

Probabilistic metaplasticity for continual learning with memristors

Fatima Tuz Zohora^{1,*}, Vedant Karia¹, Nicholas Soures¹, and Dhireesha Kudithipudi¹

¹University of Texas at San Antonio, Department of Electrical and Computer Engineering, San Antonio, TX 78249, USA

*fatimatuz.zohora@my.utsa.edu

ABSTRACT

Edge devices operating in dynamic environments critically need the ability to continually learn without catastrophic forgetting. The strict resource constraints in these devices pose a major challenge to achieve this, as continual learning entails memory and computational overhead. Crossbar architectures using memristor devices offer energy efficiency through compute-in-memory and hold promise to address this issue. However, memristors often exhibit low precision and high variability in conductance modulation, rendering them unsuitable for continual learning solutions that require precise modulation of weight magnitude for consolidation. Current approaches fall short to address this challenge directly and rely on auxiliary high-precision memory, leading to frequent memory access, high memory overhead, and energy dissipation. In this research, we propose probabilistic metaplasticity, which consolidates weights by modulating their update *probability* rather than magnitude. The proposed mechanism eliminates high-precision modification to weight magnitudes and, consequently, the need for auxiliary high-precision memory. We demonstrate the efficacy of the proposed mechanism by integrating probabilistic metaplasticity into a spiking network trained on an error threshold with low-precision memristor weights. Evaluations of two continual learning benchmarks show that probabilistic metaplasticity achieves state-of-the-art performance while consuming $\sim 67\%$ lower memory for additional parameters and up to two orders of magnitude lower energy during parameter updates compared to an auxiliary memory-based solution. The proposed model shows potential for energy-efficient continual learning with low-precision emerging devices.

Introduction

Autonomous edge devices constitute a major application domain of artificial intelligence—ranging from indoor robots to autonomous drones/vehicles¹. These devices often need on-device training for real-time adaptation to dynamic environments². However, training neural networks requires energy-intensive data movement and memory accesses, which can strain the energy budget of these resource-constrained platforms³. The challenge exacerbates in real-world scenarios where data is non-stationary and does not maintain independent and identical distribution. Conventional artificial intelligence fails to learn effectively in this scenario and catastrophically forgets previously learned information⁴. The ability to learn and adapt to non-stationary data streams requires continual learning mechanisms, which entail additional memory and computational overhead, further straining the strict resource budget on edge devices⁵. The constraints in autonomous edge applications present a unique set of challenges, which lead to certain desirable criteria for continual learning solutions. First, the continual learning solution should be task-agnostic, i.e., it should be effective with no supervision or task-boundary awareness for autonomous applications. Second, edge devices have limited on-chip memory, which prohibits numerous iterative training over a large dataset. Hence, these devices should learn in an online scenario where data are encountered only once¹. Finally, the memory and computational overhead of the continual learning mechanism should be minimized considering the strict energy constraints of edge devices⁶.

Recent research primarily focuses on alleviating catastrophic forgetting to achieve continual learning while overlooking the above-mentioned criteria. In our research, we strive to achieve continual learning in the online task-agnostic scenario while minimizing memory and computational overhead. We address this multifaceted challenge by taking inspiration from the biological brain, which shows the ability to learn continually throughout its lifetime with remarkable energy efficiency⁷. Neuromorphic or brain-inspired approaches can help in two ways. First, we can take inspiration from the neural mechanisms deemed to help memory consolidation in the brain to address catastrophic forgetting⁸. One such mechanism is metaplasticity or plasticity of plasticity⁹, whereby synapses modify their ability to change based on their history of activity. Recent research in continual learning abstracted this concept by detecting important weights and restricting their update in magnitude by modulating weight gradients. Metaplasticity-inspired weight consolidation mechanisms have shown promising performance in addressing continual

learning in neural networks^{10–15}. Second, we can harness the computational and structural principles of the brain to optimize the energy efficiency of on-device continual learning. This can translate to adopting spiking models that emulate the computational primitives of the brain with sparse binary encoding. This feature replaces energy-intensive floating-point operations with low-power event-driven computations¹⁶. Structural primitives of the brain are emulated with architectures featuring co-located memory and processing with nanoscale emerging memristor devices¹⁷. In a crossbar arrangement, the conductance of these devices can emulate weights in a network and perform vector-matrix multiplication in a single timestep¹⁸. Such architectures reduce data movement, enable parallel computations and can lead to up to two orders of magnitude greater energy efficiency compared to conventional architectures¹⁹.

Although spiking models incorporated in memristor crossbar architectures are promising, they show non-ideal characteristics which can limit their suitability for continual learning mechanisms. The weight consolidation mechanism mentioned above is a prominent method of continual learning that relies on precise modulation of the weight magnitude. This is extremely challenging to achieve with memristors since these nanoscale devices show low precision ($\sim 3 - 6$ bits^{20,21}) and inherent variability during conductance modulation. Recent research has circumvented this issue by incorporating auxiliary high-precision memory (usually implemented with SRAM) alongside memristor crossbars²². Binary memristor devices have been used to accelerate forward pass in continual learning tasks while training with backpropagation is carried out in high-precision memory²³. Memristors with multi-level conductance have also been utilized for continual learning with metaplasticity, where auxiliary high-precision memory accumulates weight gradients, and memristor weights are updated when the accumulated gradient is equivalent to the device conductance resolution²⁴. These approaches have shown promising performance with low-precision memristor weights in continual learning benchmarks. However, these solutions rely on multi-epoch training and knowledge of task progression in order to either assign separate output neurons for different tasks or to modulate hyperparameters. These prerequisites are not suitable for autonomous applications. Moreover, high-precision memory with SRAMs lead to a large silicon footprint and idle standby leakage power (tens to hundreds of picoWatts per bit)²⁵. Finally, while crossbar architectures with auxiliary high-precision memory retain energy efficiency during inference, excessive memory access during training can potentially void the benefits of the compute-in-memory architecture.

In this research, we explore a novel weight consolidation mechanism suitable for low-precision memristor weights. Neuroscience literature has shown that binary metaplastic synapses characterized by probabilistic transitions help memory retention²⁶. Inspired by this, we propose the probabilistic metaplasticity model, where a weight's plasticity translates to its probability of update rather than its magnitude of update. The probability of a weight's update evolves during training such that important weights are assigned lower update probability to prevent catastrophic forgetting. Since this model does not need precise modulation in weight magnitude, the need for auxiliary high-precision memory is eliminated. Previous research explored probabilistic metaplasticity in the context of binary synapses in models of memory²⁶ and binarized networks that only support training a single-layer network²⁷. We generalize this concept to memristor weights with multi-level conductance incorporated in a multilayer spiking network trained based on an error threshold. We evaluate the proposed solution on two image detection continual learning benchmarks considering single and multi-memristor weights. Evaluations show that low-precision memristor weights with probabilistic metaplasticity can achieve performance equivalent to state-of-the-art networks with high-precision weights in the challenging online task-agnostic continual learning scenario. Compared to an equivalent solution with high-precision auxiliary memory, probabilistic metaplasticity reduces memory overhead for additional parameters by $\sim 67\%$ while requiring up to two orders of magnitude lower energy during the parameter update phase. The proposed model is promising for memory- and energy-efficient continual learning for autonomous edge applications.

Probabilistic metaplasticity with error threshold-based training

Metaplasticity, or plasticity of plasticity, refers to the modification of a synapse or weight's ability/tendency to change based on previous activity⁹. It is deemed a key mechanism for memory consolidation in the brain²⁸. Fusi et al. proposed a cascade metaplasticity model, where binary synapses contain an additional discrete metaplastic state variable²⁶. The metaplastic state of a synapse changes probabilistically in response to input stimuli. Synapses with higher metaplastic states are assigned a lower probability of change to signify greater stability.

Similar to the discrete metaplastic states in the cascade model, memristors can be programmed to a limited number of distinguishable conductance levels (Figure 1b). In our work, we consolidate memristor weights in a similar probabilistic manner. In a typical network, weights realized with memristors are updated deterministically by increasing or decreasing their conductances based on the error. In the proposed probabilistic metaplasticity model, the weights are updated probabilistically, where the important weights exhibit a lower probability. The

