


Hardware-Software Co-Design of Analog Recurrent Neural Networks for Sub-Microwatt Inference



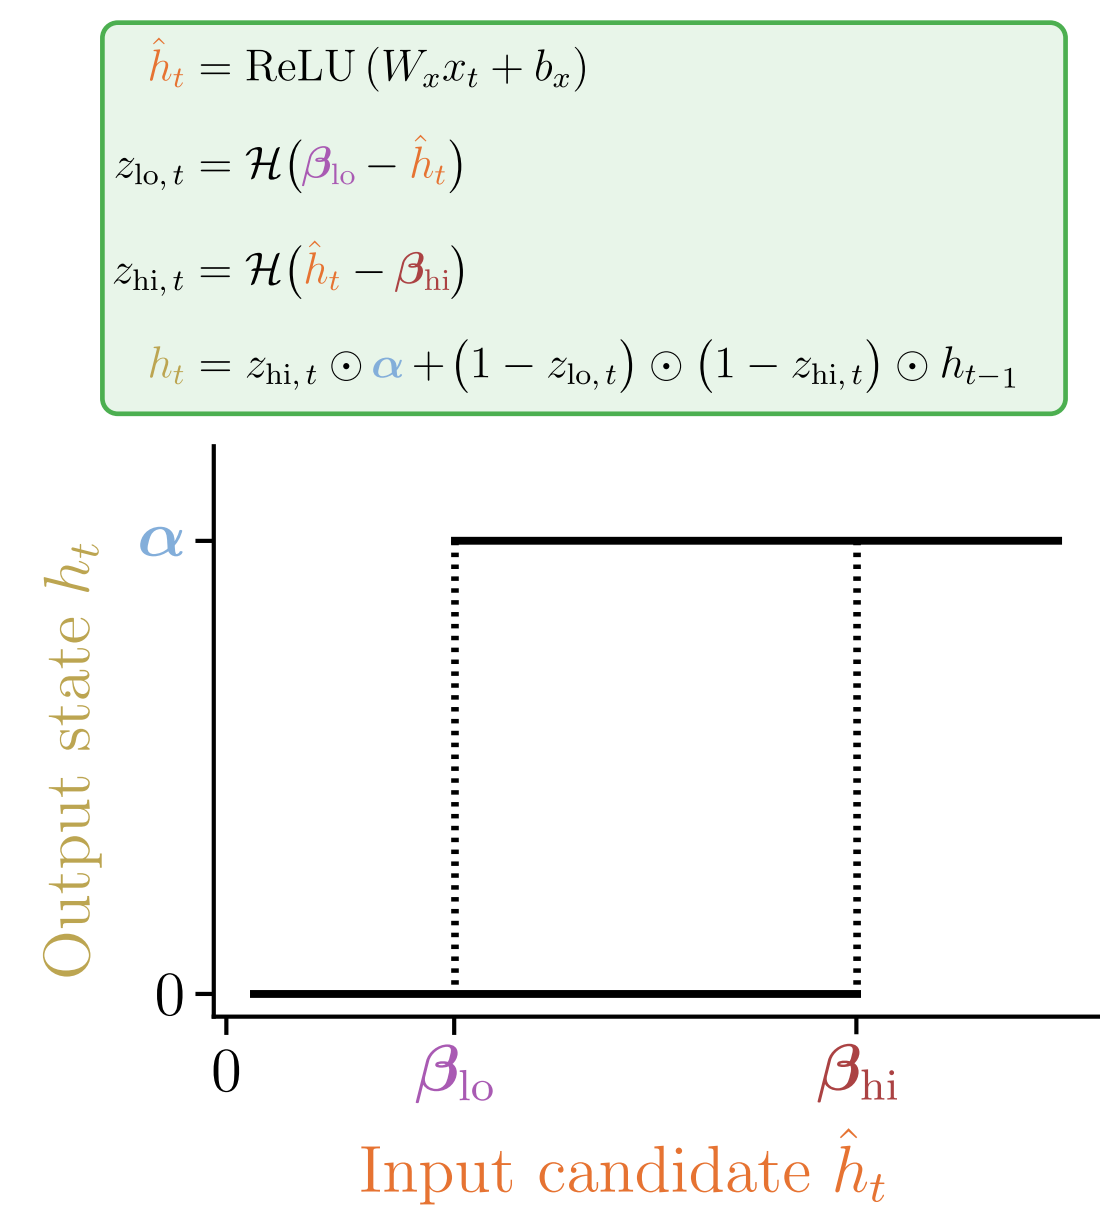
J. Brandoit · A. Fyon · L. Mendolia · D. Ernst · J.