

# Integrating Curricula with Replays: Its Effects on Continual Learning

Ren Jie Tee<sup>1,2</sup>, Mengmi Zhang<sup>2,3</sup>

<sup>1</sup> Hwa Chong Institution, Singapore

<sup>2</sup> CFAR and I2R, Agency for Science, Technology, and Research, Singapore

<sup>3</sup> Nanyang Technological University, Singapore  
mengmi@i2r.a-star.edu.sg

## Abstract

Humans engage in learning and reviewing processes with curricula when acquiring new skills or knowledge. This human learning behavior has inspired the integration of curricula with replay methods in continual learning agents. The goal is to emulate the human learning process, thereby improving knowledge retention and facilitating learning transfer. Existing replay methods in continual learning agents involve the random selection and ordering of data from previous tasks, which has shown to be effective. However, limited research has explored the integration of different curricula with replay methods to enhance continual learning. Our study takes initial steps in examining the impact of integrating curricula with replay methods on continual learning in three specific aspects: the interleaved frequency of replayed exemplars with training data, the sequence in which exemplars are replayed, and the strategy for selecting exemplars into the replay buffer. These aspects of curricula design align with cognitive psychology principles and leverage the benefits of interleaved practice during replays, easy-to-hard rehearsal, and exemplar selection strategy involving exemplars from a uniform distribution of difficulties. Based on our results, these three curricula effectively mitigated catastrophic forgetting and enhanced positive knowledge transfer, demonstrating the potential of curricula in advancing continual learning methodologies. Our code and data are available: <https://github.com/ZhangLab-DeepNeuroCogLab/Integrating-Curricula-with-Replays>

## Introduction

Continual learning enables consecutive task acquisition without forgetting previously trained tasks (Kirkpatrick et al. 2017). This adaptability is vital for autonomous systems in dynamic environments, such as updating a grocery classification model with new products without retraining it on previous products. However, a significant challenge in continual learning is catastrophic forgetting, where knowledge from recent tasks interferes with earlier ones (Parisi et al. 2019), leading to performance degradation on earlier tasks after training on a task sequence.

To resolve this problem, there are three primary types of continual learning methods commonly employed in the field: regularization-based methods introduce regularization

terms to mitigate catastrophic forgetting by preserving important parameters during training (Chaudhry et al. 2018; Zenke, Poole, and Ganguli 2017; Aljundi et al. 2018; Benzing 2022); rehearsal-based methods store and replay a subset of previous data during training to maintain knowledge from previous tasks (Rolnick et al. 2019; Chaudhry et al. 2019; Riemer et al. 2018; Vitter 1985; Rebuffi et al. 2017; Castro et al. 2018) and parameter isolation methods isolate specific parameters for each task to prevent interference between tasks (Yoon et al. 2017; Hung et al. 2019; Ostapenko et al. 2019). Rehearsal-based methods have proven highly effective in continual learning. However, existing approaches typically involve randomly selecting and rehearsing data from previous tasks. Limited research explores the incorporation of meaningful curricula into replay methods.

In parallel, in the curriculum learning literature, various approaches have focused on weakly supervised (Guo et al. 2018), unsupervised (Azad et al. 2023), and reinforcement learning tasks (Narvekar et al. 2020). These studies demonstrate that curricula improve generalization abilities, task performances, and convergence speed (Wu, Dyer, and Neyshabur 2021; Soviany et al. 2022) during training. However, they primarily address intra-class difficulty and example scheduling within a single task, neglecting the impact of class presentation sequences across multiple tasks. Recent research has explored curricula in continual learning scenarios without data replays (Singh et al. 2022). In complement to this work, our study investigates the role of curricula specifically during replay in continual learning, while keeping the curricula consistent for the feed-forward training process.

Exploring optimal curricula offers countless possibilities, and in our study, we take initial steps to investigate a limited set of potential curricula. We draw inspiration from two sources to guide the design of these curricula. Firstly, neuroscience research has revealed that neural activity patterns associated with past experiences are replayed in specific orders during rest or sleep, which is believed to contribute to memory consolidation and spatial navigation (Drieu and Zugaro 2019). Secondly, pedagogy studies indicate that repetitive practice and revisiting previous knowledge with increasing difficulty enhance long-term memory integration in students (Zhan et al. 2018).

