Does Replay Suffice for Online Continual Learning in Spiking Networks?

Raghav Patel* **(raghav.patel@my.utsa.edu)**, **Nicholas Soures*** **(nicholas.soures@utsa.edu)**,
Dhireesha Kudithipudi (dk@utsa.edu)
Neuromorphic AI Lab, One UTSA Circle
San Antonio, TX 78249 USA
* denotes equal contribution

Abstract

Continual learning models must learn sequentially from a changing data distribution, while accumulating previously learned knowledge without forgetting. Replay and parameter regularization are two prominent mechanisms that have shown promise in deep learning models but have only been explored for spiking neural networks in few works. In this work, we study the application of replay to domain-incremental learning in spiking neural networks and see if metaplasticity and synaptic consolidation can help efficient continual learning. We demonstrate that simple replay schemes can achieve state-of-the-art performance and that the incorporation of metaplasticity can increase performance when using low buffer sizes. This approach can serve as a baseline for efficient continual learning that can reduce memory and training overhead of replay.

Keywords: Spiking neural network; continual learning; replay; regularization.

Introduction & Motivation

Continual learning, where the model incrementally learns from a non-stationary stream of data, faces the issue of catastrophic interference. Several methods address this issue, such as task-specific architectures, regularization, replay/rehearsal, and template-based classification (Ven, Tuytelaars, & Tolias, 2022). Replay in particular is the recall and training of previous experiences and neuroscience experiments show that sleep can be related to memory consolidation (Wilson & McNaughton, 1994). This is an effective technique for preserving past knowledge while learning new information between tasks, either by storing past data in a memory buffer or a generative model to generate synthetic samples from prior data to replay.

Replay alone has been shown to outperform other approaches in both domain-incremental and class-incremental scenarios (Ven et al., 2022). However, this performance requires additional compute and memory overhead, requiring offline training phases and increasing the number of training iterations. Unlike replay, regularization methods rely on local information to preserve knowledge. Although these approaches aid in continual learning in different ways, there have been limited studies exploring the potential benefits of their integration, particularly in spiking networks (Proietti, Ragno, & Capobianco, 2023). In this work, we investigate a two-stage process for replay with regularization in spiking networks.

Experimental Set-up

The network architecture is described in Fig. 1 and learning is event-driven random backpropagation (eRBP) (Neftci, Augustine, Paul, & Detorakis, 2017). The replay buffer stores raw input samples in memory. The samples selected for replay will be random and will be the raw image data. The buffer update will be proportional to how many tasks have been trained so far. For example, after training on N tasks, there will be a replay buffer Q consisting of N sets of size $\frac{K}{N}$, where each set contains randomly selected samples to be replayed from the N^{th} task. The discarded samples will be randomly selected. The training methodology used here is a two-stage process. In the first stage, the network will learn from a given task and then update the replay buffer once the task has completed training. Next, the network enters a replay stage where the data is rehearsed from all prior tasks before moving onto subsequent tasks.

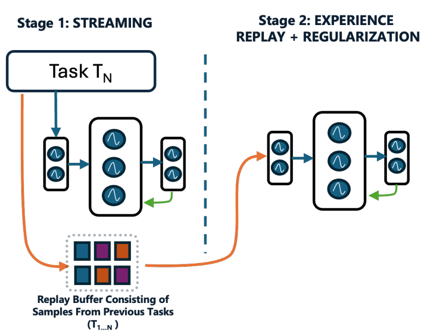

Figure 1 Proposed two stage continual learning model with experience replay combined with parameter regularization in a 3-layer spiking network. Local surrogate gradient learning on two part LIF neurons (network size: 784x200x2). Orange arrow in Stage 1 show buffer update. Orange arrow in Stage 2 show buffer complements training sequence. Green arrows show error feedback. Blue arrows show forward propagation of samples

There are two regularization techniques explored for integration with replay. First is metaplasticity, the plasticity of plasticity, which regulates the ability of a synapse to change its synaptic strength. Specifically, plasticity is computed $f(m,w) = e^{-|mw|}$, where m is a metaplastic parameter and w is the weight. The parameter m is incrementally updated to estimate a neuron's importance. The second regularization technique is consolidation which introduces a

Buffer Size (K)	500	1000	2000	4000
Mean Accuracy (MA)	66.23%	72.41%	79.55%	85.03%
Std. Deviation	$\pm$ 0.66	$\pm$ 1.04	$\pm$ 1.67	$\pm$ 0.21

Table 1 Split-MNIST Mean Accuracy on testing set for varying replay buffer sizes.

slow moving synaptic component, w^{ref}, and balances learning and retention through decay such that whenever a post-synaptic neuron spikes the synapses are changed according to $\Delta w_{i,j}(t+1) = -\alpha(w_{ij}(t) - w_{ij}^{ref}(t))$ (Soures, Helfer, Daram, Pandit, & Kudithipudi, 2021). To integrate these techniques with replay we propose modified regularization during the replay, while the training proceeds as normal with regularization. The change in metaplastic parameter m will be scaled by c while the decay rate for consolidation is scaled by v. The loss function can be described as the following: $\mathcal{L}_{total}(\theta) = \mathcal{L}_{curr}(\theta) + \mathcal{L}_{replay}(\theta, c, v)$, where θ is the parameters to the network.

