

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

CORE: Mitigating Catastrophic Forgetting in Continual Learning through Cognitive Replay

Permalink

<https://escholarship.org/uc/item/00t0z4k7>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Zhang, Jianshu

Fu, Yankai

Peng, Ziheng

et al.

Publication Date

2024

Peer reviewed

CORE: Mitigating Catastrophic Forgetting in Continual Learning through Cognitive Replay

Jianshu Zhang[†] and Yankai Fu[†] and Ziheng Peng[†] and Dongyu Yao and Kun He*
School of Cyber Science and Engineering, Wuhan University, China
{jianshu.zhang, yankaifu, ziheng peng, dongyu.yao, hekun}@whu.edu.cn

Abstract

This paper introduces a novel perspective to significantly mitigate catastrophic forgetting in continuous learning (CL), which emphasizes models’ capacity to preserve existing knowledge and assimilate new information. Current replay-based methods treat every task and data sample equally and thus can not fully exploit the potential of the replay buffer. In response, we propose **CO**gnitive **RE**play (CORE), which draws inspiration from human cognitive review processes. CORE includes two key strategies: Adaptive Quantity Allocation and Quality-Focused Data Selection. The former adaptively modulates the replay buffer allocation for each task based on its forgetting rate, while the latter guarantees the inclusion of representative data that best encapsulates the characteristics of each task within the buffer. Our approach achieves an average accuracy of 37.95% on split-CIFAR10, surpassing the best baseline new method by 6.52%. Additionally, it significantly enhances the accuracy of the poorest-performing task by 6.30% compared to the top baseline. Code is available at <https://github.com/sterzhang/CORE>.

Keywords: continual learning; catastrophic forgetting; data replay; replay buffer; cognitive strategies

Introduction

Deep neural networks have shown remarkable capabilities across many tasks (Sze, Chen, Yang, & Emer, 2017). However, in real-world applications, models must continuously adapt to new tasks rather than remain static. *Continual learning* (Van de Ven & Tolias, 2019) has emerged as a way to dynamically update neural networks, where models continuously acquire knowledge from a stream of tasks (Chen & Liu, 2020). Yet this learning paradigm faces a significant challenge of *Catastrophic Forgetting* (McCloskey & Cohen, 1989; R. M. French, 1999), where models tend to lose previous knowledge when learning new tasks (R. French, 1999), as shown in Figure 1. This problem highlights the urgent need for effective approaches that can help retain old knowledge when assimilating new information.

Numerous studies propose replay-based methods that attempt to mitigate catastrophic forgetting by replaying old data to review old tasks (Mnih et al., 2013; Rolnick, Ahuja, Schwarz, Lillicrap, & Wayne, 2019; Tiwari, Killamsetty, Iyer, & Shenoy, 2022). The core component of these methods is the “replay buffer”, which contains data from previous tasks. However, current methods fail to fully exploit the potential of the replay buffer. They allocate the replay buffer space indiscriminately across all tasks, failing to differentiate the varying rates of forgetting unique to each task. In addition, they

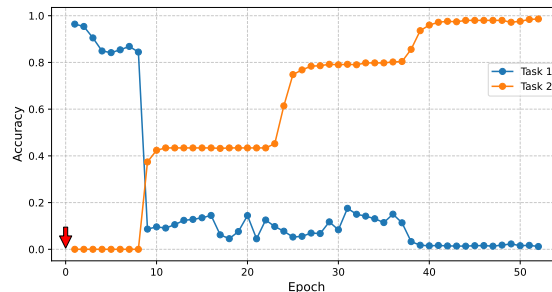


Figure 1: Illustration of Catastrophic Forgetting, which shows Task 1’s rapid decline in accuracy when Task 2 starts to train (marked by the red arrow).

overlook the quality of the data samples in the buffer, which compromises the effectiveness in data replay.

Inspired by studies that highlight similarities between the human brain and neural networks (Schmidgall et al., 2023; McCulloch & Pitts, 1943), and the remarkable success of applying insights from the human brain to improve the performance of neural networks (Mayo et al., 2023; Mnih et al., 2015), we introduce **CO**gnitive **RE**play (CORE) to maximize the effectiveness of the replay buffer. CORE aligns factors of human forgetting with the model’s forgetting, specifically cognitive overload and interference. It evaluates task-specific accuracy to determine the forgetting rate of previous tasks considering both factors. For the newly trained current task, which has not yet experienced cognitive overload, CORE anticipates its potential forgetting by evaluating its interference impact on previous tasks. We then introduce our innovative Adaptive Quantity Allocation (AQA) strategy, which dynamically allocates buffer space to each task based on the calculated attention derived from their respective forgetting and interference rates. We also propose the Quality-Focused Data Selection (QFDS) strategy to address the problem of random sample selection in previous methods (Rebuffi, Kolesnikov, Sperl, & Lampert, 2017; Rolnick et al., 2019). QFDS ensures that the samples within the replay buffer are more representative and balanced for each task. These strategies in buffer allocation and data selection enable CORE to outperform traditional baseline methods and establish a more advanced and effective approach to replay-based methods.