Specifically, we propose three types of curricula for replays and examine their impact on catastrophic forgetting and positive knowledge transfer: (1) the interleaved frequency of replayed exemplars with training data, (2) the replay sequence of exemplars, and (3) the strategy for selecting exemplars into the replay buffer. The experimental findings align with cognitive psychology principles, highlighting the advantages of frequently interleaving between training data and replayed exemplars, incorporating easy-to-hard rehearsals, and selecting exemplars from a uniform distribution of difficulties for replay. These observations present a promising avenue for advancing continual learning methods. It also provides insights into the underlying mechanisms of replay strategies in mitigating forgetting and facilitating knowledge transfer across tasks.

## Related Works

### Replay Methods in Continual Learning

Extensive research has focused on utilizing replay methods to address the issue of catastrophic forgetting. Conventional replay methods, such as iCaRL (Rebuffi et al. 2017) and ER (Rolnick et al. 2019), involve explicit training on previously saved data, while several variants, like DGR (Shin et al. 2017) and Pseudo-Recursal (Atkinson et al. 2018), replay on artificially synthesized samples by generative models, resembling data from previous tasks. Although these replay methods have made significant contributions in reducing catastrophic forgetting, they paid little attention to the incorporation of meaningful curricula into replay methods. Most methods randomly interleave the replay samples with the training data, without exploring the optimal mixing strategies (Rolnick et al. 2019; van de Ven, Siegelmann, and Tolias 2020; Ketz, Kolouri, and Pilly 2019). In our work, we systematically studied the effect of interleaving curricula, which involves mixing training data and replay samples within a pre-defined interleave interval.

### Curriculum Learning

Curriculum learning methods can be broadly categorized into two groups. The first group involves manual curriculum design by humans before training (Soviany et al. 2020; Chen and Gupta 2015), but these methods typically rely on human expertise and struggle to generalize to new domains. The second group consists of models that can autonomously design curricula without human intervention (Graves et al. 2017; Fan et al. 2018). However, the application of these methods to enhance model performance has received limited attention in the continual learning setting. Here, we highlight two factors to consider when applying curricula on the replay methods in continual learning. Firstly, while curriculum learning has demonstrated efficacy in enhancing generalization and training speed within a single task, the objective of curriculum learning in the context of continual learning is to retain knowledge from previous tasks while acquiring new knowledge from the current task. Secondly, unlike within-task curriculum learning, models in continual learning only have access to data from the current task, making it challenging to create a

comprehensive between-task curriculum that encompasses the entire dataset. Here, we took initial steps in this direction by exploring automated methods to determine the sequence of replay samples and introducing the sample selection strategy which finds the best replay samples for building a curriculum.

## Experiments

We investigated the effect of three types of replay curricula in the class incremental learning (CIL) setting. We first introduce CIL, and then elaborate on the three replay curricula individually.

**Problem Setting.** The objective of CIL is to teach a unified classification model  $\Theta$  to recognize sets of object classes incrementally over time. Specifically, an image dataset  $D$ , consisting of  $N$  object classes, is split into subsets  $\{D_1, \dots, D_t, \dots, D_T\}$  of images and presented over a sequence of  $T$  tasks. In each task  $t$ , the model only has access to training data in  $D_t$ , consisting of samples from distinct classes  $C_t$ , and  $(x_{i,t}, y_{i,t})$  is the  $i$ -th (image, label) pair in  $D_t$ . The model  $\Theta$  can run multiple passes over  $D_t$  in task  $t$ . The model stops training on  $D_t$  after its performance on the validation set saturates, considering the five most recent epochs.