In this evaluation setting, the task identity is not provided during inference and does not need to be inferred, known as domain-incremental learning. Specifically, there are five tasks with two classes per task for each dataset. Samples are presented to the network in a streaming fashion (only one sample at a time and only one epoch), and after each task is learned by the network, the network is evaluated on test sets from all tasks (both seen and unseen). Performance is measured by the Mean Accuracy (MA) which is defined as $MA = \frac{1}{N}\sum_{t=1}^{N} R^{t,N}$, where $R^{t,N}$ is the accuracy of the task T^t after training on task T^N. We also use Forward Transfer, $FWT = \frac{1}{N-1}\sum_{K=1}^{N-1}\sum_{t=K+1}^{N}\frac{R^{t,k}-R^{t,k-1}}{N-K}$ which describes the average change in accuracy across the tasks $t > k$ after it has learned T^k. And finally, we use Backwards Transfer $BWT = \frac{2}{N(N-1)}\sum_{k=2}^{N}\sum_{t=1}^{k-1}(R^{t,k} - R^{t,k-1})$ which describes the average change in accuracy across tasks $t < k$ after it has learned T^k (Kudithipudi et al., 2023).

The baseline model uses only replay in between tasks. We evaluate the performance of the baseline model for various buffer sizes ($K = \{4000, 2000, 1000, 500\}$). To study the integration of replay with regularization, we use a buffer size of $K = 2000$, similar to (Rebuffi, Kolesnikov, Sperl, & Lampert, 2017). In these experiments the regularization techniques are added to the model during training and scaled during replay stages.

Results and Discussion

Replay alone showed a significant increase in mean accuracy, as expected. By introducing regularization with replay, specifically metaplasticity shown in Table 2, we observe $\approx 10\%$ increase in mean accuracy with a buffer of 500 samples. An interesting observation is that the model shows higher forward transfer (FWT) which shows how well the network retains and transfers knowledge across tasks. We also extended the test to FMNIST dataset which only showed a marginal increase in performance (0.20%). The important conclusion from this result is that metaplasticity can increase performance with lower buffer sizes. For example, replay with metaplasticity was able to achieve 84.34% using 2000 samples, which was significantly higher than the 79.55% baseline and close to the 85.03% achieved with double the buffer size. However, it is also important to note that as the buffer size grows, the benefit of metaplasticity during replay decreases, as shown by the $\approx 1.5\%$ increase in performance with a buffer of 2000 (Table 3). In contrast to metaplasticity, consolidation did not help performance under any circumstances and does not integrate well with replay.

	MA		FWT		BWT	
c	F-MNIST	MNIST	F-MNIST	MNIST	F-MNIST	MNIST
0	88.26%	65.03%	12.66	0.2	-4.89	-15.03
0.1	87.23%	73.66%	11.91	0.62	-5.26	-9.84
0.5	87.97%	75.31%	12.24	1.21	-5.05	-8.54
1	88.46%	74.98%	12.58	1.02	-4.88	-9.36
v	F-MNIST	MNIST	F-MNIST	MNIST	F-MNIST	MNIST
0	80.02%	74.80%	12.57	4.26	-9.43	-11.17
0.1	83.89%	65.59%	10.06	1.54	-7.38	-16.04
0.5	80.25%	68.02%	13.11	3.30	-8.69	-14.66
1	7973%	65.85%	14.06	3.96	-9.31	-15.27

Table 2 Mean Accuracy (MA), Forward Transfer (FWT), Backward Transfer (BWT) for split-FMNIST (Xiao et al., 2017) and split-MNIST (Lecun et al., 1998) testing sets. Buffer size = 500 samples and v varies the amount of decay in consolidation, while c varies the amount of metaplastic synaptic weight change.

c	0	0.1	0.5	1.0
K=2000	82.84%	72.02%	84.34%	83.50%
K=500	65.03%	73.66%	75.31%	74.98%

Table 3 split-MNIST performance with varying metaplasticity (c) and buffer size (K).

This preliminary work investigates regularization during replay in spiking networks. We show that there is a possibility of saving memory space by using a smaller buffer size with metaplasticity, to increase mean accuracy. Further analysis is needed by including several trials with a wider range of metaplasticity parameter c. Although metaplasticity improvement on a small buffer size can be observed in split-MNIST, the same amount of improvement is not shown in split-FMNIST; the metaplastic parameter could be sensitive to the training data distribution. Future work can include combining both metaplasticity and consolidation with replay, since the two regularization techniques together have been shown to support continual learning in spiking networks (Soures et al., 2021).

Acknowledgments

This work is partially supported by the NSF EFRI BRAID grant #2317706.

References

Kudithipudi, D., Daram, A., Zyarah, A. M., Zohora, F. T., Aimone, J. B., Yanguas-Gil, A., . . . Epstein, B. (2023). Design principles for lifelong learning ai accelerators. *Nature Electronics*, *6*, 807-822.

Neftci, E. O., Augustine, C., Paul, S., & Detorakis, G. (2017). Event-driven random back-propagation: Enabling neuromorphic deep learning machines. *Frontiers in Neuroscience*, *11*.

Proietti, M., Ragno, A., & Capobianco, R. (2023). Memory replay for continual learning with spiking neural networks. In *2023 ieee 33rd international workshop on machine learning for signal processing (mlsp)* (p. 1-6).

Rebuffi, S.-A., Kolesnikov, A., Sperl, G., & Lampert, C. H. (2017). iCaRL: incremental classifier and representation learning. In *Cvpr.*

Soures, N., Helfer, P., Daram, A., Pandit, T., & Kudithipudi, D. (2021). Tacos: Task agnostic continual learning in spiking neural networks. *ICML*.

Ven, G. M. V. D., Tuytelaars, T., & Tolias, A. S. (2022). Three types of incremental learning. *Nat Mach Intell*, *4*, 1185–1197.

Wilson, M. A., & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, *265*, 676-679.