-M. Redouté · G. Drion

Department of Electrical Engineering and Computer Science, University of Liège, Belgium · ICNCE 2026

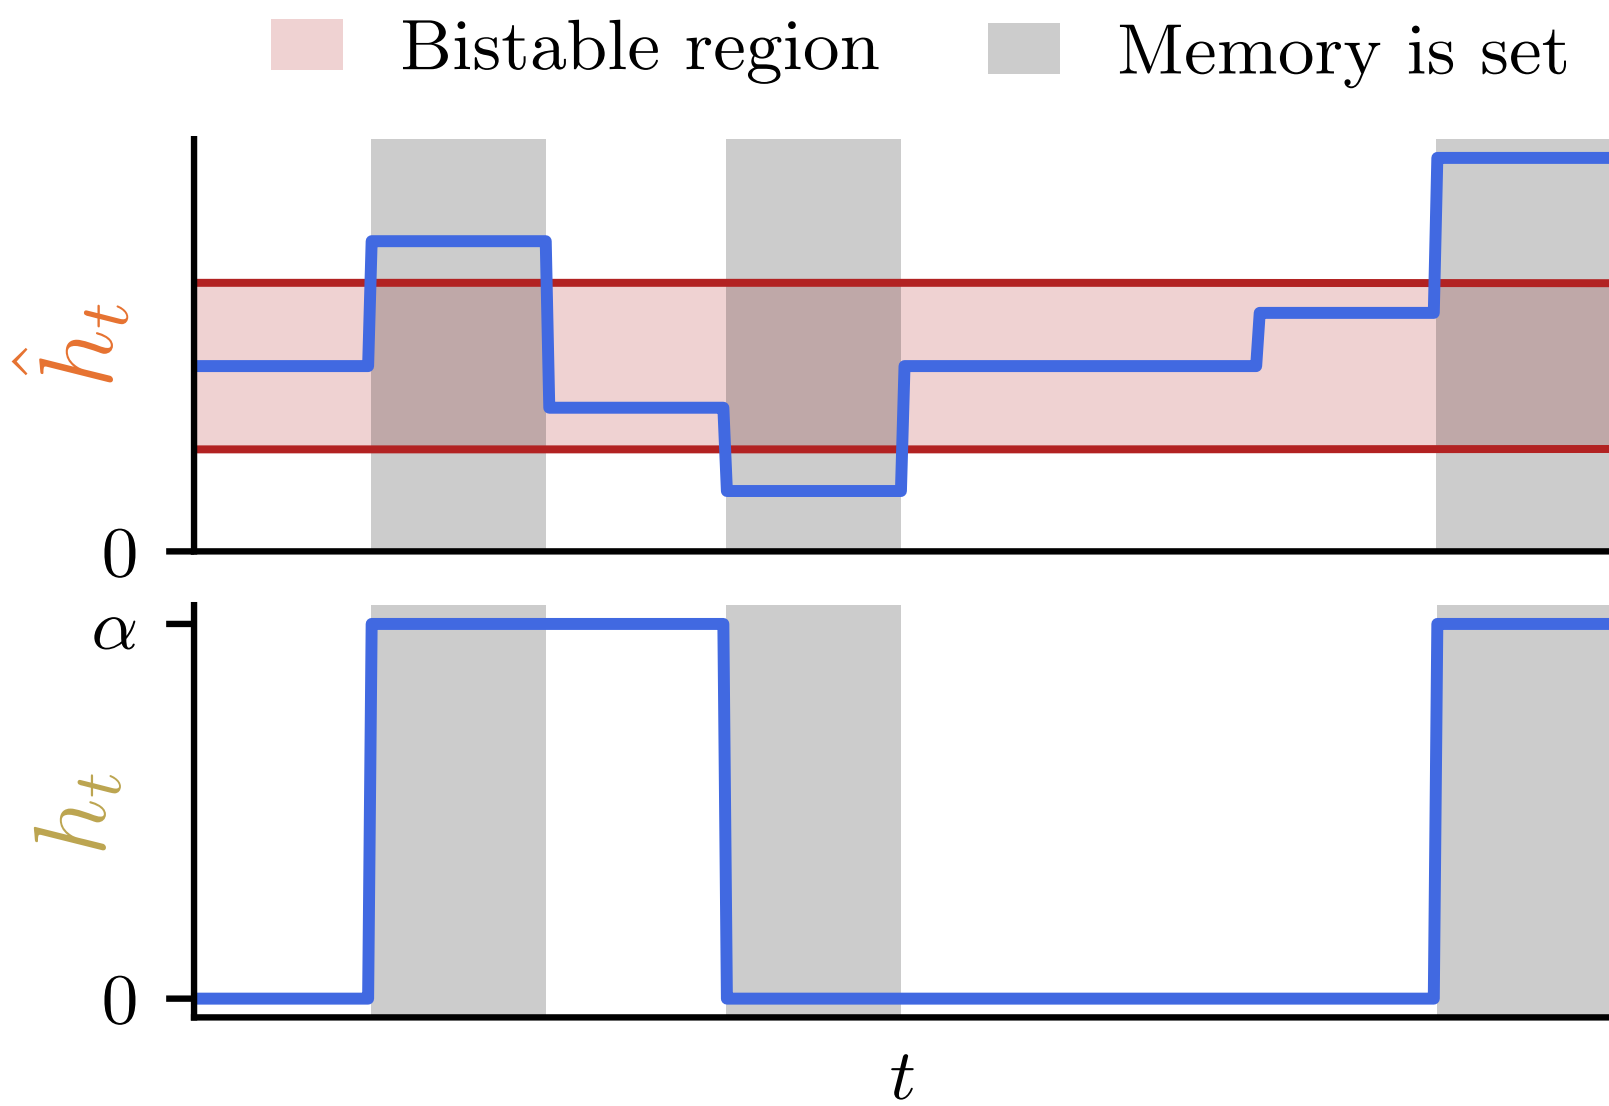
Always-on AI needs **sub-microwatt** inference, but digital burns milliwatts and analog recurrence drowns in **noise**.

Recurrence gives sequential