Our contributions can be concluded as follows:

[†] Equal contribution, *Corresponding author

- We are the first to adopt both minimum and average task accuracy values, along with continual monitoring of each task’s performance, which ensures a thorough assessment of model efficacy in continual learning.
- Drawing inspirations from human memory mechanisms, CORE implements two innovative strategies, Adaptive Quantity Allocation and Quality Feature-based Data Selection. These strategies collectively refine the utility and efficacy of replay buffers in continual learning models.
- Our extensive experiments underscore the effectiveness of CORE. For example, on split-CIFAR10, CORE achieves an average accuracy of 37.95%, which is about 6.52% higher than the best baseline method. Remarkably, CORE boosts the accuracy of the most challenging task, exceeding the best baseline by a significant margin of 6.30%.

Theoretical Background

Forgetting in Human Lifelong Learning

In the lifelong learning process of humans (illustrated in Figure 2a), the continual acquisition of new tasks frequently results in the gradual forgetting of previously learned tasks. According to (Sadeh, Ozubko, Winocur, & Moscovitch, 2016), there are two primary memory types: Recall-based Memory, which involves orthogonal representation and is primarily rooted in the hippocampus (Flesch, Juechems, Dumbalska, Saxe, & Summerfield, 2022), and Familiarity-based Memory, which is associated with extra-hippocampal structures and employs non-orthogonal representation.

In Recall-based Memory, characterized by distinct and segregated storage, cognitive overload forgetting occurs when new tasks compete for finite capacity of the system. This orthogonal representation leads to a scenario where new information directly replaces older memories within the limited available space, as illustrated in Figure 2b. In Familiarity-based Memory, the vulnerability to interference arises primarily from its non-orthogonal representation, which features overlapping neural pathways. This structure allows new information to interfere with established memory. Such interference can lead to the deterioration of older memories due to the shared neural connections, as depicted in Figure 2c.

Continual Learning against Forgetting

In response to the catastrophic forgetting in continual learning, current methods predominantly fall into three categories: *Parameter Isolation*, *Regularization*, and *Data Replay*.

Methods based on parameter isolation, such as Progressive Neural Networks (PNN) (Rusu et al., 2016) and context-dependent gating (XdG) (Masse, Grant, & Freedman, 2018), segregate parameters for distinct knowledge sets, akin to the orthogonal encoding in Recall-based Memory to minimize task interference. While effective in reducing interference, these approaches face scalability challenges as the knowledge base expands (Wang, Zhang, Su, & Zhu, 2023).

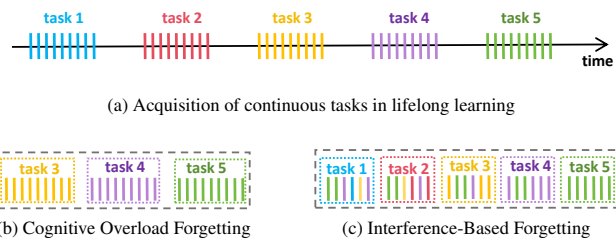


Figure 2: Illustration of continuous task acquisition and forgetting mechanisms in lifelong learning.

Regularization approaches, such as Elastic Weight Consolidation (EWC) (Huszár, 2018) and Learning Without Forgetting (LwF) (Li & Hoiem, 2017), strategically limit important parameter updates to protect previously acquired knowledge. However, the effectiveness of these methods largely depends on the precise determination of parameter significance in maintaining old knowledge (De Lange et al., 2021).

Methods based on data replay, like Experience Replay (ER) (Rolnick et al., 2019) and Incremental Classifier and Representation Learning (iCaRL) (Rebuffi et al., 2017) utilize historical data to review through a ‘replay buffer’. The challenge for these approaches lies in effectively utilizing the replay buffer, specifically in how to allocate buffer space and select representative samples strategically.

Review with Cognitive Strategies

In cognitive science, the synergistic application of the *Targeted Recall* and *Spaced Repetition* strategies (Carpenter, Cepeda, Rohrer, Kang, & Pashler, 2012), along with the *Levels of Processing framework* (Craik & Lockhart, 1972), provides a more effective approach to review.

Targeted Recall emphasizes a more focused and personalized approach that directs heightened attention and resources towards tasks exhibiting higher rates of forgetting. On the other hand, Spaced Repetition emphasizes the importance of regular, periodic review for sustained memory retention. Furthermore, the Levels of Processing framework highlights the role of information quality in learning (Craik & Lockhart, 1972). It distinguishes two types of processing: *Shallow Processing* (Sanford & Graesser, 2006), often resulting in ephemeral memory retention due to its superficial engagement; and *Deep Processing* (Lockhart & Craik, 1990), which encourages an in-depth engagement, leading to the formation of durable and readily retrievable memories.

