text{Buildings}, \text{Boats})$.

For the third type of bias, we assume that ϕ_{ideal} is already trained on classes $\mathcal{D} = (\text{Dogs}, \text{Ocean})$. If this ϕ_{ideal} has to be trained on $\mathcal{D} = (\text{Cats})$, the second term of equation in step 7 will of Algorithm 1 will indicate that the Dog images have zero similarity to the images of Cats. This introduces a bias against the Dog class as its class probabilities are incorrect. Which is also caused because the soft labels y' are obtained from a classifier that was trained only on the old classes.

The effect of these biases can be slightly mitigated by carefully selecting the parameters k and α , however, we present an alternative solution to this problem by using an auxiliary classifier to mitigate this bias.

6 AC-Distillation

The AC-Distillation (Auxiliary Classifier Distillation) is outlined in algorithm 2. The key idea of AC-Distillation is to generate \mathcal{D}_{soft} by distilling knowledge from auxiliary classifier ψ_{AC} of GAN instead of the old classifier ϕ . The details of differences of this approach are described below:

1. In step 6, the distillation is done using the auxiliary classifier ψ_{AC} obtained from GAN that we just trained. Therefore, ψ_{AC} contains information about the new classes present in \mathcal{D} in addition to old classes.

2. In step 7, we train the classifier ϕ on \mathcal{D}_{soft} that we generated in step 6. Note that we calculate cross-entropy only using the soft targets present in \mathcal{D}_{soft} and do not use the real targets present in \mathcal{D} . This is because in our experiments, when using a source of distillation that contains information about all the classes that the classifier ϕ is to be trained on, using only \mathcal{D}_{soft} performed at least as good as using $\mathcal{D}_{soft} \cup \mathcal{D}$. Furthermore, this allows us to get rid of the parameter α , which reduces our hyperparameter search space.

Note that in this method, we are limited to GANs that have an auxiliary classifier built into them. One example of a GAN that satisfies this property is AC-GAN [15].

One might doubt about the usefulness of separately training a new classifier when there already exists a classifier in the GAN. However, this constraints the size and type of architectures that can be used as a classifier, as the auxiliary classifier is a part of the GAN, which has to be built in a certain way.

7 Experiments

7.1 Benchmark

We use the benchmark protocol suggested by [16], that is, we train using each method in a class incremental way, where the classifier ϕ receives a new set of previously unseen classes each increment. After each increment, the classifier ϕ is evaluated on the test set consisting of new and old classes combined. We use MNIST [9] and CIFAR10 [8] as real sources \mathcal{S} . Each source \mathcal{S} contains 5 datasets \mathcal{D} where each dataset \mathcal{D} consists of 2 randomly picked previously unseen classes.

7.2 Implementation details

Classifier: We train Resnet-32 [20] with batch size of 100. Learning rate is set as 2 and decayed by 0.2. Sigmoid is used in the final layer and binary cross entropy is used as loss function. For MNIST, we train for 15 epochs and decay the learning rate after 8 and 12 epochs. For CIFAR10, we train for 70 epochs and decay the learning rate after 45, 60 and 68 epochs.

AC-GAN: We use the CIFAR10 architecture provided by Odena *et al* [15] with a few minor changes in their hyperparameters. More specifically, we train MNIST for 20 epochs each increment, decaying the learning rate at epochs 11 and 16 and CIFAR10 for 1000 epochs each increment, decaying the learning rate after 900 epochs. Learning rate is decayed by 0.1.

7.3 Approaches compared

We compare the following techniques in our experiments:

1. iCarl [16]: Distillation using old classifier ϕ on herded and new exemplars. Classification is done using the Mean of Exemplars rule. This approach is not privacy-preserving as old class data is stored.