sensing its memory, but had never run fully in analog: feedback lets noise accumulate. We solve this by **co-designing** a recurrent quantized cell that is analog by nature and compatible with ultralow-power applications.

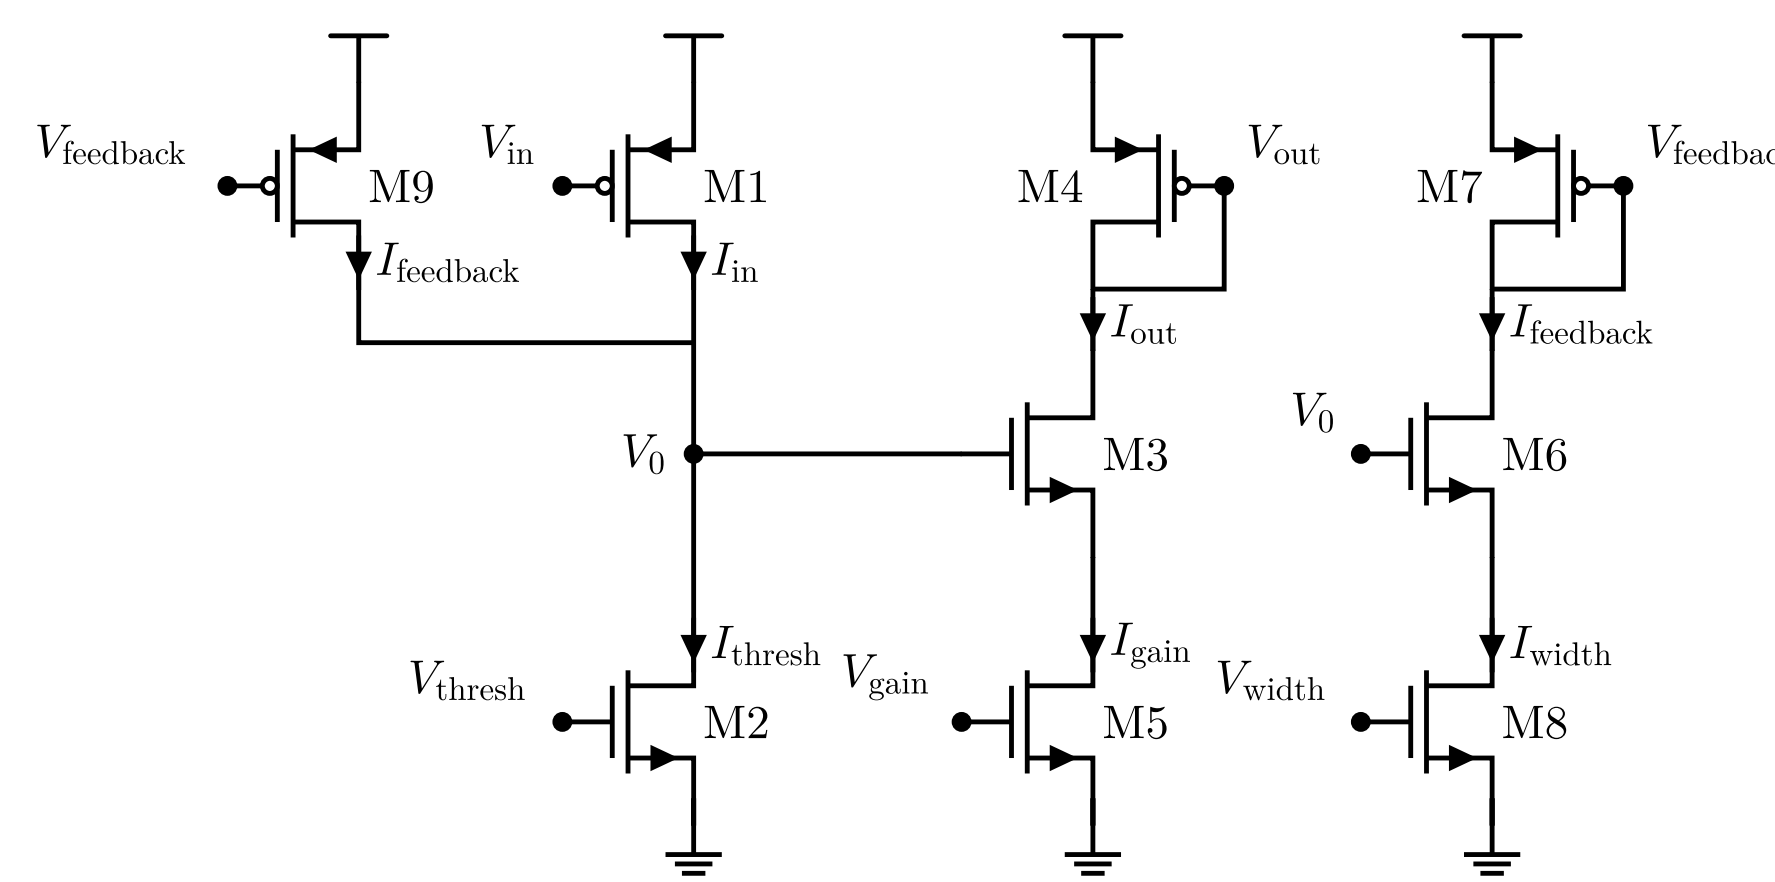
Software



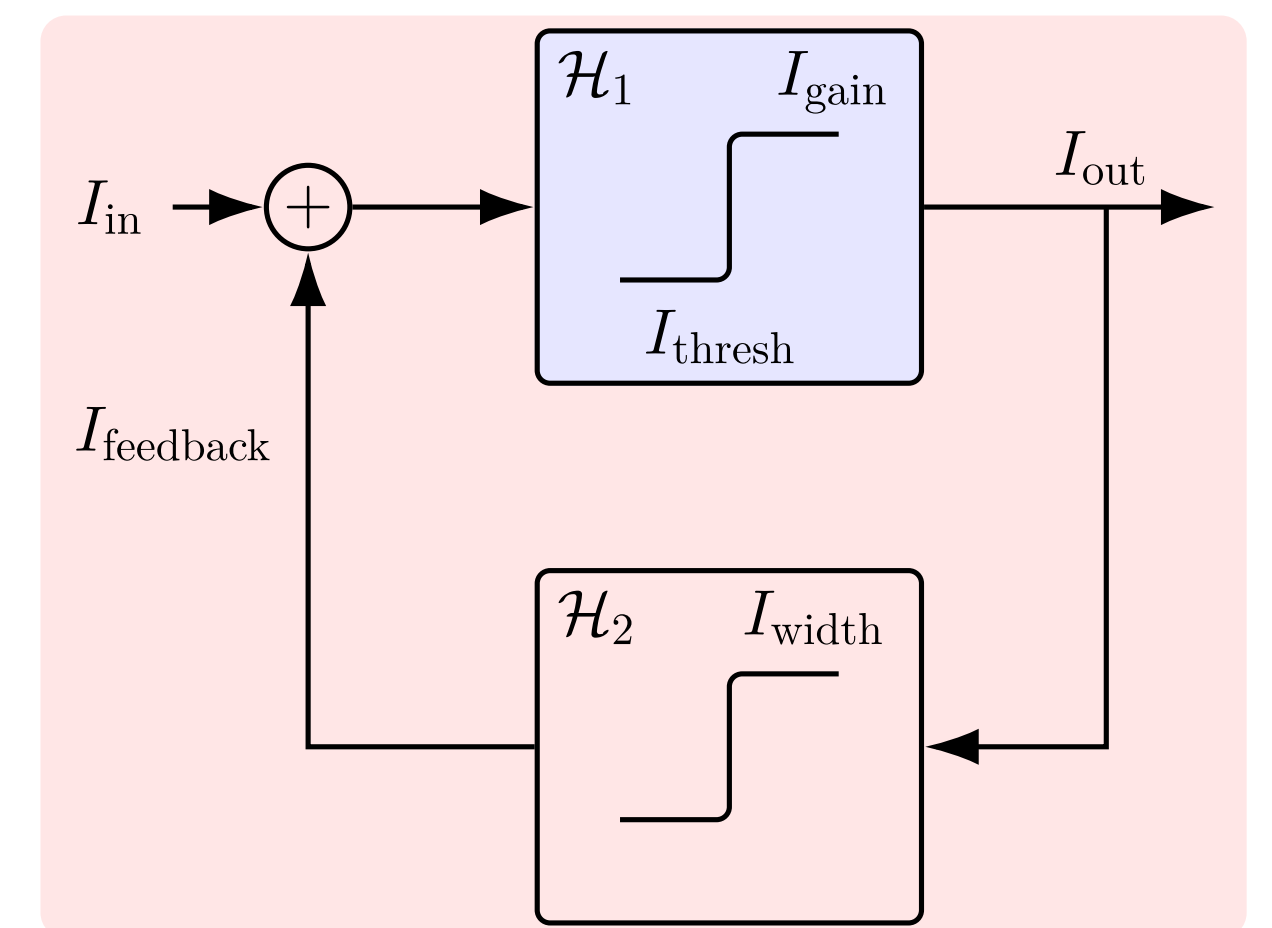
What we train !



