

Marginal Replay vs Conditional Replay for Continual Learning

Timothée Lesort^{1,2}[0000-0002-8669-0764], Alexander Gepperth³[0000-0003-2216-7808],
Andrei Stoian²[0000-0002-3479-9565], and David Filliat¹[0000-0002-5739-1618]

¹ Flowers Laboratory (ENSTA ParisTech & INRIA), France

² Thales, Theresis Laboratory, France

³ Fulda University of Applied Sciences, Germany

Abstract. We present a new replay-based method of continual classification learning that we term "conditional replay" which generates samples and labels together by sampling from a distribution conditioned on the class. We compare conditional replay to another replay-based continual learning paradigm (which we term "marginal replay") that generates samples independently of their class and assigns labels in a separate step. The main improvement in conditional replay is that labels for generated samples need not be inferred, which reduces the margin for error in complex continual classification learning tasks. We demonstrate the effectiveness of this approach using novel and standard benchmarks constructed from MNIST and FashionMNIST data, and compare to the regularization-based *elastic weight consolidation* (EWC) method [18, 38].

Keywords: Continual Learning · Generative Models · Generative Replay.

1 Introduction

This contribution is in the context of incremental, continual or lifelong learning, subject that is gaining increasing recent attention [28, 7] and for which a variety of different solutions have recently been proposed (see below). Briefly put, the problem consists of repeatedly re-training a deep neural network (DNN) model with new sub-tasks, or continual learning tasks (CLTs), (for example: new visual classes) over long time periods, while avoiding the abrupt degradation of previously learned abilities that is known under the term "catastrophic interference" or "catastrophic forgetting" [30, 5, 7]. Please see Fig. 1 for a visualization of the problem setting. It has long been known that catastrophic forgetting (CF) is a problem for connectionist models [5] of which modern DNNs are a specialized instance, but only recently there have been efforts to propose workable solutions to this problem for deep learning models [20, 18, 31, 16, 35]. A recent article [29] demonstrates empirically that most proposals fail to eliminate CF when common-sense application constraints are imposed (e.g., restricting prior access to data from new sub-tasks, or imposing constant, low memory and execution time requirements).

One aspect of the problem seems to be that gradient-based DNN training is greedy, i.e., it tries to optimize all weights in the network to solve the current task only. Previous tasks, which are not represented in the current training data, will naturally be disregarded

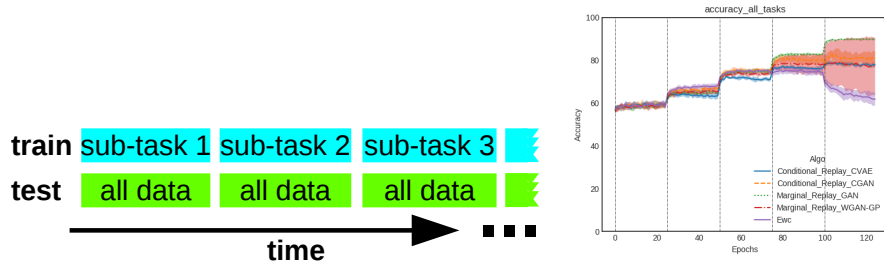


Fig. 1: Left: The problem setting of continual learning as investigated in this article. DNN models are trained one after the other on a sequence of sub-tasks (of which three are shown here), and are continuously evaluated on a test set consisting of the union of all sub-task test sets. This gives rise to results as shown exemplarily on the right-hand side of the figure, i.e., plots of test set accuracy over time for different models, where boundaries between sub-tasks (5 in this case) are indicated by vertical lines.

in this process. While approaches such as [18, 20] aim at "protecting" weights that were important for previous tasks, one can approach the problem from the other end and simply include samples from previous tasks in the training process each time a new task is introduced.