We implemented the naive replay method where some raw images and their corresponding labels are selected from previous tasks and are stored in the replay buffer  $R_t$ . These data in  $R_t$  are inter-leaved with  $D_t$  for rehearsals. There are three types of replay curricula involved in this study: (1) the interleave frequency; (2) the rehearsal sequence of  $R_t$  in CIL; and (3) the image selection for  $R_t$ .  $R_t$  is kept at a constant size of 1200 over all the tasks. See **Appendix** for more training details. As an upper bound, we also include the **offline method** where the model  $\Theta$  is trained on the entire dataset  $\Theta$  from  $\{D_1, \dots, D_T\}$  over multiple epochs without any continual learning.

**Datasets.** We conducted experiments to investigate the use of these three types of curricula in replay methods on the two image datasets ciFAIR-10 and ciFAIR-100 (Barz and Denzler 2019). ciFAIR-10 dataset contains 10 object classes. The protocol asks the model  $\Theta$  to incrementally learn 2 object classes in each task. There are a total of 5 tasks. ciFAIR-100 dataset contains 100 object classes. The CIL protocol asks the model  $\Theta$  to incrementally learn 5 object classes in each task. There are a total of 20 tasks. Both datasets have a total of 60,000 images, with 50,000 images used for training and 10,000 images used for testing. The conclusions drawn from the experiments on both datasets are consistent. Without loss of generality, we focus on all the experiments and result analysis in ciFAIR-100 in the main text. See **Appendix** for more implementation details and results on ciFAIR-10.

**Evaluation Metrics.** To assess the continual learning performance of the model  $\Theta$ , we followed (Singh et al. 2022) and introduce 2 standard evaluation metrics. We define **Forgetfulness (F)** as the percentage decrease in classification accuracy on the test instances from  $C_1$  between the  $\Theta_t$  (model after being trained on  $D_t$ ) and  $\Theta_1$ . An ideal  $\Theta_t$  could maintain the same classification accuracy

on  $C_1$  over tasks; i.e.  $\forall t, F_t = 0$ . The higher  $F$  is, the more  $\Theta$  suffers from catastrophic forgetting. To assess the overall classification performance of  $\Theta$  over tasks, we also report the continual average classification accuracy (**Avg. Accu.**). Avg. Accu. is computed as the average accuracy on all the test instance from  $C_i$ , where  $i \in \{1, 2, \dots, t\}$ . For simplicity, we report the averaged **F** and **Avg. Accu.** over all the tasks. **Experimental Controls.** Within each experiment, only one variable of interest changes while the rest of the experiment conditions are fixed as control variables. As the previous study has shown that the sequence of class presentations affects the continual learning performance (He, Wang, and Chen 2022), we use the same class presentation sequence in all three experiments. The same MobileNetV3 (small) network architecture is used as the backbone for the model  $\Theta$  for all experiments. In every experiment, the total number of training samples and the total number of replay samples exposed to  $\Theta$  remain the same across all experiment variables. Each experiment is conducted with 4 runs initialized with 4 random seeds, where the seeds are used to vary the controlled variables. The average performance across all 4 runs is reported.

### Interleave Divisions during Rehearsals

The number of interleaving divisions refers to the number of splits of in  $D_t$  and  $R_t$ . It indicates how often the model  $\Theta$  rehearses on  $R_t$ , while learning on a subset of  $D_t$ . For example, for interleaving division number 400,  $D_t$  is split into 400 groups where each group contains an equal number of  $(x_{i,t}, y_{i,t})$  (image, label) pairs, and these (image, label) pairs are randomly selected from  $D_t$  without replacement. Correspondingly,  $R_t$  is also split into 400 groups with the same splitting criteria as  $D_t$ . At each training epoch, the model  $\Theta_t$  at task  $t$  is repeatedly trained with one group of  $D_t$  followed by one group of  $R_t$ , until the entire  $D_t$  and  $R_t$  are exhaustively seen by  $\Theta_t$ . We titrate the interleave division numbers with the range of 1, 8, 60, 120, and 300. The training data is interleaved with replay data and then presented to the model in sequence. Different interleave division numbers result in different data presentation sequences; hence, different curricula.

### Rehearsal Sequence of Replay Samples

We use the interleave divisions 1 and 600 for all the experiments in this subsection and vary the rehearsal sequence of data samples in  $R_t$  by taking into account the two factors: the sample difficulty levels and the increasing or decreasing directions of sample difficulty levels.