Methodology

Overview

To alleviate catastrophic forgetting in continual learning, we introduce a novel replay-based method named **CO**gnitive **RE**play (CORE). Figure 3 provides a comprehensive illustration of CORE. Initially, the model is trained on \mathcal{D}_tCL , which contains the current task’s data and data from the replay buffer. After training, an evaluation is conducted to

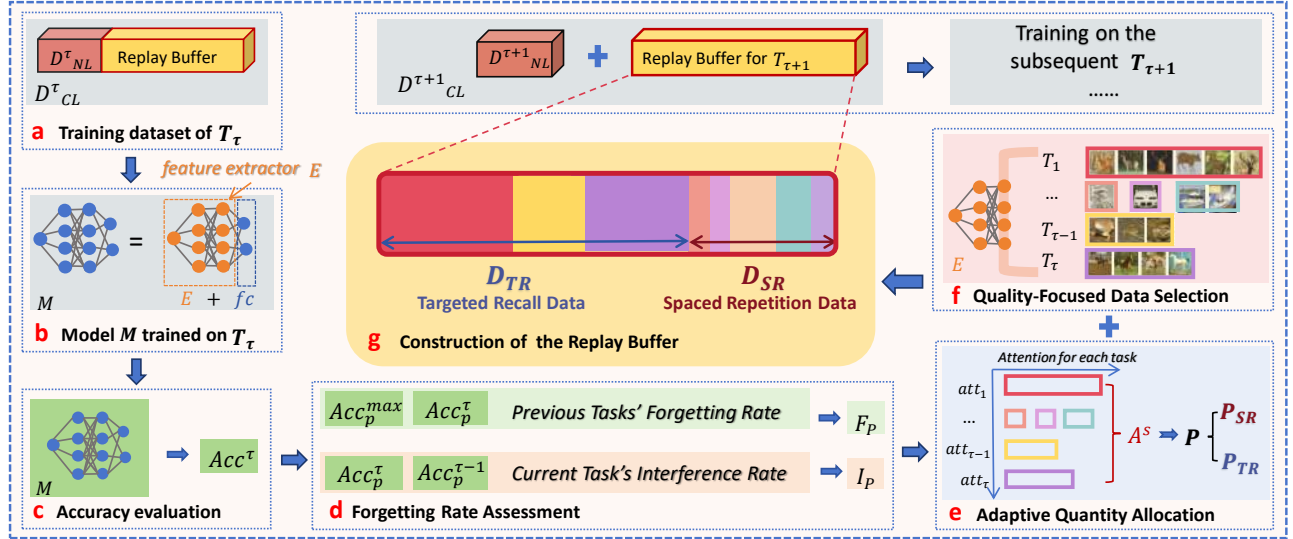


Figure 3: Pipeline of CORE. Box.a-g illustrate the process of strategically constructing the replay buffer after the introduction of current task \mathcal{T}_τ to prepare for the subsequent task $\mathcal{T}_{\tau+1}$.

gauge the forgetting rate. Utilizing our two principal strategies, we then dynamically allocate buffer space and select high-quality samples for the replay buffer, which is then used for the subsequent task $\mathcal{T}_{\tau+1}$'s training.

Cognitive Overload and Interference in Models

Forgetting in models parallels human cognitive processes, notably cognitive overload and interference. Cognitive overload in models occurs when the fixed size of the model, including its architecture and parameter capacity, becomes saturated with new information, hindering its ability to effectively retain and process previously learned information. This resembles human cognitive constraints under substantial information loads. Interference, inherent in the neural network's interconnected architecture, means even minor parameter adjustments can drastically impact task recall (Yosinski, Clune, Nguyen, Fuchs, & Lipson, 2015). This phenomenon reflects interference-based forgetting in human memory, where new learning can obstruct the retention of established knowledge.

Forgetting Rate Assessment

Our methodology strategically constructs the replay buffer, guided by the specific forgetting rates of different tasks. To quantify forgetting in previous tasks \mathcal{P} , we consider both cognitive overload and interference (Figure 3 box.d). For the current task \mathcal{T}_τ , which is newly trained and has not yet exhibited signs of forgetting, we estimate its potential for future forgetting by examining its interference impact on previous tasks.