This is the *generative replay* approach, which is in principle model-agnostic, as it can be performed with a variety of machine learning models such as decision trees, support vector machines (SVMs) or deep neural networks (DNNs). It is however unfeasible for, e.g., embodied agents or embedded devices performing object recognition, to store all samples from all previous sub-tasks. Because of this, generative replay proposes to train an additional machine learning model (the so-called *generator*). Thus, the "essence" of previous tasks comes in the form of trained generator parameters which usually require far less space than the samples themselves. A downside of this and similar approaches is that the time complexity of adapting to a new task is not constant but depends on the number of preceding tasks that should be replayed. Or, conversely, if continual learning should be performed at constant time complexity, only a fixed amount of samples can be generated, and thus there will be forgetting, although it won't be catastrophic.

In this paper we decide to investigate two different types of generative models : Generative adversarial networks (GAN) and variational auto-encoder (VAE). On one hand GAN are known to generate samples of high quality but on the other hand VAE directly maximize the likelihood of the learned distribution while training. It was then interesting to experiment both of them to compare their performances.

This article proposes and evaluates a particular method for performing replay using DNNs, termed "conditional replay", which is similar in spirit to [38] but presents important conceptual improvements (see next section). The main advantage of conditional replay is that samples can be generated conditionally, i.e., based on a provided label. Thus, labels for generated samples need not be inferred in a separate step as other replay-based approaches, e.g., [38], which we term *marginal replay* approaches. Since inferring the label of a generated sample inevitably requires the application of a

possibly less-than-perfect classifier, avoiding this step conceivably reduces the margin for error in complex continual learning tasks. The paper organization is the following, first we introduce the paper contributions and the related works, secondly we describe the methods used as well as the benchmarks, thirdly we present the paper experiments, fourthly we show and discuss our results and in a last section we conclude the paper.

1.1 Contribution

The original contributions of this article can be summarized as follows:

- **Conditional replay as a method for continual classification learning** We experimentally establish the advantages of conditional replay in the field of continual learning by comparing conditional and marginal replay models on a common set of benchmarks.
- **Improvement of marginal learning** We furthermore propose an improvement of marginal replay as proposed in [38] by using generative adversarial networks (GANs, see [9]).
- **New experimental benchmarks for generative replay strategies** To measure the merit of these proposals, we use two experimental settings that have not been previously considered for benchmarking generative replay: rotations and permutations. In addition, we promote the "10-class-disjoint" task as an important benchmark for continual learning as it is impossible to solve for purely discriminative methods (at no time, samples from different classes are provided for training so no discrimination can happen).
- **Comparison of generative replay to EWC** We show the principled advantage that generative replay techniques have with respect to regularization methods like EWC in a "one class per task" setting, which is after all a very common setting in practice and in which discriminatively trained models strongly tend to assign the same class label to every sample regardless of content.

Table 1: Hyperparameters for MNIST and Fashion MNIST all models (all CL settings have the same training hyper parameters with Adam)

Method	Epochs	LR Classifier	LR Generator	beta1	beta2	Batch Size
Marginal Replay	25	0.01	2e-4	5e-1	0.999	64
Conditional Replay	25	0.01	2e-4	5e-1	0.999	64
EWC	25	0.01	-	5e-1	0.999	64

1.2 Related work

The field of continual learning is growing and has been recently reviewed in, e.g., [28, 7]. In the context of neural networks, principal recent approaches include ensemble methods [33, 4, 26, 34, 45, 2, 36, 24], regularization approaches [18, 20, 31, 16, 41, 12, 1, 25,




Fig. 2: MNIST training data for rotation sub-tasks.




Fig. 3: MNIST training data for permutation-type CLTs.