To measure whether a sample is easy or hard to learn, we introduce two difficulty measures: (1) the confidence score difficulty metrics and (2) the distance vector difficulty metrics. The confidence score difficulty metrics were used to assess whether a teacher network with full knowledge of the entire dataset  $D$  predicted high or low confidence of the given sample belonging to its ground truth class label. Specifically, each image within  $R_t$  was input to a teacher network. The teacher network is based on a MobileNetV3 (small) architecture, pre-trained on the entire dataset  $D$ . After this, the confidence score for the ground truth class

of each sample was extracted from the teacher network’s output.  $R_t$  was then sorted according to its individual sample’s confidence score, where a higher confidence score means that the sample is easier to learn for  $\Theta$ .

However, in CIL setting, having a teacher network with full access to the whole dataset is impractical, as the data is incrementally available over tasks. Hence, we employed the distance vector difficulty metrics, used widely in literature (Singh et al. 2022; Huang et al. 2022). Intuitively, if the sample is closer to other samples in the memory buffer, it is easier for  $\Theta$  to learn and generalize to other samples as well. The 2nd last layer from a ResNet-50 model (He et al. 2016), pretrained on the ImageNet dataset, was used to extract the feature vector of each sample in  $R_t$ . A Euclidean distance matrix was created, where the pairwise Euclidean distance for all the samples based on their feature vectors was computed. We then compute the sum of each row of the matrix and denote this column vector as a distance vector. Each element in this distance vector represents how much a particular sample differs from all other samples in the feature space. A smaller value in the distance vector means that the particular replay sample is easier to learn for  $\Theta$ .

We introduce a series of rehearsal sequences in the orders of either easy-to-hard samples or hard-to-easy samples, where the difficulty levels of each sample are determined by either the confidence score difficulty metrics or the distance vector difficulty metrics. As the previous study has shown that the class orders are also essential for continual learning (He, Wang, and Chen 2022), here we also explore the effect of the class orders during replays. When we design the rehearsal sequence based on class difficulties in  $R_t$ , we adapt the two sample-level difficulty measures above to compute class-level difficulty measures by taking the average over all samples of the same class in  $R_t$ . We then sort all the samples in  $R_t$  by their class difficulty metrics, regardless of their individual sample difficulty scores. Samples in  $R_t$  sorted by their class difficulties were then presented to the model  $\Theta$  in either the easy-to-hard or hard-to-easy orders.

### Selection of Samples for Replay Buffer

In common practice, selecting samples for  $R_{t+1}$  from task  $t$  is often conducted in a random manner (Rolnick et al. 2019; Bagus and Geppert 2021). In contrast to the previous works, we vary the sample selection criteria for  $R_{t+1}$  as follows: selecting only the easiest samples from task  $t$  for  $R_{t+1}$ , selecting the hardest samples from task  $t$  for  $R_{t+1}$ , and selecting samples that are uniformly distributed across difficulty levels from task  $t$  for  $R_{t+1}$ . The difficulty levels are judged based on the confidence scores and the distance vectors defined in the previous subsection. We use interleave division numbers 1 and 600 for all the experiments in this subsection.

## Results

### Frequent Replays Enhance Performances

We report **F** and **Avg. Accu.** as a function of interleave divisions in Table 1. Notably, we observed that interleave divisions are important factors influencing the continual

Inter. Div.	1	8	60	120	300
F	63.9	62.6	57.4	<b>55.1</b>	56.1
Avg. Accu	39.9	40.7	44.6	<b>46.6</b>	<b>46.6</b>

Table 1: Continual learning performance on ciFAIR-100 as a function of interleave divisions (Inter. Div.). The higher Avg. Accu., the better; the lower F, the better. The best are in bold. The offline accuracy is 62.0%. See Experiments Section for evaluation metrics and experiments in interleave divisions.