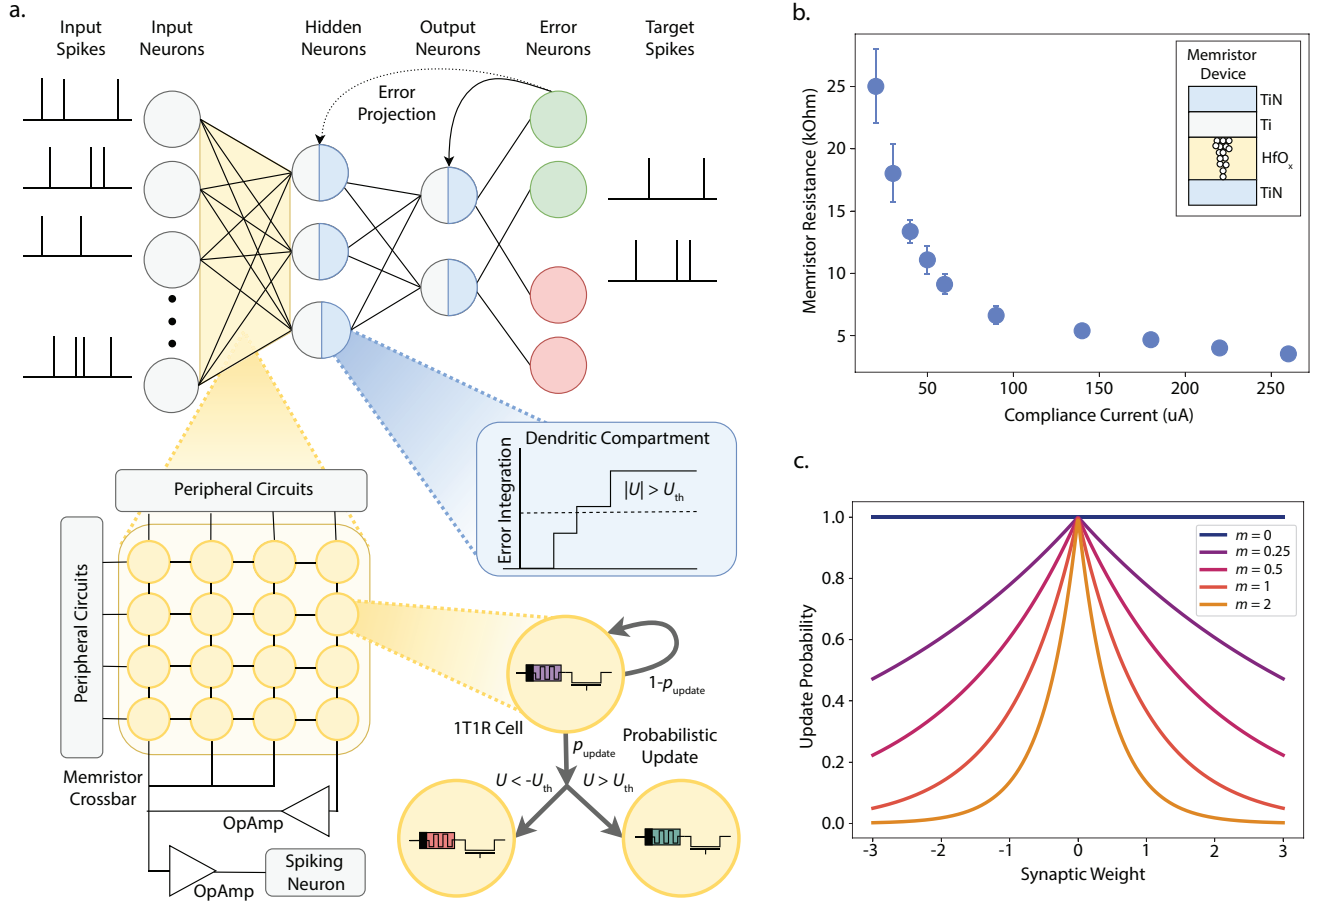


Figure 1. Probabilistic metaplasticity with error threshold-based training. **a.** Spiking network trained on error threshold where the network weights are realized with 1T1R memristor crossbar array. The dendritic compartments in the hidden and output layer neurons integrate the error. When the error reaches the threshold U_{th} , the memristor weights are updated to the next higher (negative error) or lower (positive error) conductance level with probability p_{update} . **b.** Mean and standard deviation of resistance levels versus programming compliance current for the 1T1R memristor device (shown in inset) adopted in this work²⁰. **c.** Update probability of weights for different values of the metaplastic coefficient m and the weights w . m is directly proportional to the activity level of adjacent neurons. We see that weights with highly active adjacent neurons (high m) and weight magnitude (high $|w|$) lead to low update probability.

importance of a weight is a function of its magnitude and the activity levels of adjacent neurons¹¹. High magnitude and connections to highly active pre- and post-synaptic neurons lead to a lower probability of updating or higher weight consolidation. This is captured in the metaplasticity function, which determines the update probability p_{update} of a weight w_{ij} connected between i^{th} pre-synaptic and j^{th} post-synaptic neurons -

$$p_{update} = e^{-|m_{ij}w_{ij}|}. \quad (1)$$

Here, m_{ij} is a metaplastic coefficient assigned to each weight. As the network learns, the metaplastic coefficient evolves based on the activity trace of the adjacent neurons as -

$$m_{ij}(t+1) = \begin{cases} m_{ij}(t) + \Delta m, & \text{if } X_j^{tr} \geq m_{th}^{pre} \& X_i^{tr} \geq m_{th}^{post} \\ m_{ij}(t) & \text{otherwise} \end{cases}. \quad (2)$$

We see that high activity in the adjacent neurons increases the metaplastic coefficient of a weight. Figure 1c shows the update probability for different values of weights and metaplastic coefficients. We see that weights with

Algorithm 1: Probabilistic metaplasticity with error threshold-based eRBP for continual learning in a spiking network. \mathcal{T} denotes a set of sequential tasks. S^{in} and S^{out} are the input and output spike train and W denotes weights realized with memristors. U is the error accumulated at the dendritic compartment of neurons. When $|U|$ crosses the error threshold U_{th} , the update probability p_{update} of eligible weights are calculated and compared with a random number to decide which weights should be updated. The memristor weights selected for update are programmed to the next higher or lower conductance level depending on the sign of the error. The function Program() refers to the operations required to update a memristor’s conductance. Details of the eRBP algorithm can be found in methods.

```

for  $task \in \mathcal{T}$  do
  for  $\{x^t, y^t\} \in \{X^t, Y^t\}$  do
    for  $t \in T_{\text{sim}}$  do
      Forward Pass:  $S^{\text{out}}(t) \leftarrow f(S^{\text{in}}(t), W(t))$ 
      Random Error Feedback:  $\tau_U \frac{\partial U}{\partial t} = ER_U$ 
      Update Neuron Trace:  $\frac{d}{dt} X_{\text{tr}} = -\frac{X_{\text{tr}}}{\tau_{\text{tr}}} + S$ 
      if  $|U_j| > U_{\text{th}}$  then
        for  $i \in \{X_i(t) == 1\}$  do
          if  $I_{\text{min}} < I_j < I_{\text{max}}$  then
             $\epsilon \sim \text{uniform}(0, 1)$ 
             $p_{\text{update}} = e^{-|m_{ij} w_{ij}|}$ 
            if  $p_{\text{update}} < \epsilon$  then
              Update memristor weights:  $W \leftarrow \text{Program}(R_{\text{mem}}, U)$ 
            end
             $U_j = 0$ 
          end
        end
      end
    end
    Update metaplastic coefficient:  $m \leftarrow m + \Delta m$ 
  end
end

```

higher magnitude and metaplastic coefficient exhibit lower update probability. Throughout training, a weight’s magnitude and metaplastic coefficient adapt in response to the activity of the network. This changes the update probability/plasticity of a weight, leading to a metaplastic effect to consolidate weights.

We incorporate probabilistic metaplasticity into a multi-layer spiking network trained based on an error threshold (shown in Figure 1a). The network is trained with event-driven random backpropagation (eRBP) where the network error is projected to neurons (details in Methods)²⁹. Each neuron includes a dendritic compartment that accumulates the error. In eRBP, a weight is considered eligible for update if the pre-synaptic neuron has an active spike and the post-synaptic current is within a range. When these criteria are met, the weights are updated in proportion to the accumulated error. However, low-precision weights cannot be updated with arbitrary resolution. We resolve this issue by updating the weights when the accumulated error reaches a threshold. The threshold corresponds to the magnitude of error that warrants a change in the weights equivalent to their resolution. This update criteria is similar to the gradient accumulation approach where high-precision memory accumulates the gradients, and weights are updated when their gradients reach a certain threshold²². The distinctive factor in our approach is that the weight update decision depends on the local error estimated at the neurons instead of gradients accumulated for every weight. This avoids gradient computation and storage and leads to lower memory overhead and accesses. Spiking networks trained with error-triggered learning can effectively learn a single task³⁰. We combine probabilistic metaplasticity with error threshold-based training to continually learn sequential tasks. Eligible weights are probabilistically updated, where the weight update involves programming the memristor device to its adjacent higher or lower conductance level, depending on the sign of the error. Important weights exhibit low update probability and hence retain their state and help prevent catastrophic forgetting. The detailed learning algorithm is described in Algorithm 1.

Results

Continual learning with probabilistic metaplasticity

We evaluate the proposed probabilistic metaplasticity model considering weights realized with a Hafnium Oxide-based memristor device in the IBM 65nm 10LPe CMOS/Memristor process²⁰. The device can be programmed to ten distinguishable low resistance levels which are used to map the weights (shown in Figure 1b). The average resolution of conductance is ~ 3 bits considering the average conductance at different levels. To realize both positive and negative weights, memristor crossbars are assigned a bias column whose contributions to the output are subtracted from that of the weights with inverting and summing amplifiers³¹. If g_p is the conductance of the memristor weight, g_b is the conductance of the bias memristor, then the synaptic weight can be expressed with Equation 3 where g_f is the scaling factor implemented with the conductance of the feedback resistor in the inverting amplifier.

$$w = \frac{1}{g_f}(g_p - g_b) \quad (3)$$

We evaluate a spiking network with memristor weights on image detection continual learning benchmarks. The network is trained sequentially on tasks without any knowledge of the identity of the task or its boundaries during training or inference. This scenario is known as domain incremental learning (Domain-IL), which presents two major challenges^{32,33}. First, if the task boundary is known, the network can expect a change in data distribution and take measures to reduce catastrophic forgetting. This is not possible in domain-IL. Second, all tasks share the same output layer, which leads to increased interference and catastrophic forgetting. We also consider an online scenario where the network sees each sample only once (the network is trained on each task for one epoch).

We first evaluate probabilistic metaplasticity on the split-MNIST benchmark, which consists of 5 sequential image classification tasks, each composed of two classes from the MNIST dataset³⁴. We conduct network-level simulations in Python considering a spiking network with 200 hidden neurons and 2 output neurons. The simulation emulates in situ learning with memristor weights and captures the quantized nonuniform distribution and the associated variance of memristor conductance during weight updates. Figure 2 shows the evolution of the network performance in each of the benchmark tasks as the network trains on them sequentially. When trained with error threshold-based eRBP with no probabilistic metaplasticity, we see that the network learns each task it encounters with high accuracy (Figure 2a). However, after training on all tasks, it retains high accuracy only on the recently encountered tasks, and the earlier task accuracies drop to random guess level.