3, 6], dual-memory systems [15, 32, 8], distillation-based approaches [39, 24, 17] and generative replay methods [38, 15, 22, 14, 42]. In the context of single-memory DNN methods, regularization approaches are predominant: whereas it was proposed in [11] that the popular Dropout regularization can alleviate catastrophic forgetting, the EWC method [18] proposes to add a term to the DNN energy function that protects weights that are deemed to be important for the previous sub-task(s). Whether a weight is important or not is determined by approximating and analyzing the Fisher information matrix of the DNN. A somewhat related approach is pursued with the incremental moment matching (IMM, see [20]) technique, where weights are transferred between DNNs trained on successive sub-tasks by regularization techniques, and the Fisher information matrix is used to "merge" weights for current and past sub-tasks. Other regularization-oriented approaches are proposed in [31, 41] which focus on enforcing sparsity of neural activities by lateral interactions within a layer, or in [16]. Concerning recent advances in generative replay improving upon [38]: Several works propose the use of generative models in continual learning of classification tasks [13, 43, 42, 37] but their results does not provide comparison between different types of generative models. [21] propose a conditional replay mechanism similar to the one investigated here, but their goal is the sequential learning of data generation and not classification tasks. Generally, each approach to continual learning has its advantages and disadvantages:

- ensemble methods suffer from little to no interference between present and past knowledge as usually different networks or sub-networks are allocated to different learning tasks. The problem with this approach is that, on the one hand, model complexity is not constant, and more seriously, that the task from which a sample is coming from must be known at inference time in order to select the appropriate (sub-)network.
- regularization approaches are very diverse: in general, their advantage is simplicity and (often) a constant-time/memory behavior w.r.t. the number of tasks. However, the impact of the regularizer on continual learning performance is difficult to understand,

- and several parameters need to be tuned whose significance is unclear (i.e., the strengths of the regularization terms)
- distillation approaches can achieve very good robustness and continual learning performance, but either require the retention of past samples, or the occurrence of samples from past classes in current training data to be consistent. Also, the strength of the various distillation loss regularizers needs to be tuned, usually by cross-validation.
 - generative replay and dual-memory systems show very good and robust continual learning performance, although time complexity of learning depends on the number of previous tasks for current generative replay methods. In addition, the storage of weights for a sufficiently powerful generator may prove very memory-consuming, so this approach cannot be used in all settings.

2 Methods

A basic notion in this article is that of a continual (or sequential) learning task (CLT or SLT, although we will use the abbreviation CLT in this article), denoting a classification problem that is composed of two or more sub-tasks which are presented sequentially to the model in question. Here, the CLTs are constructed from two standard visual classification benchmarks: MNIST and Fashion MNIST, either by dividing available classes into several sub-tasks, or by performing per-sample image processing operations that are identical within, and different between, sub-tasks. All continual learning models are then trained and evaluated in an identical fashion on all CLTs, and performances are compared by a simple visual inspection of classification accuracy plots.

2.1 Benchmarks

MNIST [19] is a common benchmark for computer vision systems and classification problems. It consists of gray scale 28x28 images of handwritten digits (ten balanced classes representing the digits 0-9). The train, test and validation sets contain 55.000, 10.000 and 5.000 samples, respectively.

Fashion MNIST [44] consists of grayscale 28x28 images of clothes. We choose this dataset because it claims to be a “more challenging classification task than the simple MNIST digits data [44]” while having the same data dimensions, number of classes, balancing properties and number of samples in train, test and validation sets.

2.2 Continual learning tasks (CLTs)

All CLTs are constructed from the underlying MNIST and FashionMNIST benchmarks, so the number of samples in train and test sets for each sub-task depend on the precise way of constructing them, see below.

Rotations New sub-tasks are generated by choosing a random rotation angle $\beta \in [0, \pi/2]$ and then performing a 2D in-plane rotation on all samples of the original benchmark. As both benchmarks we use contain samples of 28x28 pixels, no information loss is




Fig. 4: Samples of MNIST training data for the disjoint CLT. Each sub-task adds one more visual class, a principle which carries over identically to FashionMNIST.

introduced by this procedure. We limit rotation angles to $\pi/2$ because larger rotations could mix MNIST classes like 6 and 9. Each sub-task in rotation-based CLTs contains all 10 classes of the underlying benchmark, leading to 55.000 and 10.000 samples, respectively, in the train and test sets of each sub-task.