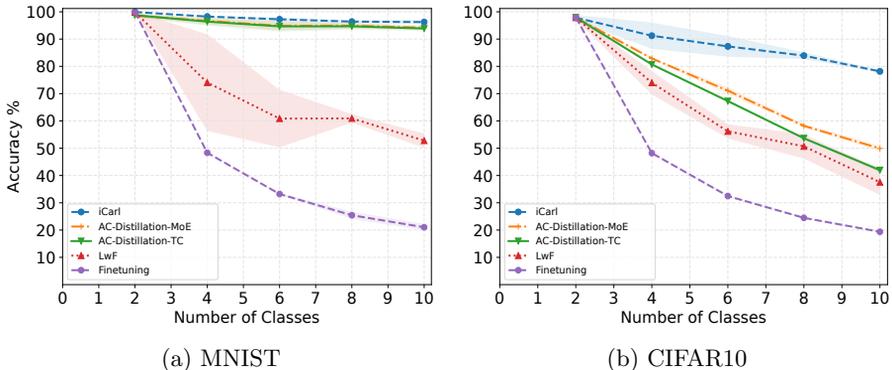


Fig. 2. All experiments were repeated 3 times. Mean and standard deviation are reported. (a) Results on MNIST with 2 classes per dataset \mathcal{D} . Here the performance AC distillation approaches is comparable to iCarl. (b) Results on CIFAR10 with 2 classes per dataset \mathcal{D} . Here the performance of AC distillation approaches is worse than iCarl. This is because the AC-GAN is unable to properly learn the dataset. In contrast, iCarl is storing the real images as exemplars. In both datasets, AC distillation has the best performance among the privacy preserving approaches.

2. LwF [10]: Distillation using old classifier ϕ only on new examples. Classification is done using newly trained classifier. This is a privacy-preserving approach.
3. Finetuning: No distillation and performing classification using the newly trained classifier. Here we simply finetune the old classifier on newly arriving classes without taking any measures to prevent catastrophic forgetting.
4. AC-Distillation-TC: Distillation using auxiliary classifier ψ_{AC} of GAN on generated and new examples. Classification is done using the newly trained classifier. This is a privacy-preserving approach.
5. AC-Distillation-MoE: Distillation using auxiliary classifier ψ_{AC} of GAN on generated and new examples. Classification is done using Mean of Exemplars classification rule. This is a privacy-preserving approach.

7.4 Results

MNIST: We can see in Fig. 2a that the performance of AC-Distillation-TC is resembles AC-Distillation-MoE. This shows that on simple datasets, the trained classifier can be used directly to obtain a good classification accuracy. When comparing to iCarl, the performance is slightly lower, which is due to the limitation of AC-GAN to perfectly model the dataset. However, if only the privacy-preserving approaches are compared, AC-Distillation approaches beat other approaches by a large margin. The finetuning approach forgets the existence of earlier classes.

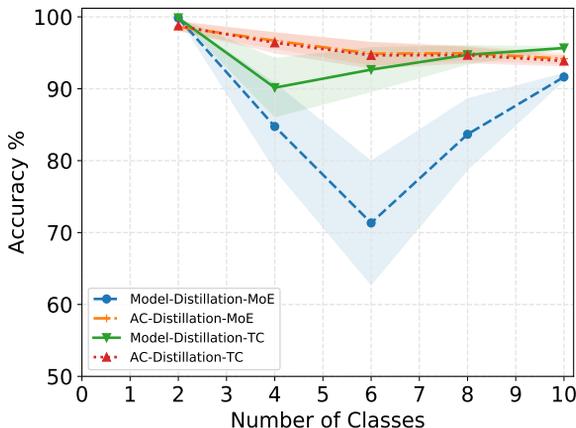


Fig. 3. Comparison of AC-Distillation methods with Model distillation on MNIST using 2 classes per dataset \mathcal{D} . All experiments were repeated 3 times. Mean and standard deviation are reported. Y-axis is truncated to highlight the differences. We can see that the AC distillation approaches have lower variance and perform significantly better than the Model distillation approaches.

CIFAR10: CIFAR10 is a much more complicated dataset to model as compared to MNIST. This is partly because the size of the image is much smaller when compared to contents of the image. We can see from Fig. 2b that the performance of AC-Distillation methods is significantly worse than iCarL, which is due to the inability of the AC-GAN to properly learn the dataset. However, even with the AC-GAN performing poorly on this dataset, AC-distillation methods perform better than the other privacy-preserving approaches.

If we compare performance of the two AC-distillation methods, we see that AC-Distillation-TC performs slightly worse than AC-Distillation-MoE. This shows that on more complicated datasets, Mean of Exemplars classification rule should be used to achieve the best performance as it slightly compensates for the inability of GANs properly model the dataset.