Probabilistic metaplasticity consolidates important weights to alleviate such catastrophic forgetting. The degree of consolidation depends on the activity threshold of pre-synaptic and post-synaptic neurons (m_{th}^{pre} and m_{th}^{post} in Equation 2). A low activity threshold leads to high weight consolidation, which helps retain high accuracy in earlier tasks, as shown in Figure 2b. However, the network cannot adapt when later tasks are encountered due to the lower plasticity of the weights, leading to poor performance. A high activity threshold leads to the opposite effect, where the network can learn later tasks better with a compromise in earlier task accuracies (Figure 2c). Optimizing the activity threshold balances the stability-plasticity trade-off to retain high accuracy in earlier tasks with minimal degradation in the ability to learn later tasks (Figure 2d).

We further evaluate probabilistic metaplasticity on the split-Fashion MNIST benchmark, which also consists of five sequential image classification tasks, each with two classes³⁵. To observe the effect of weight resolutions, we consider both single- and multi-memristor weights³⁶. Multi-memristor weights are configured by connecting n_{mem} number of memristors in parallel. Figure 1a shows a 1T1R crossbar with multi-memristor weights consisting of three 1T1R devices. During training, only one memristor arbitrated by a global counter is updated to reduce training overhead. The parallel combination of devices increases the dynamic range of the weight conductance by a factor of n_{mem} and increases the weight resolution. Table 1 lists the mean accuracy across benchmark tasks after sequential training for probabilistic metaplasticity and state-of-the-art weight consolidation-based continual learning models in networks with full-precision weights. We also list the memory required by each model for parameters in addition to weights (details in Supplementary note S1). The baseline results are for a spiking network with full-precision weights trained with eRBP. We note an improvement in the mean accuracy across tasks with increased weight resolution, with the best performance when $n_{mem} = 7$, leading to a mean weight resolution ~ 6 -bit. *Results show that with probabilistic metaplasticity, a spiking network with noisy low-precision multi-memristor weights achieves performance equivalent to state-of-the-art networks with full-precision weights in both benchmark tasks.* For reference, we consider a spiking network that consolidates weights with activity-dependent metaplasticity (details in Methods)¹¹. Since this weight consolidation mechanism requires precise weight updates, we accumulate the weight gradients in auxiliary memory, which is reset after every sample. This model is evaluated with multi-memristor

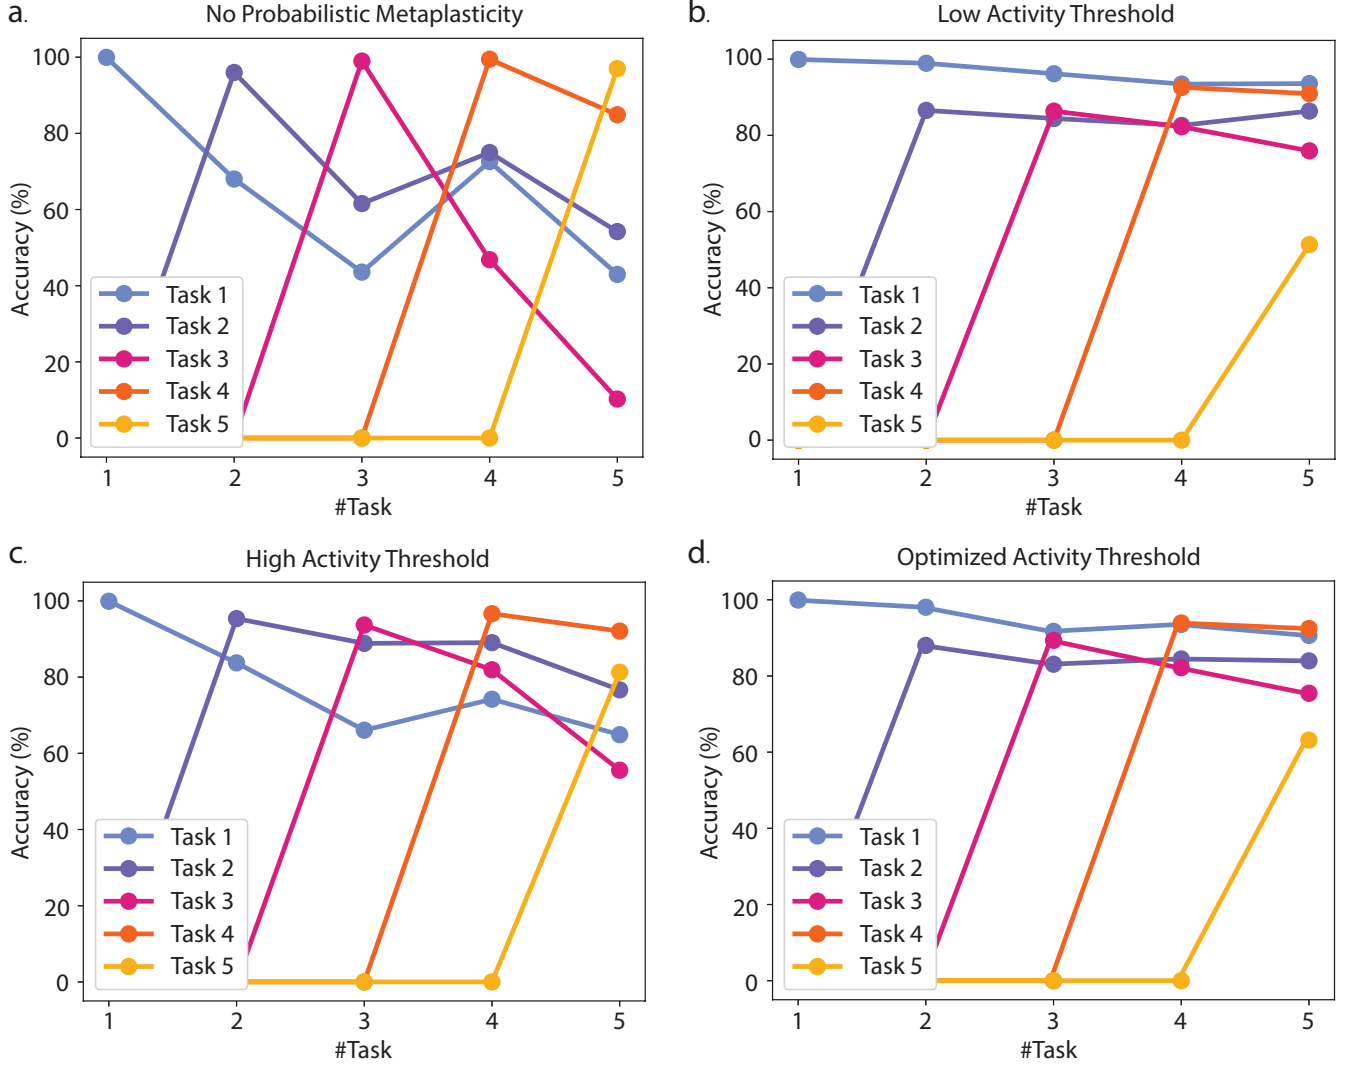


Figure 2. Probabilistic metaplasticity with error threshold-based training for memristor weights.

Evolution of task accuracies as the network is trained sequentially on the split-MNIST tasks. The x-axis shows the latest task the network has learned. We evaluate the performance of task n only after the network encounters it and denote the accuracies as 0 before that. **a.** With no probabilistic metaplasticity, the network learns the current task well, but forgets the initial tasks after sequentially learning multiple tasks. **b.** Probabilistic metaplasticity with low activity threshold leads to high rigidity in the network, so it remembers previous tasks but cannot learn the last task. **c.** High activity threshold can lead to loss in previous task accuracy while the network remains plastic to learn the new tasks. **d.** Optimized activity threshold balances plasticity and rigidity such that the network maintains high initial task accuracies while maintaining the ability to learn new tasks.

weights ($n_{\text{mem}} = 7$) considering the same network dimensions. We see that activity-dependent metaplasticity leads to performance similar to probabilistic metaplasticity in benchmark tasks. However, the latter requires $\sim 67\%$ lower memory overhead.

Finally, we perform two controlled experiments to show the contribution of probabilistic metaplasticity for continual learning. First, we study the contribution of consolidating weights with high magnitude and high adjacent neural activity. For this study, we explore random consolidation in which the update probabilities of the weights are calculated with equation 1, but the results are randomly shuffled among the weights. Second, we study the contribution of probabilistic metaplasticity in contrast to probabilistic plasticity. The controlled experiment evaluates this by updating all weights with the same update probability. We also reduce the weight update probability or plasticity by a factor with each incoming task to help retain previous knowledge. Table 2 shows the results of the controlled

experiments on the split-MNIST benchmark. We carried out the experiments considering multi-memristor weights ($n_{\text{mem}} = 7$) in the same spiking network as in the previous experiments. The baseline network is trained based on an error threshold with no additional mechanism. The table lists the accuracies of each task in the benchmark and the mean accuracy across tasks after sequential training. We see that random consolidation fails to learn continually and suffers from catastrophic forgetting. The performance of the decaying probabilistic plasticity varies with the update probability reduction factor. A high reduction factor leads to consolidation in weights, which retains high accuracies in earlier tasks. However, the network fails to learn the last task. A low reduction factor shows the opposite characteristic, where the network catastrophically forgets the earlier tasks. Controlled experiments on the split-Fashion MNIST benchmark show the same response (Supplementary note S2). Contrary to the experiments, we see that probabilistic metaplasticity retains high accuracy in earlier tasks with much lower degradation in the accuracy of the last task compared to decaying probabilistic plasticity. Since the proposed mechanism consolidates weights selectively based on their magnitude and adjacent neural activity, it retains high plasticity in weights assigned low importance, leading to a better balance of stability and plasticity.