Permutations New sub-tasks are generated by defining a random pixel permutation scheme, and then applying it to each data sample of the original benchmark. Each sub-task in permutation-based CLTs contains all 10 classes of the underlying benchmark, leading to 55.000 and 10.000 samples, respectively, in the train and test sets of each sub-task.

Disjoint classes For each benchmark, this CLT has as many sub-tasks as there are classes in the benchmark (10 in this article). Each sub-task contains the samples of a single class, i.e., roughly 6.000 samples in the train set and 1.000 samples in the test set. As the classes are balanced for both benchmarks, this does not unduly favor certain classes. This CLT presents a substantial challenge for machine learning methods since a normal DNN would, for each sub-task, learn to map all samples to a single class label irrespective of content. Selective discrimination between any two classes is hard to obtain except if replay is involved, because then a classifier actually "sees" samples from different classes at the same time.

2.3 Models

In this article, we compare a considerable number of deep learning models: unless otherwise stated, we employ the Rectified Linear Unit (ReLU) transfer function, cross-entropy loss for classifier training, and the Adam optimizer.

EWC We re-implemented the algorithm described in [18], choosing two hidden layers with 200 neurons each.


Marginal replay In the context of classification, the *marginal replay* [21, 38, 42] method works as follows : For each sub-task t , there is a dataset D_t , a classifier C_t , a generator G_t and a memory of past samples composed of a generator G_{t-1} and a classifier C_{t-1} . The latter two allow the generation of artificial samples D_{t-1} from previous sub-tasks. Then, by training C_t and G_t on D_t and D_{t-1} , the model can learn the new sub-task t without forgetting old ones. At the end of the sub-task, C_t and G_t are frozen and replace C_{t-1} and G_{t-1} . In the default setting, we use the generator for marginal replay in a way that ensures a balanced distribution of classes from past sub-tasks D_{t-1} , see also Fig. 7. This is achieved by choosing a predetermined number of samples N to be added for all sub-tasks t , and letting the generator produce tN previous samples at sub-task t . Thus, the number of generated samples increases linearly over time. We choose to evaluate two different models for the generator: WGAN-GP as used in [38] and the original GAN model [10] since it is a competitive baseline [23].

Conditional replay The conditional replay method is derived from *marginal replay*: instead of saving a classifier and a generator, the algorithm only saves a generator that can generate conditionally (for a certain class). Hence, for each sub-task t , there is a dataset D_t , a classifier C_t and two generators G_t and G_{t-1} . The goal of G_{t-1} is to generate data from all the previous sub-tasks during training on the new sub-task. Since data is generated conditionally, samples automatically have a label and do not require a frozen classifier. We follow the same strategy as for marginal replay (previous paragraph) for choosing the number of generated samples at each sub-task. However, conditional replay does not require this: it can, in principle, keep the number of generated samples constant for each sub-task since it is trivially possible to generate a balanced distribution of $\frac{N}{t}$ samples per class, from t different classes, via conditional sample generation. C_t and G_t learn from generated data D_{t-1} and D_t . At the end of a sub-task t , C_t is able to classify data from the current and previous sub-tasks, and G_t is able to sample from them also. We choose to use two different popular conditional models : CGAN described in [27] and CVAE [40].