Previous Tasks' Forgetting Rate After completing the training of the current task \mathcal{T}_τ (Figure 3 box.b), we perform an accuracy evaluation for each task in the set of previous tasks \mathcal{P} . All accuracy values during the introduction round of \mathcal{T}_τ are recorded in the set Acc^τ (Figure 3 box.c). To compute

the forgetting rate \mathcal{F}_p for previous tasks (Wang et al., 2023), we compare the historical accuracy of each task with its current performance after training on \mathcal{T}_τ . The forgetting rate f_p for each task is determined as follows:

$$F_p = \left\{ f_p = \max_{i \in \{1, \dots, \tau-1\}} Acc_p^i - Acc_p^\tau \mid p \in \mathcal{P} \right\} \quad (1)$$

Here, f_p quantifies the extent of performance degradation for each previous task p , which effectively reveals the impact of both cognitive overload and interference on forgetting.

Current Task's Interference Rate Since the cognitive overload of the recently trained task \mathcal{T}_τ has not yet occurred, we indirectly assess its future forgetting potential by examining how it affects performance on previous tasks \mathcal{P} . This assessment provides insight into how \mathcal{T}_τ might be influenced by the data of previous tasks stored in the buffer during the training of $\mathcal{T}_{\tau+1}$, which helps us anticipate and address the susceptibility of \mathcal{T}_τ to interference-induced forgetting. The interference rate I_p is calculated as follows:

$$I_p = \left\{ i_p = \frac{e^{Acc_p^{\tau-1} - Acc_p^\tau}}{\sum_{p' \in \mathcal{P}} e^{Acc_{p'}^{\tau-1} - Acc_{p'}^\tau}} \mid p \in \mathcal{P} \right\} \quad (2)$$

Here, i_p quantifies the interference effect on task p from training \mathcal{T}_τ , which is calculated by the exponential difference in task p 's accuracy before ($Acc_p^{\tau-1}$) and after (Acc_p^τ) the training of \mathcal{T}_τ . This method emphasizes the variation in performance, and the normalization across all tasks in \mathcal{P} ensures a uniform assessment of interference impacts.

Adaptive Quantity Allocation

Adaptive Quantity Allocation (AQA) dynamically manages replay buffer space by calculating each task's attention needs,

\mathcal{A} , based on its forgetting and interference rates. Previous tasks \mathcal{P} are then categorized into two subsets, \mathcal{P}_{SR} for spaced repetition and \mathcal{P}_{TR} for targeted recall, guided by the softmax-transformed attention scores \mathcal{A}^s (Figure 3 box.e).

Attention Calculation For previous tasks, the attention is quantified based on their forgetting rate $\mathcal{F}_{\mathcal{P}}$ as follows:

$$\mathcal{A} = \mathcal{A} \cup \left\{ att_p = -\log(1 - f_p) \mid p \in \mathcal{P}, f_p \in \mathcal{F}_{\mathcal{P}} \right\} \quad (3)$$

Here, att_p represents the attention allocated to task p according to its forgetting rate f_p . This ensures that tasks with more significant forgetting get more attention during review.

For the current task \mathcal{T}_{τ} , the attention is calculated based on its interference with previous tasks:

$$att_{\tau} = -\log\left(1 - \sum_i^{\tau-1} f_p \cdot i_p\right), \quad f_p \in \mathcal{F}_{\mathcal{P}}, i_p \in I_{\mathcal{P}} \quad (4)$$

In this equation, att_{τ} integrates the forgetting rate $\mathcal{F}_{\mathcal{P}}$ and the interference rate $I_{\mathcal{P}}$ over previous tasks \mathcal{P} . This integrated measure assigns appropriate attention to \mathcal{T}_{τ} to prevent potential forgetting. Furthermore, att_{τ} is subsequently included in the attention set \mathcal{A} for comprehensive buffer allocation.

Attention-based Allocation After calculating the attention \mathcal{A} , we employ the softmax transformation to normalize \mathcal{A} into a probabilistic distribution \mathcal{A}^s :

$$\mathcal{A}^s = \left\{ att_i^s = \frac{e^{att_i}}{\sum e^{att_i}} \mid att_i \in \mathcal{A} \right\} \quad (5)$$

Then, we incorporate two key cognitive strategies, Targeted Recall (TR) and Spaced Repetition (SR), to enhance the efficiency of the replay buffer. To execute this, we categorize all previous tasks \mathcal{P}^1 into two subsets $\mathcal{P}_{SR}, \mathcal{P}_{TR}$. This categorization is based on a threshold $\frac{1}{\lambda|\mathcal{P}|}$:

$$p \in \begin{cases} \mathcal{P}_{SR}, & \text{if } att_p^s \leq \frac{1}{\lambda|\mathcal{P}|} \\ \mathcal{P}_{TR}, & \text{if } att_p^s > \frac{1}{\lambda|\mathcal{P}|} \end{cases} \quad (6)$$

Here $|\mathcal{P}|$ represents the number of tasks in the set \mathcal{P} . \mathcal{P}_{TR} , designated for Targeted Recall, contains tasks that require increased attention; while \mathcal{P}_{SR} , designated for Spaced Repetition, contains tasks that require minimal attention.