In summary, probabilistic metaplasticity demonstrates continual learning with low-precision memristor weights in the challenging online task-agnostic continual learning scenario, which has not been addressed before in the literature. It achieves state-of-the-art performance with low memory overhead, which is highly suitable for edge applications.

Energy analysis

We assess energy consumption for continual learning with probabilistic metaplasticity by analyzing the computations and memory accesses required to train on the benchmark tasks. We then determine the energy to train one sample by analyzing the average number of operations and the individual operation energy requirements (details in Supplementary note S3). We also perform the analysis for activity-dependent metaplasticity with gradient accumulation as a reference. Figure 3 shows the architecture and computational flow for the analysis. We focus on the energy consumption during the parameter update phase since the forward pass and error computations are common across both approaches. We analyze the activities of a spiking network with 200 hidden neurons and multi-memristor weights ($n_{\text{mem}} = 7$). The mixed-signal and digital components for different network operations were designed in the IBM 65nm 10LPe process. We designed a 16×16 1T1R crossbar array with peripherals to estimate the energy for reading and writing memristor weights. The metaplastic coefficients and the gradients for the activity-dependent metaplasticity mechanism were modeled as on-chip SRAM with 16-bit and 32-bit bus widths respectively (more details in Supplementary note S3).

Figure 4a shows the energy consumption per sample during the parameter update phase, and Figure 4b shows the fraction of energy required for computations and memory accesses for the split-MNIST task. We observe that probabilistic metaplasticity leads to a significant reduction in energy. The reason behind this is mainly twofold. First, metaplasticity with gradient accumulation follows the eRBP criteria to assess a weight’s eligibility to update, while probabilistic metaplasticity adds an additional condition of error threshold to update weights. This leads to much fewer weights eligible for update in the latter (up to three orders of magnitude reduction on average), which consequently leads to less frequent execution of the operations needed for weight update. Second, metaplasticity with gradient accumulation needs to compute and store the gradients for the weights eligible for update. This involves three memory accesses for reading the metaplastic coefficient, the current weight gradient, and storing the updated weight gradient to high-precision memory. On the other hand, probabilistic metaplasticity only needs to read the metaplastic coefficient for eligible weights. This leads to $\sim 3.3\times$ more energy consumption due to memory access for the gradient accumulation approach compared to the probabilistic approach each time the weight update criteria are satisfied. Gradient computation also requires multiplication of the metaplasticity function with the gradient, which consumes more energy compared to the computations required in the probabilistic approach. In summary, gradient accumulation leads to a higher fraction of weights eligible for update, and each eligible weight requires a higher number of memory accesses and computations for gradient computation, leading to higher energy consumption. When the metaplasticity function is complex, such as the exponential function shown in Figure 1c, the gradient accumulation approach requires up to two orders of magnitude higher energy consumption per sample compared to the probabilistic approach. With a bilinear approximation of the exponential function, the energy for the gradient accumulation approach reduces. However, we still observe $\sim 60\times$ reduction in energy consumption per sample with probabilistic metaplasticity.

Memory optimization with parameter sharing

Energy analysis shows that SRAM read/write dominates energy consumption in both continual learning approaches. In the probabilistic metaplasticity model, most of the SRAM read/write operations account for updating the

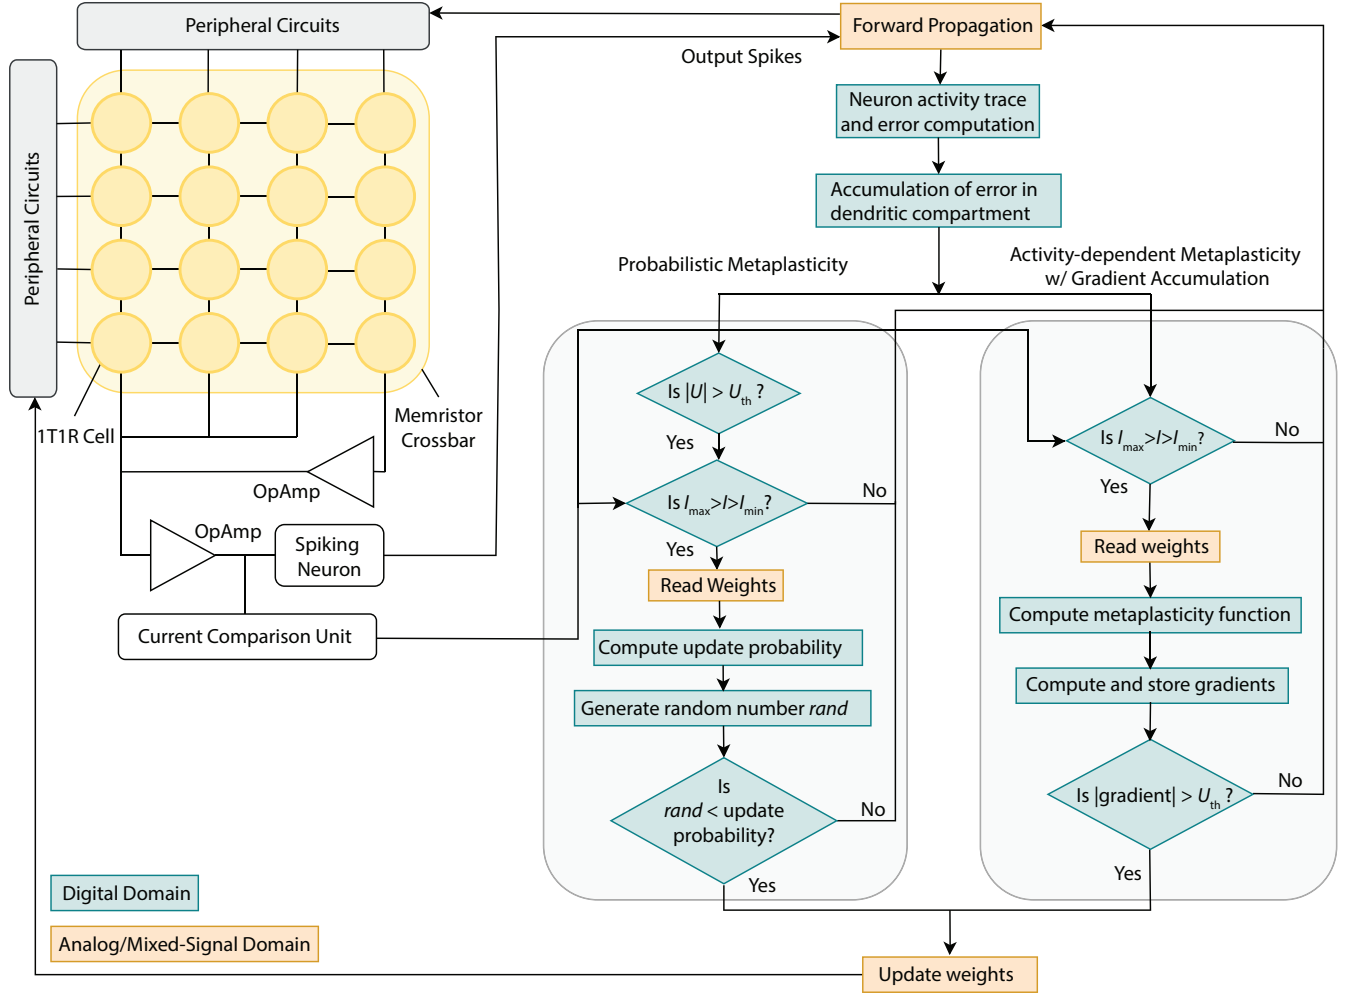


Figure 3. Mixed-signal architecture and computational flow for the continual learning mechanisms.

We consider a mixed-signal architecture where memristor crossbars carry out the forward pass of the spiking model and the error computation and weight update decisions are carried out by digital circuit modules. The figure shows the computational flow for both the proposed mechanism with probabilistic metaplasticity and the gradient accumulation approach with activity-dependent metaplasticity.

metaplastic coefficients assigned to the weights. Here, we optimize the memory overhead by sharing metaplastic coefficients among the weights connected to the same post-synaptic neuron. The metaplastic coefficient associated with the weights connected to the j^{th} post-synaptic neuron evolves as -

$$m_j(t+1) = \begin{cases} m_j(t) + \Delta m, & \text{if } X_j^{\text{tr}} \geq m_{\text{th}}^{\text{post}} \\ m_j(t), & \text{otherwise} \end{cases}, \quad (4)$$

where X_j^{tr} is the activity trace of the j^{th} post-synaptic neuron.

