Forward-only learning in memristor arrays with month-scale stability

A. Renaudineau(1), MH. Diallo(2)*, T. Dupuis(1)*, B. Imbert(2), MA. Iftakher(1), K-E. Harabi(3), C. Turck(1), T. Hirtzlin(3), D. Bonnet(1), F. Melul(2), J-D. Aguirre-Morales(2), E. Vianello(3), M. Bocquet(2), J-M. Portal(2), and D. Querlioz(1)

(1) Université Paris-Saclay, CNRS, Centre de Nanosciences et de Nanotechnologies, Palaiseau, France,
(2) Aix-Marseille Université, CNRS, Institut Matériaux Microélectronique Nanosciences de Provence, Marseille, France,
(3) Université Grenoble-Alpes, CEA, LETI, Grenoble, France.
*These authors contributed equally to this work

Memristor arrays implement analog multiply-accumulate (MAC) operations by Ohm's and Kirchhoff's laws, enabling remarkably low-energy inference in dense crossbars [1]. However, turning memristor arrays from inference engines into systems capable of on-chip learning has proved difficult. Weight updates have a high energy cost and cause device wear, analog states drift, and backpropagation requires a backward pass with reversed signal flow. Here we experimentally demonstrate learning on standard filamentary HfOx/Ti arrays that addresses these challenges with two design choices [2]. First, we rely on forward-only training algorithms in the Forward-Forward family [3] that use only inference-style operations. Second, we use sub-1 V reset-only, single-pulse updates that cut energy and yield stable analog states. We train two-layer classifiers on an ImageNet-resolution four-class task using arrays up to 8,064 devices. Two forward-only variants (a two-pass supervised Forward-Forward and a single-pass competitive rule) achieve test accuracies of 89.5% and 89.6%, respectively; a reference experiment using backpropagation reaches 90.0%. Across five independent runs per method, these accuracies are indistinguishable within statistical uncertainty. Trained models retain accuracy for at least one month under ambient conditions, consistent with the stability of reset-only states. Sub-1 V reset updates use 460 times less energy than conventional program-and-verify programming and require just 46% more energy than inference-only operation. Together, these results establish forward-only, sub-1 V learning on standard filamentary stacks at array scale, outlining a practical, pulse-aware route to adaptive edge intelligence.

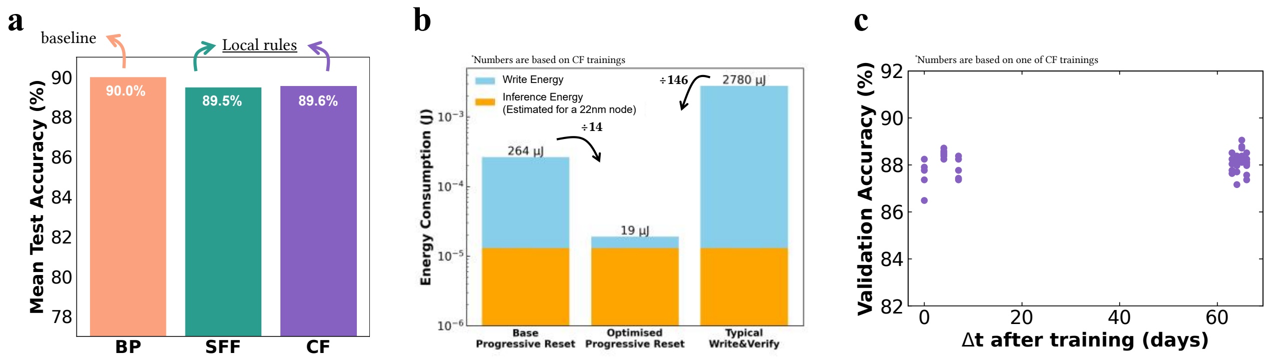

Figure 1: **a**. Test accuracy comparing the experimental results with a Backprop (BP) reference and two Forward-only learning rules ("Supervised Forward-Forward", SFF and "Cluster-Forward", CF). **b**. Energy results comparing the standard Write&Verify memristor programming scheme to our proposed single-pulse, reset-only updates. **c**. Retention test showcasing accuracy stable up to two months after training.

[1] M. Le Gallo et al., Nat. Elecron., 6, 680-693, 2023.

[2] A. Renaudineau et al., arxiv:2601.09903, 2026.

[3] G. Hinton, arxiv:2212.13345, 2022.

This work benefited from France 2030 grants (ANR-22-PEEL-0010, ANR-23-PEIA-0002).