To prevent tasks in \mathcal{P}_{SR} from receiving ineffectively low att_p^s , we set a minimum attention threshold:

$$\hat{att}_p = \frac{1}{\lambda|\mathcal{P}|}, \quad \text{if } p \in \mathcal{P}_{SR} \quad (7)$$

After ensuring the minimum allocation of attention to \mathcal{P}_{SR} , we focus on proportionally allocating the remaining attention to \mathcal{P}_{TR} based on Targeted Recall (TR):

¹ \mathcal{T}_{τ} is now contained in previous tasks set \mathcal{P} .

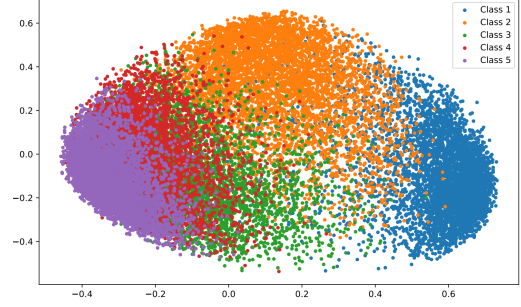


Figure 4: A dimension-reduced feature space for five classes.

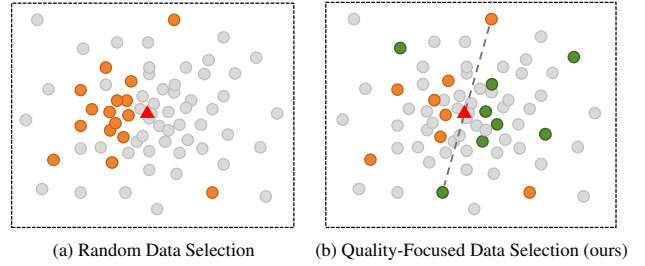


Figure 5: Comparison of data selection methods. The red triangle: the feature center of the class; Orange dots: samples randomly selected; Green dots: samples chosen to balance the sample distribution, which maintain a symmetrical relation (blue dotted line) around the feature center.

$$\hat{att}_p = \left(1 - \frac{|\mathcal{P}_{SR}|}{\lambda|\mathcal{P}|\right) \cdot \frac{att_p^s}{\sum_{p' \in \mathcal{P}_{TR}} att_{p'}^s}, \quad \text{if } p \in \mathcal{P}_{TR} \quad (8)$$

In this equation, $|\mathcal{P}_{SR}|$ represents the number of tasks in \mathcal{P}_{SR} . This allocation is meticulously designed to effectively implement the Targeted Recall strategy, guaranteeing that tasks more prone to forgetting are allocated with the requisite attention for efficient review and memory reinforcement.

Quality-Focused Data Selection

In addition to our allocations for \mathcal{D}_{TR} and \mathcal{D}_{SR} , we place a strong emphasis on the quality of the data in the replay buffer. Inspired by the latent space utilization (Zeiler & Fergus, 2014), our approach employs a segment of the model as a feature extractor \mathcal{E} , which maps data into its latent feature space (Figure 3 box.f). The latent feature space (Figure 4) aligns with the model’s cognitive representation, which enables a more effective data selection. Additionally, this strategy leverages the model’s existing capabilities without additional computational burden.

We introduce an algorithm to select more representative data for each class by repeatedly performing two steps: 1) randomly selecting a sample from the dataset of the specific class; 2) selecting a subsequent sample that minimizes the distance between the average feature representation of all pre-

viously selected samples and the average feature representation of the entire class in the feature space. A detailed procedure is provided in Algorithm 1.

Although previous random selection indeed offers equal chances for all samples, it leads to a concentration of samples around certain areas while leaving others sparse (Figure 5a), which may not adequately represent the entire feature distribution. In contrast, our approach (Figure 5b) effectively counters the uneven sample distribution, which guarantees a more uniform coverage of the feature space and a comprehensive representation of the dataset’s key features.

Algorithm 1 Feature Based Data Selection

```

input  $num$            // Number of data to be selected
require  $\mathcal{D}$            // Dataset
require  $\mathcal{E} : \mathcal{X} \rightarrow \mathcal{R}^d$  // Feature Extractor
1:  $\mu \leftarrow \frac{1}{|\mathcal{D}|} \sum_{p \in \mathcal{D}} \mathcal{E}(p)$ 
2:  $SelectData \leftarrow \emptyset$ 
3: for  $i = 1, 3, 5, \dots, num$  do
4:    $p_i \leftarrow \text{RandomlySelect}(\mathcal{D})$ 
5:    $p_{i+1} \leftarrow \arg \min_{x \in \mathcal{D}} \left\| \frac{1}{i+1} \cdot \left( \sum_{j=1}^i \mathcal{E}(p_j) + \mathcal{E}(x) \right) - \mu \right\|$ 
6:    $SelectData \leftarrow SelectData \cup \{p_i, p_{i+1}\}$ 
7: end for
output  $SelectData$ 

```

Experiments and Results

Experimental Settings

We conducted extensive experiments on split-MNIST (Zenke, Poole, & Ganguli, 2017), split-CIFAR-10, and split-CIFAR-100 (Krizhevsky, Hinton, et al., 2009) datasets to evaluate our method. For a rigorous test of continual learning, we divide each dataset into 10 separate tasks and set the replay buffer size to 500 for all replay-based methods.