We evaluate the probabilistic metaplasticity model with shared metaplastic coefficients on a spiking network with 200 hidden neurons and multi-memristor weights ($n_{\text{mem}} = 7$). Table 3 shows the evaluation of the benchmark tasks. We see that probabilistic metaplasticity with shared metaplastic coefficients performs better compared to state-of-the-art models with weight consolidation based on neural activity^{37,38}. Furthermore, assigning metaplastic coefficients to neurons instead of weights leads to a substantial decrease in memory overhead and energy consumption during training. Energy analysis during the parameter update phase considering a bilinear metaplasticity function shows that shared parameters lead to up to $\sim 35\times$ reduction in energy consumption per sample (Figure 5a). Figure 5b shows the breakdown of energy due to computations, SRAM and memristor read/writes for the split-MNIST task. Shared metaplastic coefficients reduce the number of memory accesses by up to two orders of magnitude,

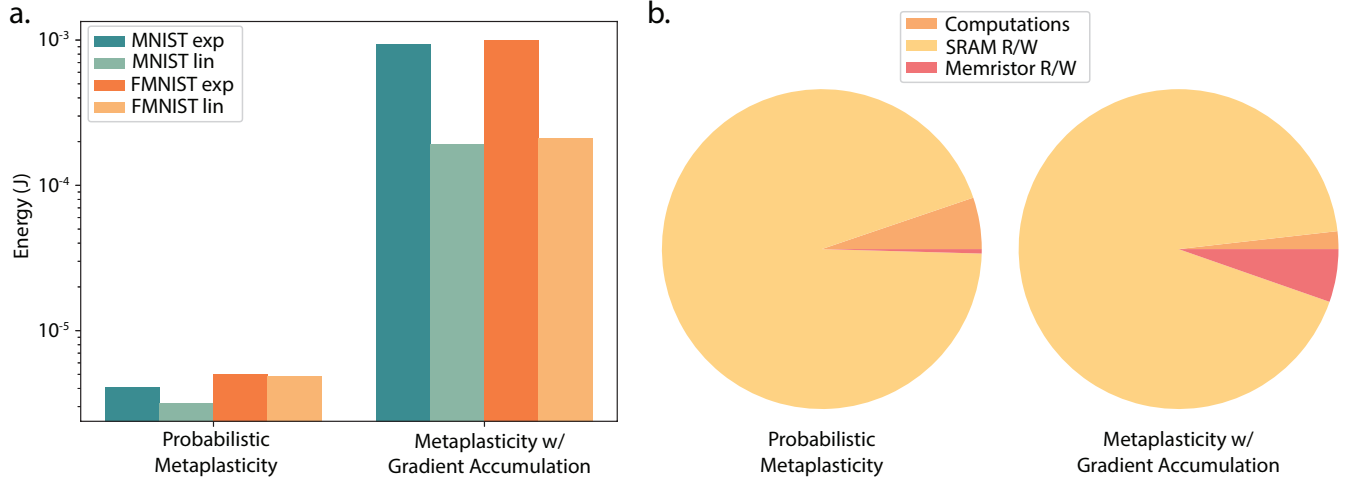


Figure 4. Energy consumption during parameter update phase. **a.** Energy required during the parameter update phase per sample with probabilistic metaplasticity and activity-dependent metaplasticity with gradient accumulation considering exponential (exp) and bilinear (lin) metaplasticity functions. **b.** Breakdown of energy consumption during parameter update in the split-MNIST task. In the diagram, "Computation" denotes all operations that do not involve SRAM or memristor read/write.

and due to the smaller size of the memory block, each access requires up to $\sim 23\times$ lower energy. This leads to up to three orders of magnitude reduction in energy required for memory accesses. Sharing metaplastic coefficients also improves the scalability of the probabilistic metaplasticity model, as it reduces the space complexity of the algorithm from $O(n^2)$ to $O(n)$.

The reduction in energy consumption, however, is accompanied by a minor degradation in performance compared to the probabilistic metaplasticity model with individual metaplastic coefficients. This can be attributed to the limited flexibility of the network to consolidate weights, which leads to a poorer balance of the stability-plasticity dilemma. In the split-MNIST task, shared metaplastic parameters lead to reduced plasticity after learning the five sequential tasks. This can be observed from the update probability of the weights, which is inversely proportional to the product of the metaplastic coefficient and the weight magnitude ($|mw|$). The distribution of $|mw|$ in hidden weights (Figure 5c) shows that individual metaplastic coefficients lead to a higher fraction of weights with $|mw|$ close to zero. Individual task accuracies after sequential training also show that shared parameters result in poorer accuracies in later tasks. Despite the reduction in plasticity, the substantial gain in energy consumption incentivizes shared metaplastic coefficients in extreme-edge applications. Introducing mechanisms to reduce the magnitude of metaplastic coefficients based on neural activity can potentially improve the balance of stability and plasticity³⁹.

Discussion

Emerging nonvolatile memristor devices offer immense opportunities for energy-efficient continual learning on the edge. However, their limited precision and high variability may constrain their suitability for continual learning solutions. Our research takes inspiration from the computational principles of the brain to devise a continual learning mechanism suited to non-ideal low-precision memory devices. Computations in the brain are characterized by stochastic and low-precision synapses^{40, 41}. Research in computational memory models demonstrates that stochastic binary synapses with low-precision metaplastic states can help memory retention²⁶. Our work draws a parallel between the low-precision hidden metaplastic states of a cascade metaplastic synapse and the conductance levels of a memristor device with two key distinctions. First, instead of binarized weights, we considered multi-level weights represented by memristors. Second, in addition to the state of the synapse, our approach considers the activity of the adjacent neurons to determine the plasticity of a weight. Previous research showed that such activity-dependent metaplasticity leads to a better balance of the stability-plasticity dilemma¹¹.

In this work, we demonstrated the efficacy of probabilistic metaplasticity in a spiking network trained with error threshold-based eRBP, where weights were realized with a Hafnium oxide-based memristor. However, the proposed mechanism is not limited to the setting above. Probabilistic metaplasticity naturally extends to other surrogate gradient-based training algorithms for spiking networks where the local error at the neurons can be estimated⁴². It

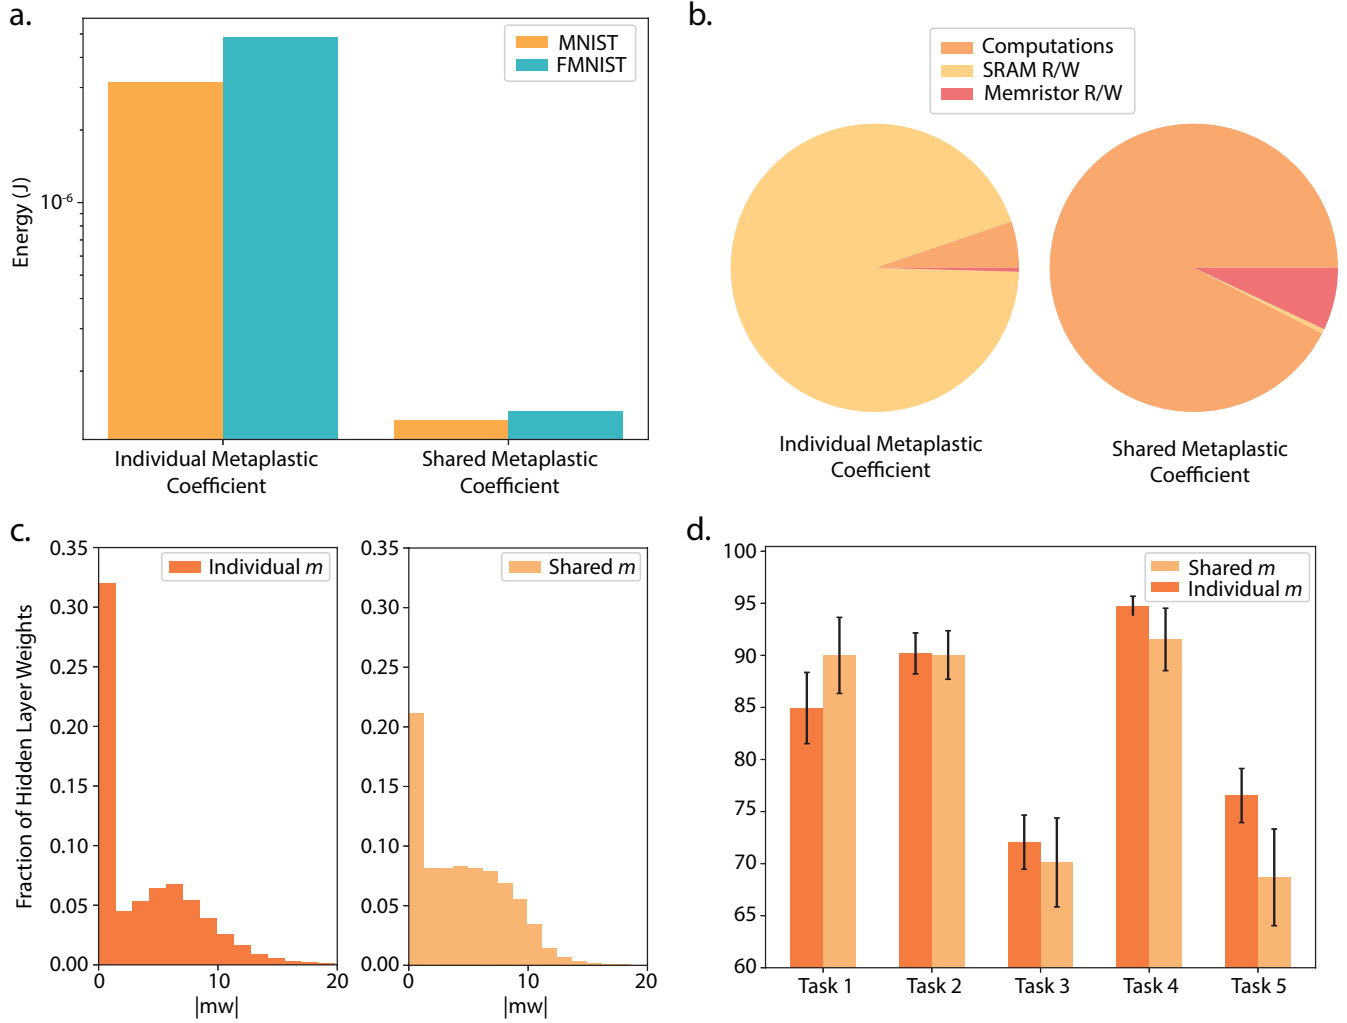


Figure 5. Effect of individual and shared metaplastic coefficients on energy and stability-plasticity **a.** Energy consumption during the parameter update phase of the split-MNIST and split-Fashion MNIST tasks for probabilistic metaplasticity with individual and shared metaplastic coefficients. **b.** The breakdown of energy for computations, SRAM read/write, and memristor read/write for the split-MNIST task. **c.** Distribution of the product of metaplastic coefficient m and weight magnitude $|w|$ of the hidden layer of the network trained on the split-MNIST task. We see that the shared m leads to lower fraction of weights with low $|mw|$ value. **d.** The individual task accuracies in the split-MNIST benchmark after sequential training.

can also be explored in conjunction with gradient accumulation-based training for low-precision weights. In this setting, probabilistic metaplasticity will eliminate the need to multiply metaplastic factors with gradients; however, the high frequency of memory accesses may lead to a lower gain in energy consumption. The proposed mechanism can also be explored with other emerging nonvolatile memory devices. Recent research has explored probabilistic training algorithms for emerging low-precision memory devices in deep neural networks considering learning settings with independent and identically distributed data^{43,44}. Probabilistic metaplasticity can be explored in these networks to alleviate catastrophic forgetting during sequential learning with non-stationary data distribution. It is also promising to enable continual learning in neuromorphic accelerators with low-precision weights⁴⁵.