Figure 1: Continual learning performance on ciFAIR-100 dataset with different rehearsal sequences. Average classification accuracy over all tasks for models trained with different rehearsal sequences is reported. The higher the average accuracy, the better. See the legend on the right for conventions of different rehearsal sequence conditions. For example, a light grey bar with a red boundary without the texture of slanted bar patterns (the leftmost bar) indicates that the model is trained with an easy-to-hard rehearsal sequence sorted by the instance-level difficulties based on distance vector metrics. See Experiments Section for details. See Figure 4 in Appendix for the results for F.

learning performance of the replay method with the larger interleave divisions leading to better performances, as indicated by the decreasing **F** and increasing **Avg. Accu.** over all the tasks. It is possible that the model parameters at large division numbers are updated more frequently for both the current task and all previous tasks, resulting in minimal forgetting. However, we also note that the continual learning performance saturates at interleave division number 120. This implies that increasing interleave divisions beyond optimal values brings no extra benefits in continual learning.

### Easy-to-Hard Rehearsal Sequences Are Beneficial

We studied the models trained with different rehearsal sequences sorted in easy-to-hard or hard-to-easy curricula based on sample-level or class-level difficulty measures computed from either the confidence scores or distance vectors. We reported the Avg. Accu. results in Figure 1 and F scores in **Appendix** and made four key observations. First, aligning with the observations in Table 1 and the discussion from the previous subsection, large interleave divisions benefit continual learning models with higher average accuracy and less forgetting. Second, rehearsal sequences sorted by instance-level difficulties lead to much better continual learning performances (compare red bars





Figure 2: Continual learning performance on ciFAIR-100 dataset under different replay sample selection strategies. Average classification accuracy over all tasks for models trained with different replay samples selected using different strategies. The interpretation of legends and design conventions are similar to Figure 1. See Experiments Section for details. See Figure 6 in Appendix for the results for F.


versus blue bars). Third, the confidence score is a better evaluation metric measuring instance-level difficulties, as shown by the bars with and without texture patterns. Finally, the models trained with the easy-to-hard rehearsal sequences outperform the ones with reversed rehearsal sequences (compare light versus dark grey bars). It is possible that easy-to-hard rehearsal sequences make the models converge faster on the previous tasks due to more stable gradient updates; hence, the sequences lead to minimal forgetting and higher classification accuracy. We also compared the continual learning performance for both the offline method and the continual learning method with various curricula and observed that there still exists a large performance gap between these two.

### Replays with Only Hard Data Hurt Performances

Here, we explored the effect of different sample selection strategies for replay samples in terms of the sample difficulty levels based on distance vectors or confidence scores. From Figure 2. Our observations indicate that exclusively choosing the most challenging replay samples leads to inferior performance compared to selecting the easiest samples or incorporating samples with a balanced distribution of difficulty levels. Selecting samples with a uniform distribution of difficulty levels yields the best continual learning performance. This outcome may be attributed to the fact that difficult replay samples result in less flat loss landscapes, which in turn make the training process more challenging and slower to converge (Wu, Dyer, and Neyshabur 2021). A curriculum for training the models to rehearse from the easiest to the hardest samples is the best, as it balances the greater precision in data fitting due to the hardest samples and the fast convergence speed during training due to the easier samples. Similar to the previous subsection, we also noted that the confidence score is a better measure of sample difficulty levels than the distance vectors.



(a) Avg. Classification Accuracy




(b) Forgetfulness

Figure 3: Continual learning performance on ciFAIR-10 dataset under different rehearsal sequences. Average classification accuracy in (a) and Forgetfulness (F) in (b) over all tasks for models trained with different rehearsal sequences are reported. The higher the average accuracy, the better. The lower the F, the better. The design conventions follow the ones in Figure 1.


## Conclusion

Our study examines the role of curricula during replays in the class-incremental learning setting in continual learning. We designed and conducted a series of controlled experiments to study the three key questions on replays: how often is the replay, what data should be replayed, and in what sequence to replay. Across the two common image datasets, our experimental results shed light on the underlying principles of replay methods in continual learning and reveal the good curricula design choices for replay methods. These curricula designs not only facilitate positive knowledge transfers (which has been explored in existing curriculum learning literature), but also mitigate catastrophic forgetting (a significant problem we need to solve in continual learning). Specifically, we found that (1) replays should happen frequently; (2) only rehearsing on the most difficult exemplars hurts continual learning performances; and (3) rehearsals on samples with increasing difficulty eliminate forgetting more than its reversed difficulty orders.