To provide a comprehensive evaluation of our approach, we utilize the standard metric of average task accuracy (De Lange et al., 2022). Moreover, we consider a new metric, the lowest task accuracy across all tasks, which is crucial for identifying a balanced performance across all tasks, rather than excelling in some while neglecting others.

For a fair comparison with baselines, our experiments are based on the framework provided by (van de Ven, Tuytelaars, & Tolias, 2022), which integrates a range of baseline methodologies in continual learning. We compare our approach with a range of methodologies in continual learning for a comprehensive assessment. This includes replay-based methods such as iCaRL (Rebuffi et al., 2017), ER (Rolnick et al., 2019), and DGR (Shin, Lee, Kim, & Kim, 2017), which also utilize a replay buffer like ours. To offer a broader perspective on CORE’s effectiveness, we compare our approach with baselines based on regularization such as LwF (Li & Hoiem, 2017) and EWC (Huszár, 2018). Moreover, we include Joint (De Lange et al., 2021) in our comparison to serve

as an ideal upper bound that all tasks are trained concurrently.

Evaluation Across Baselines

Table 1 shows the final average accuracy Acc_{avg} of the model, along with the lowest accuracy Acc_{min} across all tasks. Across three datasets, we found that regularization methods like LwF and EWC struggle to prevent the forgetting of older tasks, mainly due to the lack of available data. This highlights the importance of the replay buffer in knowledge retention. Therefore, our analysis focuses mainly on replay-based baselines that employ the replay buffer like ours.

Experiments on split-MNIST We find that though replay-based baselines such as iCaRL and DGR have relatively high average accuracy values, their lowest accuracy values are extremely low. The significant difference between DGR’s lowest accuracy of 63.85% and its average accuracy of 85.46% reveals substantial imbalances in task-specific performance. This underscores the importance of the metric Acc_{min} for a holistic evaluation of performance. In contrast, CORE outperforms the baselines with an average accuracy of 94.52%, closely approaching the upper-bound of 97.81%. Moreover, the lowest accuracy of all tasks still maintains 89.94%, largely proving the effectiveness of our approach.

Experiments on split-CIFAR10 On a more complex dataset, our approach still achieves an impressive average accuracy of 37.95%, significantly higher than iCaRL’s 32.62%. On this dataset, ER’s lowest accuracy is only 10.00%, akin to random guessing. Notably, DGR’s lowest accuracy plunges to 0.00%, indicating severe forgetting in the specific task. Additionally, iCaRL exhibits a staggering disparity of 26.82% between its average accuracy (32.62%) and its lowest task accuracy (5.80%). In contrast, CORE maintains a consistent performance, with the lowest task accuracy of 16.30%, signifying a more balanced approach to knowledge retention.

Experiments on split-CIFAR100 It is worth noting that on the larger split-CIFAR100 dataset, where each task contains more classes (i.e. more inference), a more pronounced forgetting occurs. This original data distribution sort of alleviates the drawbacks of iCaRL and ER’s less strategic buffer allocation while undermines the effectiveness of our AQA strategy. Nevertheless, our proposed QFDS ensures the data samples in the replay buffer are representative, thus guaranteeing CORE’s superior performance on split-CIFAR100.

Analysis of Progressive Learning Scenarios

For a comprehensive assessment of the effectiveness of our method, it is insufficient to only consider the final accuracy. It is crucial to understand how individual tasks maintain their performance as new tasks are introduced in the process. This section aims to show this dynamic by observing the accuracy of the first task \mathcal{T}_1 as it trains on nine subsequent tasks.

Figure 6 visualizes the accuracy trajectory of \mathcal{T}_1 from the introduction of the fifth task. Focusing on this stage is crucial as it marks a point where differences become more apparent.

Table 1: The average and minimum accuracy (%) of 10 tasks on three datasets. Results are reported as an average of 3 runs. ‘ Acc_{avg} ’ refers to the average accuracy, and ‘ Acc_{min} ’ refers to the lowest accuracy of tasks. The best results are in bold.