We explored the effect of different weight resolutions by employing multi-memristor weights alongside single-memristor weights. Although we observed improved performance with higher resolution weights, it increases the area footprint of memristor crossbars. Improving the resolution of a single memristor device can lead to better performance with lower area overhead. Recent research has shown memristor devices with up to ~ 6 -bit precision²¹. In addition to the resolution of conductance modulation, the nonlinearity and variability in memristor devices can affect the performance of error threshold-based training. Due to these non-ideal characteristics, weight updates

vary in response to the same local error depending on the current conductance level of the memristor. To minimize this effect, we restricted weight mapping to only ten low-resistance states of the 1T1R device. This choice avoids the higher variability observed in the high-resistance states along with the substantial difference in conductance between the highest low-resistance and high-resistance levels during weight updates²⁰. Extreme non-uniformity in the distribution of conductance levels can potentially degrade the performance of error threshold-based training, which can be explored in future studies.

In summary, we propose probabilistic metaplasticity for continual learning with low-precision memristor weights. The proposed approach consolidates important weights by modulating the probability of weight update rather than its magnitude. This eliminates the need for auxiliary high-precision memory for weight consolidation, leading to energy-efficient training with fewer memory accesses and computations. Benchmark evaluations demonstrated that probabilistic metaplasticity effectively prevents catastrophic forgetting in the online continual learning scenario with no task supervision, satisfying the criteria for autonomous applications. Moreover, it enables a spiking network with low-precision memristor weights to achieve accuracies equivalent to state-of-the-art continual learning models with high-precision weights while requiring lower memory overhead. Probabilistic metaplasticity can thus lead to efficient solutions for autonomous continual learning on edge.

Methods

Memristor device

We evaluate the proposed probabilistic metaplasticity model considering Hafnium Oxide-based 1T1R cells as weights. The Resistive Random Access Memory (RRAM) device consists of a TiN/Ti/HfO₂/Ti stack connected to a transistor in the IBM 65nm 10LPe CMOS/Memristor process²⁰. The device can be used in binary mode by setting and resetting it to high and low conductive states. During SET operation, modulating the compliance currents by controlling the gate voltage of the transistor gives rise to 10 distinct levels in the high-conductive state, with mean conductance in the range of $\sim 40 \mu\text{S}$ - $283 \mu\text{S}$ with an average conductance resolution of $\sim 27 \mu\text{S}$ ²⁰. When used as a synaptic weight, the resolution of this device (calculated as (maximum conductance - minimum conductance)/average conductance resolution) is close to 3 bits of precision. In the multi-memristor weight, due to the parallel combination of memristors, the maximum and minimum conductance of the weight increases. However, since only one device is updated during training, the average conductance resolution during weight update is the same as a single memristor weight. We use scaling to map the multi-memristor weights to the desired range. With four and seven memristors per weight, the weight precision is close to 4 and 6 bits.

Spiking neural network training algorithm

We train the multilayer spiking neural network with event-driven random backpropagation (eRBP), which uses surrogate gradients²⁹. The input stimuli to the network are encoded as Poisson spike trains, where the spike rate is proportional to the strength of the input. The spike train propagates through the weights and accumulates as the input current to the Leaky Integrate and Fire neurons. For the j^{th} neuron, the accumulated current can be expressed as -

$$I_j(t+1) = I_j(t) + \frac{\Delta t}{\tau_{\text{syn}}} \left(\sum_{i=1}^N w_{ij} S_i(t) - I_j(t) \right). \quad (5)$$

Here $S_i(t)$ is the input spike at the i^{th} pre-synaptic neuron at timestep t , N is the total number of input neurons, τ_{syn} is the synaptic time constant, and w_{ij} is the synaptic weight between i^{th} and j^{th} pre- and post-synaptic neurons. This current is integrated at the neuron membrane potential as shown in Equation 6, where τ_{mem} , R , C and V_{rest} are the membrane time constant, membrane resistance, capacitance and the resting potential of the neuron.

$$V_j(t+1) = V_j(t) + \frac{\Delta t}{\tau_{\text{mem}}} \left((V_{\text{rest}} - V_j(t)) + I_j(t) R \right) \quad (6)$$

When the membrane potential reaches its threshold V_{th} , the neuron emits a spike S_j , and its membrane potential is reset to the resting potential. A refractory period follows every time a neuron spikes, during which the neuron remains inactive.

The spike train progresses through the network to produce the output layer spike train S^{out} . It is compared to the target spike train L to produce the error current I^{err} -

$$I^{\text{err}}(t) = S^{\text{out}}(t) - L(t). \quad (7)$$

Two sets of error neurons, for false positive and false negative errors, spike in response to I^{err} . The output spikes of error neurons $S^{\text{fp/fn}}$ are used to compute the error in the output and hidden layers. At the output layer o , the error at the j^{th} neuron is -

$$E_j^o(t) = S_j^{\text{fp}}(t) - S_j^{\text{fn}}(t). \quad (8)$$

While for the hidden layer h , the error is backpropagated to the i^{th} neuron by -

$$E_i^h(t) = \sum_{j=1}^{n_{\text{out}}} \left(w_{ij}^{\text{fp},h} S_j^{\text{fp}}(t) - w_{ij}^{\text{fn},h} S_j^{\text{fn}}(t) \right), \quad (9)$$

where n_{out} is the number of output neurons, and $w_{ij}^{\text{fp},h}$ and $w_{ij}^{\text{fn},h}$ are random feedback weights from error neurons encoding false positives and false negatives, respectively.

The neurons in the hidden and output layers integrate the errors as -

$$U(t+1) = U(t) + \frac{\Delta t}{\tau_{\text{mem}}} \left(-U(t) + E(t)R \right). \quad (10)$$

As the dendritic compartment accumulates the error, the weights are updated in proportion to it whenever there is an active pre-synaptic spike, and the post-synaptic current is within a range. The update in weight can be expressed as -

$$\Delta w_{ij} = -\eta S_i U_j \Theta(I_j), \quad (11)$$

where η is the learning rate, S_i is the pre-synaptic spike, U_j is the error integrated at the dendritic compartment of the j^{th} post-synaptic neuron, and Θ is a boxcar function, which indicates that the current magnitude at post-synaptic neuron j falls within the threshold.

Activity-dependent metaplasticity with gradient accumulation for continual learning

Metaplasticity in neural networks is usually abstracted with a metaplastic factor, which modulates the magnitude of weight gradients^{11, 15}. The factor assumes a low value for important weight, which restricts subsequent changes and preserves previously learned information. In this work, we adopt activity-dependent metaplasticity¹¹ for continual learning with gradient accumulation. According to this mechanism, the importance of a weight is determined by its magnitude and the activity of the adjacent neurons. The function is expressed as -

$$f(m_{ij}, w_{ij}) = e^{-|m_{ij} w_{ij}|} \quad (12)$$

Here, w_{ij} is the weight value, and m_{ij} is a metaplastic coefficient assigned to a weight that captures the activity of its adjacent neurons. It is updated by computing the trace of the neuron activity as -

$$m_{ij}(t+1) = \begin{cases} m_{ij}(t) + \Delta m, & \text{if } X_j^{\text{tr}} \geq m_{\text{th}}^{\text{pre}} \& X_i^{\text{tr}} \geq m_{\text{th}}^{\text{post}} \\ m_{ij}(t), & \text{otherwise} \end{cases}, \quad (13)$$

The weight gradients are modulated as -

$$\Delta w_{ij} = -\eta S_i U_j \Theta(I_j) f(m_{ij}, w_{ij}) \quad (14)$$

The required weight updates are highly precise and cannot be directly incorporated into low-precision memristor devices. Therefore, the weight gradients are accumulated in high-precision memory during sample spike-train presentations. Whenever the gradient reaches a threshold corresponding to the memristor resolution, the weights are updated. The gradients are reset after each sample presentation. The complete training procedure is described in Supplementary Algorithm S1.

Data availability statement

All datasets used in this work (MNIST³⁴ and Fashion-MNIST³⁵) are publicly available.

Code availability statement

The source code supporting the figures and experiments in this work is available from the corresponding author upon reasonable request.