3 Experiments

We conduct experiments using all models and CLTs described in the previous section. Each class (regardless of the CLT) is presented for 25 epochs, Results are presented either based on the time-varying classification accuracy on the *whole* test set, or on the class (from the test set) that was presented first. In the first case, accuracy should ideally increase over time and reach its maximum after the last class has been presented. In the second case, accuracy should be stable if the model does not forget or decrease over time, reflecting that some information about the first class is forgotten. We distinguish two major experimental goals or questions:


- Establishing the performance of the newly proposed methods (marginal replay with GAN, conditional replay with CGAN or CVAE) w.r.t. the state of the art. To this effect, we conduct experiments that increase the number of generated samples over time in a way that ensures an effectively balanced class distribution (see Fig. 7). We do this both for marginal and conditional replay in order to ensure a fair comparison,




(a) accuracy for MNIST disjoint CLT




(b) accuracy for Fashion MNIST disjoint CLT




(c) accuracy for MNIST permutation CLT



(d) accuracy for Fashion MNIST permutation CLT



(e) accuracy for MNIST rotation CLT



(f) accuracy for Fashion MNIST rotation CLT

Fig. 5: Test set accuracies during training on different CLTs, shown for all sub-tasks (indicated by dotted lines).

although technically conditional replay can generate balanced distribution even with a constant number of generated samples.




Fig. 6: We compare final accuracy when the ratio between size of old task and size of new task is 1 (balanced) or 1/10 (unbalanced, factor was chosen empirically).

- demonstrating the advantages of conditional w.r.t. marginal replay strategies, especially when only few samples can be generated, thus obtaining a skewed class distribution for marginal replay (see Fig. 7).

Results shedding light on the first question are presented in Fig. 5 (showing classification accuracy on whole test set over time, see Fig. 8 for accuracy on first sub-task), whereas the second question is addressed in Fig. 6 for the disjoint CTL only due to space limitations.

4 Results and discussion

From the experiments described in the previous section, we can state the following principal findings:

Replay methods outperform EWC As can be observed from Fig. 5, the novel methods we propose (marginal replay with GAN and WGAN-GP, conditional replay with CGAN and conditional replay with CVAE) outperform EWC, on all CLTs, sometimes by a large margin. Particular attention should be given to the performance of EWC: while generally




Fig. 7: Why marginal replay must linearly increase the number of generated samples: distribution of classes produced by the generator of a marginal replay strategy after sequential training of 10 sub-tasks (of 1 class each). This essentially corresponds to the "disjoint" type of CLTs. Shown are three cases: "*balanced: tN* " (blue bars) where tN samples are generated for each sub-task t , "*unbalanced: N* " (orange bars) where the number of generated samples is constant and equal to the number of newly trained samples N for each sub-task, and "*unbalanced: $0.1tN$* " where $0.1tN$ samples are generated. We observe that, in order to ensure a balanced distribution of classes, the number of generated samples must be re-scaled, or, in other words, must increase linearly with the number of sub-tasks.

acceptable for rotation and permutation CLTs, it completely fails for the disjoint CLT. This is due to the fact that there is only one class in each sub-task, making EWC try to map all samples to the currently presented class label regardless of input, since no replay is available to include samples from previous sub-tasks (as outlined before in Sec. 1.1).

Marginal replay with GAN outperforms WGAN-GP The clear advantage of GAN over WGAN-GP is the higher stability of the generative models. This is not only observable in Fig. 5, but also when measuring performance on the first sub-task only during the course of continual learning (see Fig.8).

Conditional replay can be run at constant time complexity A very important point in favour of conditional replay is run-time complexity, as expressed by the number of samples that need to be generated each time a new sub-task is trained. Since the generators in marginal replay strategies generate samples regardless of class, the distribution of classes will be proportional to the distribution of classes during the last training of the generator, which leads to an unbalanced class distribution over time, with the oldest classes being strongly under-represented (see Fig. 7). This is avoided by increasing the number of generated samples over time for marginal replay, leading to a balanced class distribution (see also Fig. 7) while vastly increasing the number of samples. Conditional replay, on the other hand, can selectively generate samples from a defined class, thus constructing a class-balanced dataset without needing to increase the number of generated samples over time. In the interest of accuracy, it can of course make sense to increase the number of generated samples over time, just as for marginal replay. This, however, is a deliberate




Fig. 8: Comparison of the accuracy of each approach on the first sub-task. This is another, very intuitive measure of how much is forgotten during continual learning. Means and standard deviations computed over 8 seeds.

choice and not something required by conditional replay itself.