There are numerous other possible choices of curricula designs for replay methods, such as a unified difficulty



(a) Avg. Classification Accuracy



(b) Forgetfulness

Figure 4: Continual learning performance on ciFAIR-100 dataset under different rehearsal sequences. Complementary to the average accuracy results in Figure 1, we also report the results of F scores on the ciFAIR-100 dataset in (b). The results of Avg. Accu. in (a) are duplicated from Figure 1 and are shown here again for completeness.


metric considering both confidence scores and distance vectors or the use of a student feedback loop to update the difficulty scores. In the future, we will look into the role of curricula under stringent continual learning conditions, such as learning with limited training time or noisy data. We will also conduct experiments on other large-scale datasets and apply our replay curriculum to existing replay-based continual learning algorithms.

## A Experimental Details


For experiments on both ciFAIR-10 and ciFAIR-100, PyTorch’s default implementation of cross entropy loss was used for object classification tasks. The SGD algorithm was used as the optimizer. The learning rate was set at a constant of 0.001. Momentum was fixed at 0.9. A batch size of 32 is used.

For ciFAIR-10, we employ a 2-layer 2D-convolutional network with 6 and 16 channels in the successive layers, followed by 3 fully connected layers with 400, 120 and 84 hidden units respectively. ReLU was used as the activation function.

We follow the standard training and testing data splits



(a) Avg. Classification Accuracy



(b) Forgetfulness

Figure 5: Continual learning performance on ciFAIR-10 dataset under different replay sample selection strategies. Average classification accuracy in (a) and Forgetfulness scores (F) in (b) over all the tasks for models trained with different rehearsal samples based on different selection strategies are reported. The higher the average accuracy, the better. The lower the F, the better. The design conventions follow the ones in Figure 2.

from the original ciFAIR-10. In every task, the model is trained for 250 epochs. Each experiment is conducted with 20 runs initialized with 20 random seeds, where the seeds are used to vary the controlled variables. The average performance across all 20 runs is reported.

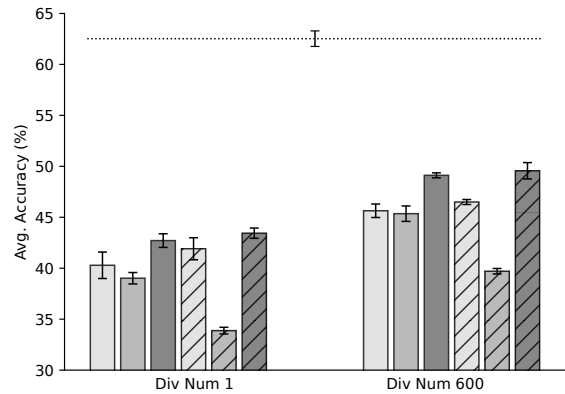
For ciFAIR-100, PyTorch’s implementation of MobileNetV3 (small) was used, including the default layers and activation function.

We used a custom training, validation, and test data splits with a ratio of 9:1:2, and a stopping criteria for training depending on the validation loss. The ciFAIR-100 images were upscalded to 72x72 using PyTorch’s bicubic interpolation function before training.


## B More Results and Analysis

We reported the continual learning performance on ciFAIR-10 dataset of the models trained with the three types of curricula as elaborated in **Experiments Section**.

See Table 2 for interleave divisions, Figure 3 for rehearsal sequences, and Figure 5 for sample selections. All the



(a) Avg. Classification Accuracy



(b) Forgetfulness

Figure 6: Continual learning performance on ciFAIR-100 dataset under different replay sample selection strategies. Complementary to the average accuracy results in Figure 2, we also report the results of F scores on the ciFAIR-100 dataset in (b). The results of Avg. Accu. in (a) are duplicated from Figure 2 and are shown here again for completeness.