References

1. Hayes, T. L. & Kanan, C. Online Continual Learning for Embedded Devices. *arXiv preprint arXiv:2203.10681* (2022).
2. Kukreja, N. *et al.* Training on the Edge: The why and the how. In *2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, 899–903, DOI: [10.1109/IPDPSW.2019.00148](https://doi.org/10.1109/IPDPSW.2019.00148) (2019).
3. Dally, W. On the Model of Computation: Point. *Commun. ACM* **65**, 30–32 (2022).
4. French, R. M. Catastrophic forgetting in connectionist networks. *Trends Cogn. Sci.* **3**, 128–135, DOI: [https://doi.org/10.1016/S1364-6613\(99\)01294-2](https://doi.org/10.1016/S1364-6613(99)01294-2) (1999).
5. Verwimp, E. *et al.* Continual Learning: Applications and the Road Forward. *arXiv preprint arXiv:2311.11908* (2023).
6. Kudithipudi, D. *et al.* Design principles for lifelong learning AI accelerators. *Nat. Electron.* 1–16 (2023).
7. Balasubramanian, V. Heterogeneity and Efficiency in the Brain. *Proc. IEEE* **103**, 1346–1358, DOI: [10.1109/JPROC.2015.2447016](https://doi.org/10.1109/JPROC.2015.2447016) (2015).
8. Kudithipudi, D. *et al.* Biological underpinnings for lifelong learning machines. *Nat. Mach. Intell.* **4**, 196–210, DOI: [10.1038/s42256-022-00452-0](https://doi.org/10.1038/s42256-022-00452-0) (2022).
9. Abraham, W. C. & Bear, M. F. Metaplasticity: the plasticity of synaptic plasticity. *Trends Neurosci.* **19**, 126–130, DOI: [https://doi.org/10.1016/S0166-2236\(96\)80018-X](https://doi.org/10.1016/S0166-2236(96)80018-X) (1996).
10. Benna, M. K. & Fusi, S. Computational principles of synaptic memory consolidation. *Nat. Neurosci.* **19**, 1697–1706, DOI: [10.1038/nn.4401](https://doi.org/10.1038/nn.4401) (2016).
11. Soures, N., Helfer, P., Daram, A., Pandit, T. & Kudithipudi, D. TACOS: Task Agnostic Continual Learning in Spiking Neural Network. In *Theory and Foundation of Continual Learning Workshop at ICML'2021* (July 2021).
12. Kirkpatrick, J. *et al.* Overcoming catastrophic forgetting in neural networks. *Proc. Natl. Acad. Sci.* **114**, 3521–3526, DOI: [10.1073/pnas.1611835114](https://doi.org/10.1073/pnas.1611835114) (2017). <https://www.pnas.org/doi/pdf/10.1073/pnas.1611835114>.
13. Zenke, F., Poole, B. & Ganguli, S. Continual Learning Through Synaptic Intelligence. In Precup, D. & Teh, Y. W. (eds.) *Proceedings of the 34th International Conference on Machine Learning*, vol. 70 of *Proceedings of Machine Learning Research*, 3987–3995 (PMLR, 2017).
14. Kaplanis, C., Shanahan, M. & Clopath, C. Continual Reinforcement Learning with Complex Synapses. In Dy, J. & Krause, A. (eds.) *Proceedings of the 35th International Conference on Machine Learning*, vol. 80 of *Proceedings of Machine Learning Research*, 2497–2506 (PMLR, 2018).
15. Laborieux, A., Ernault, M., Hirtzlin, T. & Querlioz, D. Synaptic metaplasticity in binarized neural networks. *Nat. Commun.* **12**, 2549 (2021).
16. Han, B., Sengupta, A. & Roy, K. On the Energy Benefits of Spiking Deep Neural Networks: A Case Study. In *2016 International Joint Conference on Neural Networks (IJCNN)*, 971–976, DOI: [10.1109/IJCNN.2016.7727303](https://doi.org/10.1109/IJCNN.2016.7727303) (2016).
17. Yu, S. Neuro-Inspired Computing with Emerging Nonvolatile memory. *Proc. IEEE* **106**, 260–285, DOI: [10.1109/JPROC.2018.2790840](https://doi.org/10.1109/JPROC.2018.2790840) (2018).
18. Xia, Q. & Yang, J. J. Memristive crossbar arrays for brain-inspired computing. *Nat. Mater.* **18**, 309–323, DOI: [10.1038/s41563-019-0291-x](https://doi.org/10.1038/s41563-019-0291-x) (2019).
19. Cheng, M. *et al.* TIME: A Training-in-Memory Architecture for Memristor-Based Deep Neural Networks. In *Proceedings of the 54th Annual Design Automation Conference 2017, DAC '17*, DOI: [10.1145/3061639.3062326](https://doi.org/10.1145/3061639.3062326) (Association for Computing Machinery, New York, NY, USA, 2017).

20. Liehr, M., Hazra, J., Beckmann, K., Rafiq, S. & Cady, N. Impact of Switching Variability of 65nm CMOS Integrated Hafnium Dioxide-based ReRAM Devices on Distinct Level Operations. In *2020 IEEE International Integrated Reliability Workshop (IIRW)*, 1–4 (IEEE, 2020).
21. Park, J. *et al.* TiOx-Based RRAM Synapse With 64-Levels of Conductance and Symmetric Conductance Change by Adopting a Hybrid Pulse Scheme for Neuromorphic Computing. *IEEE Electron Device Lett.* **37**, 1559–1562, DOI: [10.1109/LED.2016.2622716](https://doi.org/10.1109/LED.2016.2622716) (2016).
22. Nandakumar, S. R. *et al.* Mixed-Precision Deep Learning Based on Computational Memory. *Front. Neurosci.* **14**, DOI: [10.3389/fnins.2020.00406](https://doi.org/10.3389/fnins.2020.00406) (2020).
23. Li, Y. *et al.* Mixed-precision continual learning based on computational resistance random access memory. *Adv. Intell. Syst.* **4**, 2200026, DOI: <https://doi.org/10.1002/aisy.202200026> (2022). <https://onlinelibrary.wiley.com/doi/pdf/10.1002/aisy.202200026>.
24. D’Agostino, S. *et al.* Synaptic metaplasticity with multi-level memristive devices. In *2023 IEEE 5th International Conference on Artificial Intelligence Circuits and Systems (AICAS)*, 1–5, DOI: [10.1109/AICAS57966.2023.10168563](https://doi.org/10.1109/AICAS57966.2023.10168563) (2023).
25. Lu, A. *et al.* High-speed emerging memories for AI hardware accelerators. *Nat. Rev. Electr. Eng.* **1**, 24–34, DOI: [10.1038/s44287-023-00002-9](https://doi.org/10.1038/s44287-023-00002-9) (2024).
26. Fusi, S., Drew, P. J. & Abbott, L. F. Cascade Models of Synaptically Stored Memories. *Neuron* **45**, 599–611 (2005).
27. Zohora, F. T., Karia, V., Daram, A. R., Ziyarah, A. M. & Kudithipudi, D. MetaplasticNet: Architecture with Probabilistic Metaplastic Synapses for Continual Learning. In *2021 IEEE International Symposium on Circuits and Systems (ISCAS)*, 1–5, DOI: [10.1109/ISCAS51556.2021.9401262](https://doi.org/10.1109/ISCAS51556.2021.9401262) (2021).
28. Crestani, A. P. *et al.* Metaplasticity contributes to memory formation in the hippocampus. *Neuropsychopharmacology* **44**, 408–414, DOI: [10.1038/s41386-018-0096-7](https://doi.org/10.1038/s41386-018-0096-7) (2019).
29. Neftci, E. O., Augustine, C., Paul, S. & Detorakis, G. Event-Driven Random Back-Propagation: Enabling Neuromorphic Deep Learning Machines. *Front. Neurosci.* **11**, 324 (2017).
30. Payvand, M., Fouda, M. E., Kurdahi, F., Eltawil, A. M. & Neftci, E. O. On-Chip Error-Triggered Learning of Multi-Layer Memristive Spiking Neural Networks. *IEEE J. on Emerg. Sel. Top. Circuits Syst.* **10**, 522–535, DOI: [10.1109/JETCAS.2020.3040248](https://doi.org/10.1109/JETCAS.2020.3040248) (2020).
31. Ziyarah, A. M. & Kudithipudi, D. Semi-Trained Memristive Crossbar Computing Engine with In Situ Learning Accelerator. *J. Emerg. Technol. Comput. Syst.* **14**, DOI: [10.1145/3233987](https://doi.org/10.1145/3233987) (2018).
32. van de Ven, G. M. & Tolias, A. S. Three scenarios for continual learning. *arXiv:1904.07734 [cs, stat]* (2019). [1904.07734](https://arxiv.org/abs/1904.07734).
33. Hsu, Y.-C., Liu, Y.-C., Ramasamy, A. & Kira, Z. Re-evaluating Continual Learning Scenarios: A Categorization and Case for Strong Baselines. In *NeurIPS Continual learning Workshop* (2018).
34. Lecun, Y., Bottou, L., Bengio, Y. & Haffner, P. Gradient-Based Learning Applied to Document Recognition. *Proc. IEEE* **86**, 2278–2324, DOI: [10.1109/5.726791](https://doi.org/10.1109/5.726791) (1998).
35. Xiao, H., Rasul, K. & Vollgraf, R. Fashion-MNIST: A Novel Image Dataset for Benchmarking Machine Learning Algorithms. *arXiv preprint arXiv:1708.07747* (2017).
36. Boybat, I. *et al.* Neuromorphic computing with multi-memristive synapses. *Nat. Commun.* **9**, 1–12 (2018).
37. Kim, S. & Lee, S. Continual Learning with Neuron Activation Importance. In Sclaroff, S., Distant, C., Leo, M., Farinella, G. M. & Tombari, F. (eds.) *Image Analysis and Processing – ICIAP 2022*, 310–321 (Springer International Publishing, Cham, 2022).
38. Daram, A. & Kudithipudi, D. NEO: Neuron State Dependent Mechanisms for Efficient Continual Learning. In *Proceedings of the 2023 Annual Neuro-Inspired Computational Elements Conference, NICE ’23*, 11–19, DOI: [10.1145/3584954.3584960](https://doi.org/10.1145/3584954.3584960) (Association for Computing Machinery, New York, NY, USA, 2023).
39. Soures, N. *Lifelong Learning in Spiking Neural Networks Through Neural Plasticity*. Ph.D. thesis, Rochester Institute of Technology (2023).
40. Bartol, T. M. *et al.* Hippocampal Spine Head Sizes Are Highly Precise. *bioRxiv* DOI: [10.1101/016329](https://doi.org/10.1101/016329) (2015). <https://www.biorxiv.org/content/early/2015/03/11/016329.full.pdf>.