Marginal replay outperforms conditional replay when many samples can be generated From Fig. 5, it can be observed that marginal replay outperforms conditional replay by a small margin. This comes at the price of having to generate a large number of samples, which will become unfeasible if many classes are involved in the retraining.

Conditional replay is superior when few samples are generated The results of Fig. 6 show that conditional replay is superior to marginal replay when generating fewer samples at each sub-task (more precisely: $0.1tN$ samples instead of tN , for sub-task t and number of new samples per sub-task N). This can be understood quite easily: since we generate only $0.1tN$ samples instead of tN samples at each sub-task, marginal replay produces an unbalanced class distribution, see Fig. 7, which strongly impairs classification performance. This is a principal advantage that conditional replay has over marginal replay: generating balanced class distributions while having much more control over the number of generated samples.

5 Conclusions

Summary We have proposed several of performing continual learning with replay-based models and empirically demonstrated (on novel benchmarks) their merit w.r.t. the state

of the art, represented by the EWC method. A principal conclusion of this article is that conditional replay methods show strong promise because they have competitive performance, and they impose less restrictions on their use in applications. Most notably, they can be used at constant time complexity, meaning that the number of generated samples does not need to increase over time, which would be problematic in applications with many sub-tasks and real-time constraints.

Concerning the benchmarks While one might argue that MNIST and FashionMNIST are too simple for a meaningful evaluation, this holds only for non-continual learning scenarios. In fact, recent articles [29] show that MNIST-related CLTs are still a major obstacle for most current approaches to continual learning under realistic conditions. So, while we agree that MNIST and FashionMNIST are not suitable benchmarks in general anymore, we must stress the difficulty of MNIST-related CLTs in continual learning, thus making these benchmarks very suitable indeed in this particular context. The use of intrinsically more complex benchmarks, such as CIFAR, SVHN or ImageNet is at present not really possible since generative methods are not really good enough for replaying these data [21].

Next steps Future work will include a closer study of conditional replay: in particular, we would like to better understand why they exhibit better performance w.r.t marginal replay in cases where the number of generated samples is restricted to be low. In addition, it would be interesting to study the continual learning behavior of conditional replay models when a fixed number of generated samples is imposed at each sub-task, for various CLTs. The latter topic is interesting because the success of replay-based continual learning methods in applications will depend on whether the number of generated samples (and thereby time and memory complexity) can be reduced to manageable levels.

Observations An interesting point is that the disjoint type CLTs pose enormous problems to conventional machine learning architectures, and therefore represent a very useful benchmark for continual learning algorithms. If each of a CLT's sub-tasks contains a single visual class, training them one after the other will induce no between-class discrimination at all since every training step just "sees" a single class. Replay-based methods nicely bridge this gap, allowing continual learning while allowing between-class discrimination. To our mind, any application-relevant algorithm for continual learning therefore must include some form of experience replay.

Outlook Ultimately, the goal of our research is to come up with replay-based models where the effort spent on replaying past knowledge is small compared to the effort of training with new samples, which will require machine learning models that are, intrinsically, less prone to catastrophic forgetting than DNNs are.