The results on MNIST indicate that AC-Distillation is a viable approach to privacy-preserving incremental learning. However, when using more complicated datasets, AC-GAN is unable to properly model the dataset. This results in a poor performance when using AC-distillation on such datasets. In order to achieve a performance comparable to non privacy-preserving approaches on more difficult datasets, AC-GAN has to be replaced with more advanced GAN architectures.

7.5 Comparison of AC-Distillation with Model Distillation

In this section, we compare Auxiliary Classifier based knowledge distillation methods with Model distillation methods. More specifically, we compare AC-

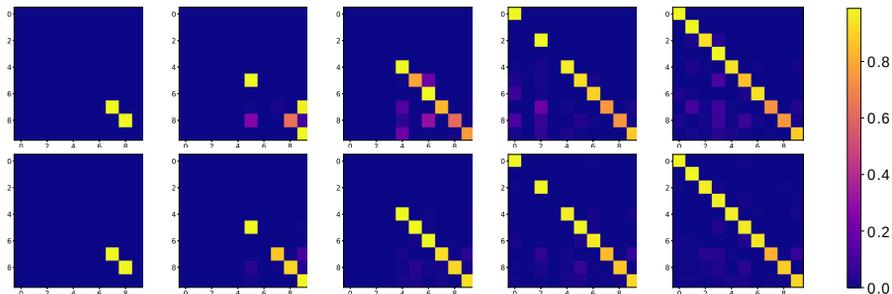


Fig. 4. Confusion matrices for incremental training session on MNIST where x-axis denotes the ground truth and y-axis denotes the predicted label. Top: Confusion matrices for Model-Distillation-MoE. Bottom: Confusion matrices for AC-Distillation-MoE. These matrices are for a single run of each experiment. Each matrix indicates a separate training session with new knowledge added to it as we go towards the right side. The new classes in the above example arrived in the this order: $\mathcal{S} = \{(7, 8), (5, 9), (4, 6), (0, 2), (1, 3)\}$. We can observe from these figures that the Model distillation approach tends to forget the knowledge about the older classes in the middle increments while the AC-distillation approach is able to retain the knowledge about old classes.

Distillation-TC and *AC-Distillation-MoE* described in Section 7.3 with the following:

1. Model-Distillation-TC: Distillation using old classifier ϕ on generated and new examples. Classification is done using the newly trained classifier.
2. Model-Distillation-MoE: Distillation using old classifier ϕ on generated and new examples. Classification is done using the Mean of Exemplars classification rule.

To make the comparison fair, all the experiments were run using AC-GAN and the parameters described in Section 7.2. The results of the comparison are shown in Fig. 3. At each increment of training, the classifier receives two new classes to train on. It can be seen that the Model distillation methods have higher variance and perform worse as compared to Auxiliary Classifier based distillation methods. The decreased performance of Model distillation methods is caused due to the biases discussed in Section 5.

To understand their behavior, we can see the comparison of confusion matrices for Auxiliary Classifier and Model distillation based Mean of Exemplar classification methods in Fig. 4. It appears that when using AC-Distillation, the Mean of Exemplars classification rule is able to classify both the new and old classes properly. However, in the case of Model distillation the Mean of Exemplars rule is unable to retain the knowledge about the old classes properly, and incorrectly classifies them as members of the new classes. As more and more new classes are fed to the classifier to learn about, the performance of Model

distillation methods becomes comparable with AC-distillation methods.

From the above analysis, we can observe that when we have a small number of classes to incrementally train on, the effect of bias is significant in Model distillation methods. This leads to poor performance of Model distillation methods when compared to AC-distillation methods. However, when the number of classes is increased, the effect of bias is mitigated to some extent and performance of the classifier becomes comparable with AC-distillation methods. In contrast to this, AC-distillation methods do not suffer from this bias at any number of classes.

8 Conclusion

We have shown that performing knowledge distillation from the auxiliary classifier of AC-GAN corrects the biases that are prevalent in GAN based incremental learning methods. We proposed the AC-distillation approach to correct the biases. On MNIST, we studied the biases in model distillation approaches and demonstrated the effectiveness of the AC-distillation approach. On more complicated datasets, the performance is significantly worse due to the inability of AC-GAN to properly model the dataset. However, if a more advanced GAN is used that can properly model the dataset, the performance of the proposed approach is comparable to state of the art rehearsal based incremental learning methods, while being privacy preserving.

