Teach it new things. Make it forget. Keep learning.

A learner with no backpropagation and no stored training photos, tested on vision, radar and speech. Every number below carries its label.

0.149 s

Delete 5,000 photos; the sums match a model that never saw them, bit for bit

89.55%

Radar learned only from simulation, tested on measured chips; six looks, 3 seeds

87.17%

CIFAR-10, one look, 3 seeds, no preview; 5 eyes in different lesson orders

The numbers

CIFAR-10	87.17 ± 0.19%, one look, three seeds, settings chosen on validation, no preview. 5 eyes, each taught in a different lesson order: 5 passes, not one stream. CIFAR-10.1: 74.38%. One strict eye: 84.93%, one look, three seeds.
CIFAR-100	64.8%, one look, 7 eyes, unlabelled preview; 65.46% with six looks; one ensemble. One strict eye: 61.7%, three seeds, 50 categories then ten lessons of 5, no preview, no pretraining; settings selected earlier on CIFAR-10 test subsets, then frozen.
MNIST	99.06%, five looks, closed form, all at once, hand-set constants, one seed.
Fashion-MNIST	91.73% with two looks, 90.83% with one, same model, one seed. Best so far, pending reruns.
Radar	Trained on 806 simulated chips, tested on 539 measured: 89.55 ± 0.28%, six looks, three seeds. Vehicle blanked: 10.39%, near chance. DenseNet + noise: 77.7. Backprop ensembles: over 95. Earlier versions were read on this test. KTH 77 GHz, whole recordings held out: 89.1% on nine classes, 98.1% drone or not, three seeds; dataset classification, not field detection.
Speech	Speech Commands v2, all 35 words: 91.2% average across seven lessons, 88.9% final, three seeds, official test list.
Continual text	71.02% across five datasets (575,000 training documents), zero replay, no backprop. Unsupervised source router: 94.18%. Counts 1.54M tokens/s in 32 MB (table counting).
Exact delete and merge	Removing 5,000 photos settles in 0.149 s; integer sums show zero difference from a clean run, and answers match it on 10,799 of 10,800. Two copies with a shared eye merged in opposite orders: bit-identical.
Readout memory	7,835 MB down to 1.887 MB live RAM and 612 s down to 11.10 s, while CIFAR-10 at six looks rose from 83.10% to 85.21%. Readout stage only; eye and on-disk codes are separate. Earlier full-pipeline runs peaked at 17.0–17.6 GiB.
Small whole model	97.40% MNIST and 89.39% Fashion in 150.46 KB (about 150 KB), whole model, no backprop, no matrix solve. Timing-based eye, local error rule on one layer, multi-pass, all at once, linear reader at test, one seed.
Timing neurons	16 timing neurons name 3,963 synthetic temporal sequences at 99% over three seeds. Graded release decided by a true firing threshold reaches 99.57% ± 0.44, cutting stray-noise false fires from 27.0% to 0.50%. Two-beat chords with check neurons hold 97.92% / 95.75% / 91.58% over 1 / 2 / 4 slots.

What the measurements establish

These are our own benchmark runs, not independent replication or customer performance. They ran on a laptop, not on a phone, robot or chip. Published references are not matched protocols. Earlier cards are superseded: an earlier CIFAR-10 headline used test subsets to choose settings and has been retired.

Try to break it

Pick a public test and keep the labels; we return predictions for you to score. On your own categories, a sealed $5,000 evaluation: you keep the held-out test labels, and we return predictions and measured accuracy, retention, model size and learning time. No source code or unpublished method details are shared.

Updated 3 November 2026. Numbers and wording from the claim register and the 3 November audit of run files. Primary logs remain private and available for a scoped evaluation.

Alexander Argeadai · mars@buildcivilization.com · +1 913-543-0947 · Lenexa, Kansas