References

1. Aljundi, R., Babiloni, F., Elhoseiny, M., Rohrbach, M., Tuytelaars, T.: Memory aware synapses: Learning what (not) to forget. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 139–154 (2018)

2. Aljundi, R., Chakravarty, P., Tuytelaars, T.: Expert gate lifelong learning with a network of experts. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3366–3375 (2017)
3. Chaudhry, A., Dokania, P.K., Ajanthan, T., Torr, P.H.: Riemannian walk for incremental learning: Understanding forgetting and intransigence. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 532–547 (2018)
4. 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 (2017)
5. French, R.: Catastrophic forgetting in connectionist networks. Trends in Cognitive Sciences (4) (1999)
6. Gepperth, A., Wiech, F.: Simplified computation and interpretation of fisher matrices in incremental learning with deep neural networks. In: ICANN, submitted (2019)
7. Gepperth, A., Hammer, B.: Incremental learning algorithms and applications. In: European Symposium on Artificial Neural Networks (ESANN) (2016)
8. Gepperth, A., Karaoguz, C.: A bio-inspired incremental learning architecture for applied perceptual problems. Cognitive Computation **8**(5), 924–934 (Oct 2016). <https://doi.org/10.1007/s12559-016-9389-5>, <https://doi.org/10.1007/s12559-016-9389-5>
9. 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)
10. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q. (eds.) Advances in Neural Information Processing Systems 27, pp. 2672–2680. Curran Associates, Inc. (2014), <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>
11. Goodfellow, I.J., Mirza, M., Xiao, D., Courville, A., Bengio, Y.: An Empirical Investigation of Catastrophic Forgetting in Gradient-Based Neural Networks (2013). <https://doi.org/10.1088/1751-8113/44/8/085201>, <http://arxiv.org/abs/1312.6211>
12. Hinton, G.E., Srivastava, N., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.R.: Improving neural networks by preventing co-adaptation of feature detectors pp. 1–18 (2012). <https://doi.org/arXiv:1207.0580>, <http://arxiv.org/abs/1207.0580>
13. Kamra, N., Gupta, U., Liu, Y.: Deep Generative Dual Memory Network for Continual Learning. ArXiv e-prints (Oct 2017)
14. Kamra, N., Gupta, U., Liu, Y.: Deep generative dual memory network for continual learning. arXiv preprint arXiv:1710.10368 (2017)
15. Kemker, R., Kanan, C.: Fearnnet: Brain-inspired model for incremental learning. arXiv preprint arXiv:1711.10563 (2017)
16. Kim, H., Kim, S., Lee, J.: Keep and learn: Continual learning by constraining the latent space for knowledge preservation in neural networks. CoRR **abs/1805.10784** (2018), <http://arxiv.org/abs/1805.10784>
17. Kim, H.E., Kim, S., Lee, J.: Keep and learn: Continual learning by constraining the latent space for knowledge preservation in neural networks. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 520–528. Springer (2018)
18. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., Hadsell, R.: Overcoming catastrophic forgetting in neural networks (2016). <https://doi.org/10.1073/pnas.1611835114>, <http://arxiv.org/abs/1612.00796>

19. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-Based Learning Applied to Document Recognition (1998), <http://ieeexplore.ieee.org/document/726791/{#}full-text-section>
20. Lee, S.W., Kim, J.H., Jun, J., Ha, J.W., Zhang, B.T.: Overcoming catastrophic forgetting by incremental moment matching. In: *Advances in Neural Information Processing Systems*. pp. 4655–4665 (2017)
21. Lesort, T., Caselles-Dupré, H., Garcia-Ortiz, M., Stoian, A., Filliat, D.: Generative Models from the perspective of Continual Learning. arXiv e-prints arXiv:1812.09111 (Dec 2018)
22. Lesort, T., Caselles-Dupré, H., Garcia-Ortiz, M., Stoian, A., Filliat, D.: Generative models from the perspective of continual learning. arXiv preprint arXiv:1812.09111 (2018)
23. Lesort, T., Goudou, J.F., Filliat, D.: Training discriminative models to evaluate generative ones (2018)
24. Li, Z., Hoiem, D.: Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence* **40**(12), 2935–2947 (2018)
25. Liu, X., Masana, M., Herranz, L., Van de Weijer, J., Lopez, A.M., Bagdanov, A.D.: Rotate your networks: Better weight consolidation and less catastrophic forgetting. In: *2018 24th International Conference on Pattern Recognition (ICPR)*. pp. 2262–2268. IEEE (2018)
26. Mallya, A., Lazebnik, S.: Packnet: Adding multiple tasks to a single network by iterative pruning. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 7765–7773 (2018)
27. Mirza, M., Osindero, S.: Conditional generative adversarial nets (2014)
28. Parisi, G.I., Kemker, R., Part, J.L., Kanan, C., Wermter, S.: Continual lifelong learning with neural networks: A review. arXiv preprint arXiv:1802.07569 (2018)
29. Pfülb, B., Gepperth, A.: A comprehensive, application-oriented study of catastrophic forgetting in dnns. In: *International Conference on Learning Representations (ICLR)* (2019), accepted
30. Pfülb, B., Gepperth, A., Abdullah, S., Krawczyk, A.: Catastrophic forgetting: still a problem for dnns. In: *International Conference on Artificial Neural Networks (ICANN)* (2018)
31. Rahaf Aljundi, M.R., Tuytelaars, T.: Selfless sequential learning. *CoRR* **abs/1806.05421** (2018), <http://arxiv.org/abs/1806.05421>
32. Rebuffi, S.A., Kolesnikov, A., Sperl, G., Lampert, C.H.: icarl: Incremental classifier and representation learning. In: *Proc. CVPR* (2017)
33. Ren, B., Wang, H., Li, J., Gao, H.: Life-long learning based on dynamic combination model. *Applied Soft Computing* **56**, 398–404 (2017)
34. 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 (2016)
35. Serrà, J., Suris, D., Miron, M., Karatzoglou, A.: Overcoming catastrophic forgetting with hard attention to the task. In: *The 35th International Conference on Machine Learning (ICML 2018)* (2018), <https://arxiv.org/abs/1801.01423>
36. Serrà, J., Suris, D., Miron, M., Karatzoglou, A.: Overcoming catastrophic forgetting with hard attention to the task. arXiv preprint arXiv:1801.01423 (2018)
37. Shah, H., Javed, K., Shafait, F.: Distillation techniques for pseudo-rehearsal based incremental learning. arXiv preprint arXiv:1807.02799 (2018)
38. Shin, H., Lee, J.K., Kim, J., Kim, J.: Continual learning with deep generative replay. In: *Advances in Neural Information Processing Systems*. pp. 2990–2999 (2017)
39. Shmelkov, K., Schmid, C., Alahari, K.: Incremental learning of object detectors without catastrophic forgetting. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 3400–3409 (2017)

40. Sohn, K., Lee, H., Yan, X.: Learning structured output representation using deep conditional generative models. In: Cortes, C., Lawrence, N.D., Lee, D.D., Sugiyama, M., Garnett, R. (eds.) *Advances in Neural Information Processing Systems 28*, pp. 3483–3491. Curran Associates, Inc. (2015)
41. Srivastava, R.K., Masci, J., Kazerounian, S., Gomez, F., Schmidhuber, J.: Compete to Compute. *Nips* pp. 2310–2318 (2013)
42. Wu, C., Herranz, L., Liu, X., Wang, Y., van de Weijer, J., Raducanu, B.: Memory replay gans: learning to generate images from new categories without forgetting. *arXiv preprint arXiv:1809.02058* (2018)
43. Wu, Y., Chen, Y., Wang, L., Ye, Y., Liu, Z., Guo, Y., Zhang, Z., Fu, Y.: Incremental classifier learning with generative adversarial networks. *CoRR* **abs/1802.00853** (2018), <http://arxiv.org/abs/1802.00853>
44. Xiao, H., Rasul, K., Vollgraf, R.: Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms pp. 1–6 (2017), <http://arxiv.org/abs/1708.07747>
45. Yoon, J., Yang, E., Lee, J., Hwang, S.J.: Lifelong learning with dynamically expandable networks. *arXiv preprint arXiv:1708.01547* (2017)