Inter. Div.	1	8	50	200	400
F	62.5	59.3	<b>57.1</b>	57.4	57.2
Avg. Accu	53.5	54.9	<b>55.8</b>	<b>55.8</b>	<b>55.8</b>

Table 2: Continual learning performance on ciFAIR-10 as a function of interleave divisions (Inter. Div.). The higher Avg. Accu., the better; the lower F, the better. The best are in bold. See Experiments Section for evaluation metrics and experiments in interleave divisions.

tables and figures on ciFAIR-10 dataset follow the same design conventions as the corresponding tables and figures on ciFAIR-100 dataset in the main text.

The conclusions from the results of ciFAIR-10 dataset are consistent with the ones on the ciFAIR-100 dataset.

## Acknowledgements

This research is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG2-RP-2021-025), its NRFF award NRF-NRFF15-2023-0001, Mengmi Zhang’s Startup Grant from Agency for Science, Technology, and Research

(A\*STAR), and Early Career Investigatorship from Center for Frontier AI Research (CFAR), A\*STAR. The authors declare that they have no competing interests.

## References

- Aljundi, R.; Babiloni, F.; Elhoseiny, M.; Rohrbach, M.; and Tuytelaars, T. 2018. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European conference on computer vision (ECCV)*, 139–154.
- Atkinson, C.; McCane, B.; Szymanski, L.; and Robins, A. V. 2018. Pseudo-Recursal: Solving the Catastrophic Forgetting Problem in Deep Neural Networks. *CoRR*, abs/1802.03875.
- Azad, A. S.; Gur, I.; Emhoff, J.; Alexis, N.; Faust, A.; Abbeel, P.; and Stoica, I. 2023. CLUTR: Curriculum Learning via Unsupervised Task Representation Learning. arXiv:2210.10243.
- Bagus, B.; and Gepperth, A. 2021. An investigation of replay-based approaches for continual learning. In *2021 International Joint Conference on Neural Networks (IJCNN)*, 1–9. IEEE.
- Barz, B.; and Denzler, J. 2019. Do we train on test data? Purging CIFAR of near-duplicates. *CoRR*, abs/1902.00423.
- Benzing, F. 2022. Unifying importance based regularisation methods for continual learning. In *International Conference on Artificial Intelligence and Statistics*, 2372–2396. PMLR.
- Castro, F. M.; Marín-Jiménez, M. J.; Guil, N.; Schmid, C.; and Alahari, K. 2018. End-to-end incremental learning. In *Proceedings of the European conference on computer vision (ECCV)*, 233–248.
- Chaudhry, A.; Dokania, P. K.; Ajanthan, T.; and Torr, P. H. 2018. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European conference on computer vision (ECCV)*, 532–547.
- Chaudhry, A.; Rohrbach, M.; Elhoseiny, M.; Ajanthan, T.; Dokania, P. K.; Torr, P. H.; and Ranzato, M. 2019. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486*.
- Chen, X.; and Gupta, A. 2015. Webly Supervised Learning of Convolutional Networks. In *2015 IEEE International Conference on Computer Vision (ICCV)*, 1431–1439.
- Drieu, C.; and Zugaro, M. 2019. Hippocampal Sequences During Exploration: Mechanisms and Functions. *Frontiers in cellular neuroscience*, 13: 232.
- Fan, Y.; Tian, F.; Qin, T.; Li, X.; and Liu, T. 2018. Learning to Teach. *CoRR*, abs/1805.03643.
- Graves, A.; Bellemare, M. G.; Menick, J.; Munos, R.; and Kavukcuoglu, K. 2017. Automated Curriculum Learning for Neural Networks. In Precup, D.; and Teh, Y. W., eds., *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, 1311–1320. PMLR.
- Guo, S.; Huang, W.; Zhang, H.; Zhuang, C.; Dong, D.; Scott, M. R.; and Huang, D. 2018. Curriculumnet: Weakly supervised learning from large-scale web images. In *Proceedings of the European conference on computer vision (ECCV)*, 135–150.
- He, C.; Wang, R.; and Chen, X. 2022. Rethinking class orders and transferability in class incremental learning. *Pattern Recognition Letters*, 161: 67–73.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.
- Huang, P.; Xu, M.; Zhu, J.; Shi, L.; Fang, F.; and Zhao, D. 2022. Curriculum Reinforcement Learning using Optimal Transport via Gradual Domain Adaptation. arXiv:2210.10195.
- Hung, C.-Y.; Tu, C.-H.; Wu, C.-E.; Chen, C.-H.; Chan, Y.-M.; and Chen, C.-S. 2019. Compacting, picking and growing for unforgetting continual learning. *Advances in Neural Information Processing Systems*, 32.
- Ketz, N.; Kolouri, S.; and Pilly, P. K. 2019. Using World Models for Pseudo-Rehearsal in Continual Learning. *CoRR*, abs/1903.02647.
- 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.; and Hadsell, R. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13): 3521–3526.
- Narvekar, S.; Peng, B.; Leonetti, M.; Sinapov, J.; Taylor, M. E.; and Stone, P. 2020. Curriculum Learning for Reinforcement Learning Domains: A Framework and Survey. *J. Mach. Learn. Res.*, 21(1).
- Ostapenko, O.; Puscas, M.; Klein, T.; Jahnichen, P.; and Nabi, M. 2019. Learning to remember: A synaptic plasticity driven framework for continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11321–11329.
- Parisi, G. I.; Kemker, R.; Part, J. L.; Kanan, C.; and Wermter, S. 2019. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113: 54–71.
- Rebuffi, S.-A.; Kolesnikov, A.; Sperl, G.; and Lampert, C. H. 2017. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2001–2010.
- Riemer, M.; Cases, I.; Ajemian, R.; Liu, M.; Rish, I.; Tu, Y.; and Tesauro, G. 2018. Learning to learn without forgetting by maximizing transfer and minimizing interference. *arXiv preprint arXiv:1810.11910*.
- Rolnick, D.; Ahuja, A.; Schwarz, J.; Lillicrap, T.; and Wayne, G. 2019. Experience replay for continual learning. *Advances in Neural Information Processing Systems*, 32.
- Shin, H.; Lee, J. K.; Kim, J.; and Kim, J. 2017. Continual learning with deep generative replay. *Advances in neural information processing systems*, 30.
- Singh, P.; Li, Y.; Sikarwar, A.; Lei, W.; Gao, D.; Talbot, M. B.; Sun, Y.; Shou, M. Z.; Kreiman, G.; and Zhang, M. 2022. Learning to Learn: How to Continuously Teach Humans and Machines. arXiv:2211.15470.
- Soviany, P.; Ardei, C.; Ionescu, R.; and Leordeanu, M. 2020. Image Difficulty Curriculum for Generative Adversarial