Dataset	Metrics	iCaRL	ER	DGR	LwF	EWC	Ours	Upper Bound
split-MNIST	Acc_{avg}	88.00	92.39	85.46	10.00	10.00	94.52	97.81
	Acc_{min}	78.97	87.47	63.85	0.00	0.00	89.94	95.64
split-CIFAR10	Acc_{avg}	32.62	29.74	13.00	10.00	10.00	37.95	65.03
	Acc_{min}	5.80	10.00	0.00	0.00	0.00	16.30	44.30
split-CIFAR100	Acc_{avg}	24.02	25.41	6.31	10.35	6.14	27.24	37.95
	Acc_{min}	15.50	18.70	0.00	0.10	0.00	20.40	32.60

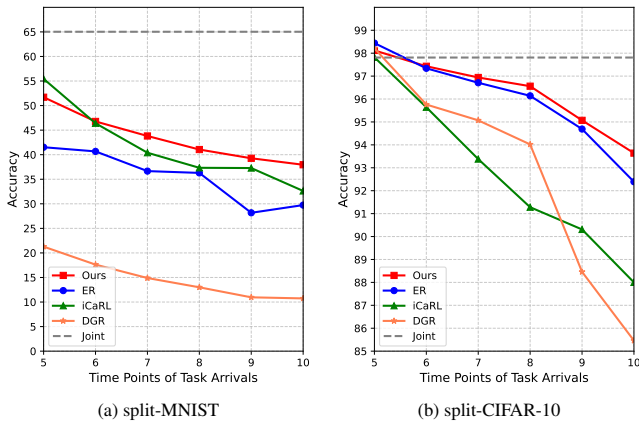


Figure 6: Accuracy trajectories for the first task from the fifth task onward on split-MNIST and split-CIFAR-10.

In the later stages, our method consistently outperforms the baselines, illustrating its robustness in maintaining previous knowledge. Notably, Figure 6b indicates a more stabilized trend in the accuracy decline for our method, suggesting the stability of our approach over time. This trend indicates that our approach can effectively manage the influx of multiple tasks without experiencing significant performance drops.

Grid Search of Parameter λ

In our methodology, the parameter λ plays a vital role in setting the minimum attention threshold $\frac{1}{\lambda|\mathcal{P}|}$, ensuring effective review of each task in the replay buffer. To determine the optimal value of λ , we conducted a grid search, as outlined in Table 2. The grid search identified $\lambda = 2$ as the optimal value, effectively balancing review efficacy.

Ablation Study

The ablation study reveals CORE’s effectiveness, with significant performance reductions when either Adaptive Quantity Allocation (AQA) or Quality-Focused Data Selection (QFDS) is omitted. AQA’s absence leads to a 5.17% decrease in average accuracy and a 5.30% drop in minimum task ac-

Table 2: Accuracy values for different λ .

λ	split-MNIST	split-CIFAR10	split-CIFAR100
1	92.39	29.74	25.41
2	93.64	37.95	27.38
3	93.62	35.55	27.24
4	93.04	36.76	26.45
5	92.57	34.14	27.19

curacy, while omitting QFDS results in a 6.43% reduction in average accuracy and a 3.90% decrease in minimum task accuracy. This indicates AQA’s crucial role in boosting minimum task accuracy and QFDS’s impact on average accuracy, with the greatest performance decline observed when both are excluded, highlighting their combined importance in combating catastrophic forgetting.

Table 3: Ablation study of CORE on split-CIFAR-10.

Method	Acc_{avg}	Acc_{min}
CORE	37.95	16.30
CORE w/o AQA	32.78	11.00
CORE w/o QFDS	31.52	12.40
CORE w/o QFDS+AQA	29.74	10.00

Conclusion

In conclusion, CORE mimics human memory mechanism to optimize the utilization of the replay buffer in continual learning, emphasizing both the quantity and quality of data samples within the buffer. Experiment results confirm our approach’s superiority, as it not only achieves the highest average task accuracy but also outperforms on the tasks that are most susceptible to forgetting. This advancement establishes CORE as a new benchmark for replay-based paradigms and highlights the benefits of integrating human cognitive strategies into machine learning models.

Acknowledgments

This research was supported in part by the National Natural Science Foundation of China under grant No. 62172303.