41. O'Connor, D. H., Wittenberg, G. M. & Wang, S. S.-H. Graded bidirectional synaptic plasticity is composed of switch-like unitary events. *Proc. Natl. Acad. Sci.* **102**, 9679–9684, DOI: [10.1073/pnas.0502332102](https://doi.org/10.1073/pnas.0502332102) (2005). <https://www.pnas.org/doi/pdf/10.1073/pnas.0502332102>.
42. Kaiser, J., Mostafa, H. & Neftci, E. Synaptic Plasticity Dynamics for Deep Continuous Local Learning (DECOLLE). *Front. Neurosci.* **14**, DOI: [10.3389/fnins.2020.00424](https://doi.org/10.3389/fnins.2020.00424) (2020).
43. Misba, W. A., Lozano, M., Querlioz, D. & Atulasimha, J. Energy Efficient Learning With Low Resolution Stochastic Domain Wall Synapse for Deep Neural Networks. *IEEE Access* **10**, 84946–84959, DOI: [10.1109/ACCESS.2022.3196688](https://doi.org/10.1109/ACCESS.2022.3196688) (2022).
44. Zhang, Y., He, G., Tang, K.-T. & Wang, G. On-chip Learning of Multilayer Perceptron Based on Memristors with Limited Multilevel States. In *2019 IEEE International Conference on Artificial Intelligence Circuits and Systems (AICAS)*, 11–12, DOI: [10.1109/AICAS.2019.8771513](https://doi.org/10.1109/AICAS.2019.8771513) (2019).
45. Davies, M. *et al.* Loihi: A Neuromorphic Manycore Processor with On-Chip Learning. *IEEE Micro* **38**, 82–99, DOI: [10.1109/MM.2018.112130359](https://doi.org/10.1109/MM.2018.112130359) (2018).
46. Li, Z. & Hoiem, D. Learning without Forgetting. *IEEE Transactions on Pattern Analysis Mach. Intell.* **40**, 2935–2947, DOI: [10.1109/TPAMI.2017.2773081](https://doi.org/10.1109/TPAMI.2017.2773081) (2018).
47. 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)* (2018).
48. Zeno, C., Golan, I., Hoffer, E. & Soudry, D. Task Agnostic Continual Learning Using Online Variational Bayes. *arXiv preprint arXiv:1803.10123* (2018).
49. Schug, S., Benzing, F. & Steger, A. Presynaptic stochasticity improves energy efficiency and helps alleviate the stability-plasticity dilemma. *eLife* **10**, e69884, DOI: [10.7554/eLife.69884](https://doi.org/10.7554/eLife.69884) (2021).

Acknowledgements

This effort is partially supported by NSF EFRI BRAID Award #2317706, NSF NAIAD Award #2332744 and Air Force Research Laboratory under agreement number FA8750-20-2-1003 through BAA FA8750-19-S-7010. We thank Dr. Maximilian Liehr and Dr. Nathaniel Cady for helpful discussions on the RRAM device characteristics. We thank Dr. Abdullah Ziyarah for helpful discussions and feedback. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of NSF, Air Force Research Laboratory or the U.S. Government.

Author contributions statement

F.Z. conceived the idea, conducted the experiments, and wrote the manuscript. V.K. contributed to the energy estimation for the digital components and memory models. N.S. helped develop the experiments. D. K. oversaw the development and execution of the research and offered critical feedback. All authors reviewed the manuscript.

Additional information

Competing interests

The authors declare the following competing interests: The authors declare no competing interests.

Figure legends

- **Figure 1: Probabilistic metaplasticity with error threshold-based training.** **a.** Spiking network trained on error threshold where the network weights are realized with 1T1R memristor crossbar array. The dendritic compartments in the hidden and output layer neurons integrate the error. When the error reaches the threshold U_{th} , the memristor weights are updated to the next higher (negative error) or lower (positive error) conductance level with probability p_{update} . **b.** Mean and standard deviation of resistance levels versus programming compliance current for the 1T1R memristor device (shown in inset) adopted in this work²⁰. **c.** Update probability of weights for different values of the metaplastic coefficient m and the weights w . m is directly proportional to the activity level of adjacent neurons. We see that weights with highly active adjacent neurons (high m) and weight magnitude (high $|w|$) lead to low update probability.

- **Figure 2: Probabilistic metaplasticity with error threshold-based training for memristor weights.** Evolution of task accuracies as the network is trained sequentially on the split-MNIST tasks. The x-axis shows the latest task the network has learned. We evaluate the performance of task n only after the network encounters it and denote the accuracies as 0 before that. **a.** With no probabilistic metaplasticity, the network learns the current task well, but forgets the initial tasks after sequentially learning multiple tasks. **b.** Probabilistic metaplasticity with low activity threshold leads to high rigidity in the network, so it remembers previous tasks but cannot learn the last task. **c.** High activity threshold can lead to loss in previous task accuracy while the network remains plastic to learn the new tasks. **d.** Optimized activity threshold balances plasticity and rigidity such that the network maintains high initial task accuracies while maintaining the ability to learn new tasks.
- **Figure 3: Mixed-signal architecture and computational flow for the continual learning mechanisms.** We consider a mixed-signal architecture where memristor crossbars carry out the forward pass of the spiking model and the error computation and weight update decisions are carried out by digital circuit modules. The figure shows the computational flow for both the proposed mechanism with probabilistic metaplasticity and the gradient accumulation approach with activity-dependent metaplasticity.
- **Figure 4: Energy consumption during parameter update phase.** **a.** Energy required during the parameter update phase per sample with probabilistic metaplasticity and activity-dependent metaplasticity with gradient accumulation considering exponential (exp) and bilinear (lin) metaplasticity functions. **b.** Breakdown of energy consumption during parameter update in the split-MNIST task. In the diagram, "Computation" denotes all the operations that do not involve SRAM or memristor read/write.
- **Figure 5: Effect of individual and shared metaplastic coefficients on energy and stability-plasticity** **a.** Energy consumption during the parameter update phase of the split-MNIST and split-Fashion MNIST tasks for probabilistic metaplasticity with individual and shared metaplastic coefficients. **b.** The breakdown of energy for computations, SRAM read/write, and memristor read/write for the split-MNIST task. **c.** Distribution of the product of metaplastic coefficient m and weight magnitude $|w|$ of the hidden layer of the network trained on the split-MNIST task. We see that the shared m leads to lower fraction of weights with low $|mw|$ value. **d.** The individual task accuracies in the split-MNIST benchmark after sequential training.

Table 1. The mean accuracies across tasks with standard deviations of different continual learning models on the split-MNIST and split-Fashion MNIST tasks after sequential training in the domain-IL scenario.

Model	Split-MNIST	Split-Fashion MNIST	Memory Overhead
Baseline	60.69 % \pm 0.6	75.52% \pm 1.31	0 kB
LwF ⁴⁶	71.50% \pm 1.63	71.02% \pm 0.46	788 kB
MAS ⁴⁷	66.42% \pm 2.47	68.57% \pm 6.85	1.5 MB
BGD ⁴⁸	80.44% \pm 0.45	89.73% \pm 0.88	1.5 MB
SS ⁴⁹	82.90% \pm 0.01	91.98% \pm 0.12	788 kB
TACOS ¹¹	82.56% \pm 1.12	93.22% \pm 0.22	1.182 MB
Probabilistic Metaplasticity			
$n_{\text{mem}} = 1$	81.14% \pm 1.54	92.56% \pm 1.49	314.4 kB
$n_{\text{mem}} = 2$	82.67% \pm 1.21	92.42% \pm 0.63	314.4 kB
$n_{\text{mem}} = 7$	83.69% \pm 0.78	93.23% \pm 0.14	314.4 kB
Gradient Accumulation w/ metaplasticity ($n_{\text{mem}} = 7$)	83.18% \pm 0.42	92.33% \pm 0.27	943.2 kB

Table 2. Performance of a spiking network on the split-MNIST task in different settings. The table shows the mean and standard deviation of the individual task accuracies and the mean accuracy across tasks over 5 runs after sequential training.

	Baseline	Random Consolidation	Decaying Probabilistic Plasticity			Probabilistic Metaplasticity
			(Factor=2)	(Factor=5)	(Factor=10)	
Task 1	44.31 \pm 3.80	48.01 \pm 2.77	47.05 \pm 1.63	50.11 \pm 1.89	73.58 \pm 2.80	84.95 \pm 3.42
Task 2	53.67 \pm 1.42	57.14 \pm 0.61	56.39 \pm 0.62	94.31 \pm 0.27	94.76 \pm 0.43	90.20 \pm 1.96
Task 3	8.63 \pm 1.04	17.73 \pm 1.09	22.42 \pm 1.53	90.51 \pm 0.40	84.59 \pm 0.59	72.06 \pm 2.60
Task 4	88.31 \pm 1.24	96.24 \pm 0.55	96.68 \pm 0.44	93.59 \pm 0.91	82.72 \pm 1.39	94.73 \pm 0.94
Task 5	97.52 \pm 0.32	95.73 \pm 0.62	95.17 \pm 0.38	48.85 \pm 1.66	40.44 \pm 1.16	76.54 \pm 2.60
Mean	58.49 \pm 0.73	62.97 \pm 0.56	63.54 \pm 0.36	75.47 \pm 0.32	75.22 \pm 0.79	83.70 \pm 0.78

Table 3. Evaluation of probabilistic metaplasticity and neural activity-dependent weight consolidation models on the split-MNIST and split-Fashion MNIST tasks in the domain-IL scenario.

Model	Split-MNIST	Split-FMNIST	Memory Overhead
Probabilistic Metaplasticity (Individual m)	83.70 \pm 0.78	93.23 \pm 0.14	314.4 kB
Probabilistic Metaplasticity (Shared m)	82.07 \pm 0.84	92.72 \pm 0.30	\sim 0.5 kB
NAI ³⁷	68.35 \pm 1.34	68.82 \pm 1.15	3.2 kB
NEO ³⁸	78.14 \pm 2.23	86.82 \pm 0.60	\sim 1 MB