References

1. Bucilua, C., Caruana, R., Niculescu-Mizil, A.: Model compression. In: Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining. pp. 535–541. ACM (2006)
2. Fernando, C., Banarse, D., Blundell, C., Zwols, Y., Ha, D., Rusu, A.A., Pritzel, A., Wierstra, D.: Pathnet: Evolution channels gradient descent in super neural networks. arXiv preprint [arXiv:1701.08734](https://arxiv.org/abs/1701.08734) (2017)
3. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in neural information processing systems. pp. 2672–2680 (2014)
4. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint [arXiv:1503.02531](https://arxiv.org/abs/1503.02531) (2015)
5. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. CVPR (2017)
6. Kemker, R., Kanan, C.: Fearnnet: Brain-inspired model for incremental learning. ICLR (2018)
7. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al.: Overcoming catastrophic forgetting in neural networks. Proceedings of the National Academy of Sciences 114(13), 3521–3526 (2017)
8. Krizhevsky, A., Nair, V., Hinton, G.: The cifar-10 dataset. online: <http://www.cs.toronto.edu/kriz/cifar.html> (2014)

9. LeCun, Y., Cortes, C., Burges, C.: Mnist handwritten digit database. AT&T Labs [Online]. Available: <http://yann.lecun.com/exdb/mnist> 2 (2010)
10. Li, Z., Hoiem, D.: Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2017)
11. Lopez-Paz, D., et al.: Gradient episodic memory for continual learning. In: *Advances in Neural Information Processing Systems*. pp. 6470–6479 (2017)
12. McCloskey, M., Cohen, N.J.: Catastrophic interference in connectionist networks: The sequential learning problem. In: *Psychology of learning and motivation*, vol. 24, pp. 109–165. Elsevier (1989)
13. Mensink, T., Verbeek, J., Perronnin, F., Csurka, G.: Distance-based image classification: Generalizing to new classes at near-zero cost. *IEEE transactions on pattern analysis and machine intelligence* 35(11), 2624–2637 (2013)
14. Mirza, M., Osindero, S.: Conditional generative adversarial nets. *arXiv preprint [arXiv:1411.1784](https://arxiv.org/abs/1411.1784)* (2014)
15. Odena, A., Olah, C., Shlens, J.: Conditional image synthesis with auxiliary classifier GANs. In: Precup, D., Teh, Y.W. (eds.) *Proceedings of the 34th International Conference on Machine Learning*. *Proceedings of Machine Learning Research*, vol. 70, pp. 2642–2651. PMLR, International Convention Centre, Sydney, Australia (06–11 Aug 2017)
16. Rebuffi, S.A., Kolesnikov, A., Lampert, C.H.: icarl: Incremental classifier and representation learning. In: *Proc. CVPR* (2017)
17. Reed, S., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., Lee, H.: Generative adversarial text to image synthesis. *arXiv preprint [arXiv:1605.05396](https://arxiv.org/abs/1605.05396)* (2016)
18. Robins, A.: Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science* 7(2), 123–146 (1995)
19. Rusu, A.A., Rabinowitz, N.C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., Hadsell, R.: Progressive neural networks. *arXiv preprint [arXiv:1606.04671](https://arxiv.org/abs/1606.04671)* (2016)
20. Szegedy, C., Ioffe, S., Vanhoucke, V., Alemi, A.A.: Inception-v4, inception-resnet and the impact of residual connections on learning. In: *AAAI*. vol. 4, p. 12 (2017)
21. Venkatesan, R., Venkateswara, H., Panchanathan, S., Li, B.: A strategy for an uncompromising incremental learner. *arXiv preprint [arXiv:1705.00744](https://arxiv.org/abs/1705.00744)* (2017)
22. Wu, Y., Chen, Y., Wang, L., Ye, Y., Liu, Z., Guo, Y., Zhang, Z., Fu, Y.: Incremental classifier learning with generative adversarial networks. *arXiv preprint [arXiv:1802.00853](https://arxiv.org/abs/1802.00853)* (2018)
23. Xiao, T., Zhang, J., Yang, K., Peng, Y., Zhang, Z.: Error-driven incremental learning in deep convolutional neural network for large-scale image classification. In: *Proceedings of the 22nd ACM international conference on Multimedia*. pp. 177–186. ACM (2014)