- Networks (CuGAN). In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 3452–3461. Los Alamitos, CA, USA: IEEE Computer Society.
- Soviany, P.; Ionescu, R. T.; Rota, P.; and Sebe, N. 2022. Curriculum Learning: A Survey. *International Journal of Computer Vision*, 130(6): 1526–1565.
- van de Ven, G. M.; Siegelmann, H. T.; and Tolias, A. S. 2020. Brain-inspired replay for continual learning with artificial neural networks. *Nature Communications*, 11(1): 4069.
- Vitter, J. S. 1985. Random sampling with a reservoir. *ACM Transactions on Mathematical Software (TOMS)*, 11(1): 37–57.
- Wu, X.; Dyer, E.; and Neyshabur, B. 2021. When Do Curricula Work? In *International Conference on Learning Representations*.
- Yoon, J.; Yang, E.; Lee, J.; and Hwang, S. J. 2017. Lifelong learning with dynamically expandable networks. *arXiv preprint arXiv:1708.01547*.
- Zenke, F.; Poole, B.; and Ganguli, S. 2017. Continual learning through synaptic intelligence. In *International conference on machine learning*, 3987–3995. PMLR.
- Zhan, L.; Guo, D.; Chen, G.; and Yang, J. 2018. Effects of Repetition Learning on Associative Recognition Over Time: Role of the Hippocampus and Prefrontal Cortex. *Frontiers in human neuroscience*, 12: 277.