References

- Carpenter, S. K., Cepeda, N. J., Rohrer, D., Kang, S. H. K., & Pashler, H. (2012, Sep). Using spacing to enhance diverse forms of learning: Review of recent research and implications for instruction. *Educational Psychology Review*, 369–378. doi: 10.1007/s10648-012-9205-z
- Chen, Z., & Liu, B. (2020, Jan). Lifelong machine learning. In *Transfer learning* (p. 196–208). doi: 10.1017/9781139061773.016
- Craik, F. I., & Lockhart, R. S. (1972, Dec). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior*, 671–684. doi: 10.1016/s0022-5371(72)80001-x
- De Lange, M., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., ... Tuytelaars, T. (2021). A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern analysis and machine intelligence*, 44(7), 3366–3385.
- De Lange, M., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., ... Tuytelaars, T. (2022). A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(7), 3366–3385. doi: 10.1109/TPAMI.2021.3057446
- Flesch, T., Juechems, K., Dumbalska, T., Saxe, A., & Summerfield, C. (2022). Orthogonal representations for robust context-dependent task performance in brains and neural networks. *Neuron*, 110(7), 1258–1270.
- French, R. (1999, Apr). Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 128–135. doi: 10.1016/s1364-6613(99)01294-2
- French, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4), 128–135.
- Huszár, F. (2018). Note on the quadratic penalties in elastic weight consolidation. *Proceedings of the National Academy of Sciences*, 115(11), E2496–E2497.
- Krizhevsky, A., Hinton, G., et al. (2009). Learning multiple layers of features from tiny images.
- Li, Z., & Hoiem, D. (2017). Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12), 2935–2947.
- Lockhart, R. S., & Craik, F. I. (1990). Levels of processing: A retrospective commentary on a framework for memory research. *Canadian Journal of Psychology/Revue canadienne de psychologie*, 44(1), 87.
- Masse, N. Y., Grant, G. D., & Freedman, D. J. (2018). Alleviating catastrophic forgetting using context-dependent gating and synaptic stabilization. *Proceedings of the National Academy of Sciences*, 115(44), E10467–E10475.
- Mayo, D., Scott, T. R., Ren, M., Elsayed, G., Hermann, K., Jones, M., & Mozer, M. (2023). Multitask learning via interleaving: A neural network investigation. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 45).
- McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation* (Vol. 24, pp. 109–165). Elsevier.
- McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5, 115–133.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M. (2013, Dec). Playing atari with deep reinforcement learning. *arXiv: Learning, arXiv: Learning*.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... others (2015). Human-level control through deep reinforcement learning. *nature*, 518(7540), 529–533.
- Rebuffi, S.-A., Kolesnikov, A., Sperl, G., & Lampert, C. H. (2017). icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 2001–2010).
- Rolnick, D., Ahuja, A., Schwarz, J., Lillicrap, T., & Wayne, G. (2019). Experience replay for continual learning. *Advances in Neural Information Processing Systems*, 32.
- Rusu, A. A., Rabinowitz, N. C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., ... Hadsell, R. (2016). Progressive neural networks. *arXiv preprint arXiv:1606.04671*.
- Sadeh, T., Ozubko, J. D., Winocur, G., & Moscovitch, M. (2016). Forgetting patterns differentiate between two forms of memory representation. *Psychological science*, 27(6), 810–820.
- Sanford, A. J., & Graesser, A. C. (2006). Shallow processing and underspecification. *Discourse Processes*, 42(2), 99–108.
- Schmidgall, S., Achterberg, J., Miconi, T., Kirsch, L., Ziaei, R., Hajiseyedrazi, S., & Eshraghian, J. (2023). Brain-inspired learning in artificial neural networks: a review. *arXiv preprint arXiv:2305.11252*.
- Shin, H., Lee, J. K., Kim, J., & Kim, J. (2017). Continual learning with deep generative replay. *Advances in neural information processing systems*, 30.
- Sze, V., Chen, Y.-H., Yang, T.-J., & Emer, J. S. (2017). Efficient processing of deep neural networks: A tutorial and survey. *Proceedings of the IEEE*, 105(12), 2295–2329.
- Tiwari, R., Killamsetty, K., Iyer, R., & Shenoy, P. (2022). Gcr: Gradient coreset based replay buffer selection for continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 99–108).
- Van de Ven, G. M., & Tolia, A. S. (2019). Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*.
- van de Ven, G. M., Tuytelaars, T., & Tolia, A. S. (2022). Three types of incremental learning. *Nature Machine Intelligence*, 4, 1185–1197.

- Wang, L., Zhang, X., Su, H., & Zhu, J. (2023). A comprehensive survey of continual learning: Theory, method and application. *arXiv preprint arXiv:2302.00487*.
- Yosinski, J., Clune, J., Nguyen, A., Fuchs, T., & Lipson, H. (2015, Jun). Understanding neural networks through deep visualization. *arXiv: Computer Vision and Pattern Recognition, arXiv: Computer Vision and Pattern Recognition*.
- Zeiler, M. D., & Fergus, R. (2014). Visualizing and understanding convolutional networks. In *Computer vision—eccv 2014: 13th european conference, zurich, switzerland, september 6-12, 2014, proceedings, part i 13* (pp. 818–833).
- Zenke, F., Poole, B., & Ganguli, S. (2017). Continual learning through synaptic intelligence. In *International conference on machine learning* (pp. 3987–3995).