Hardware



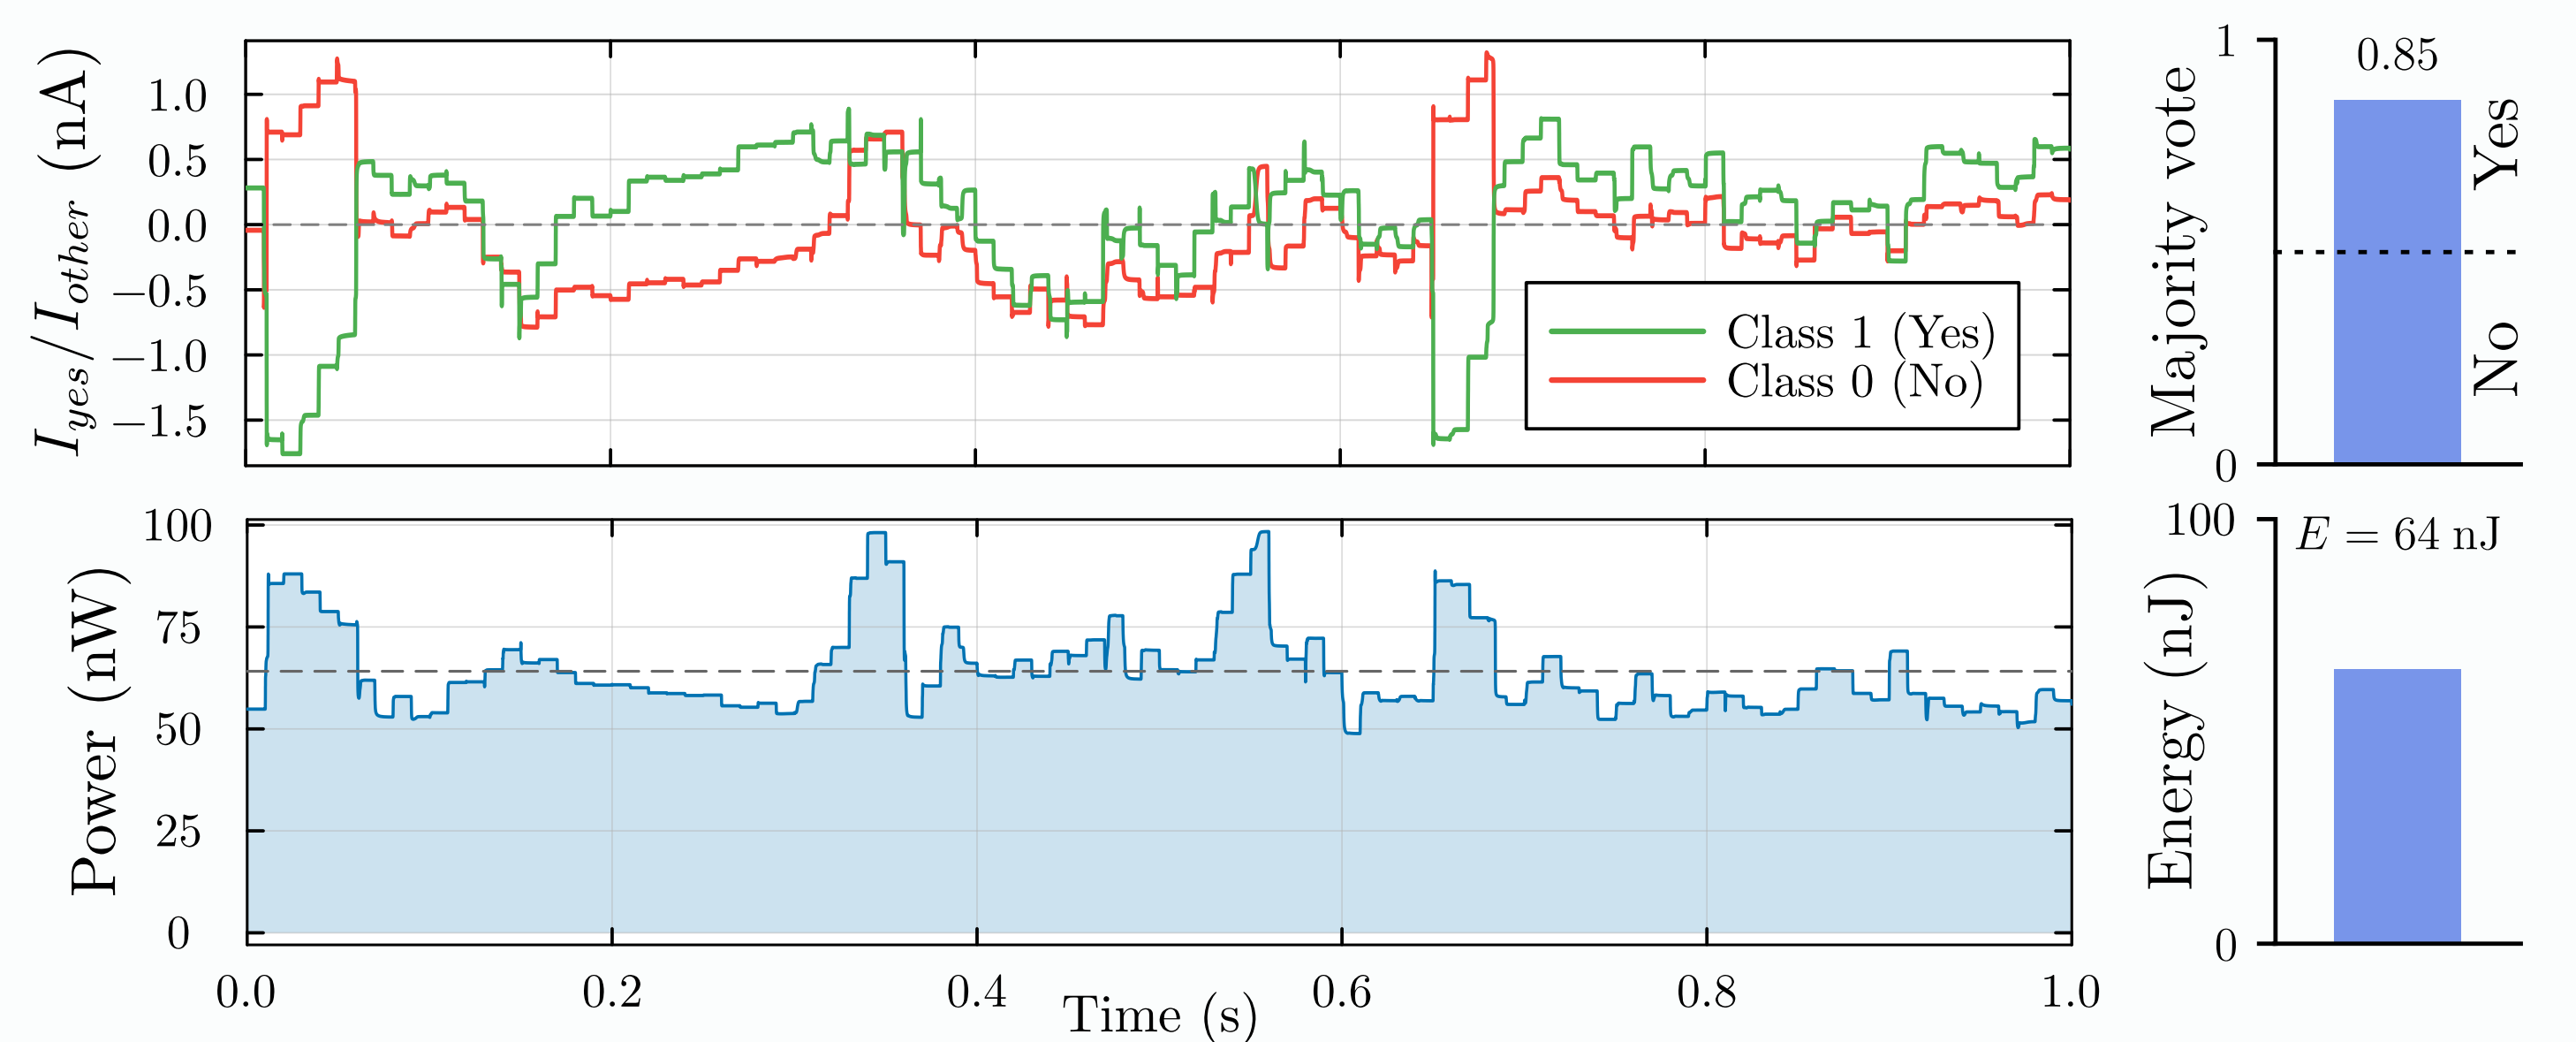
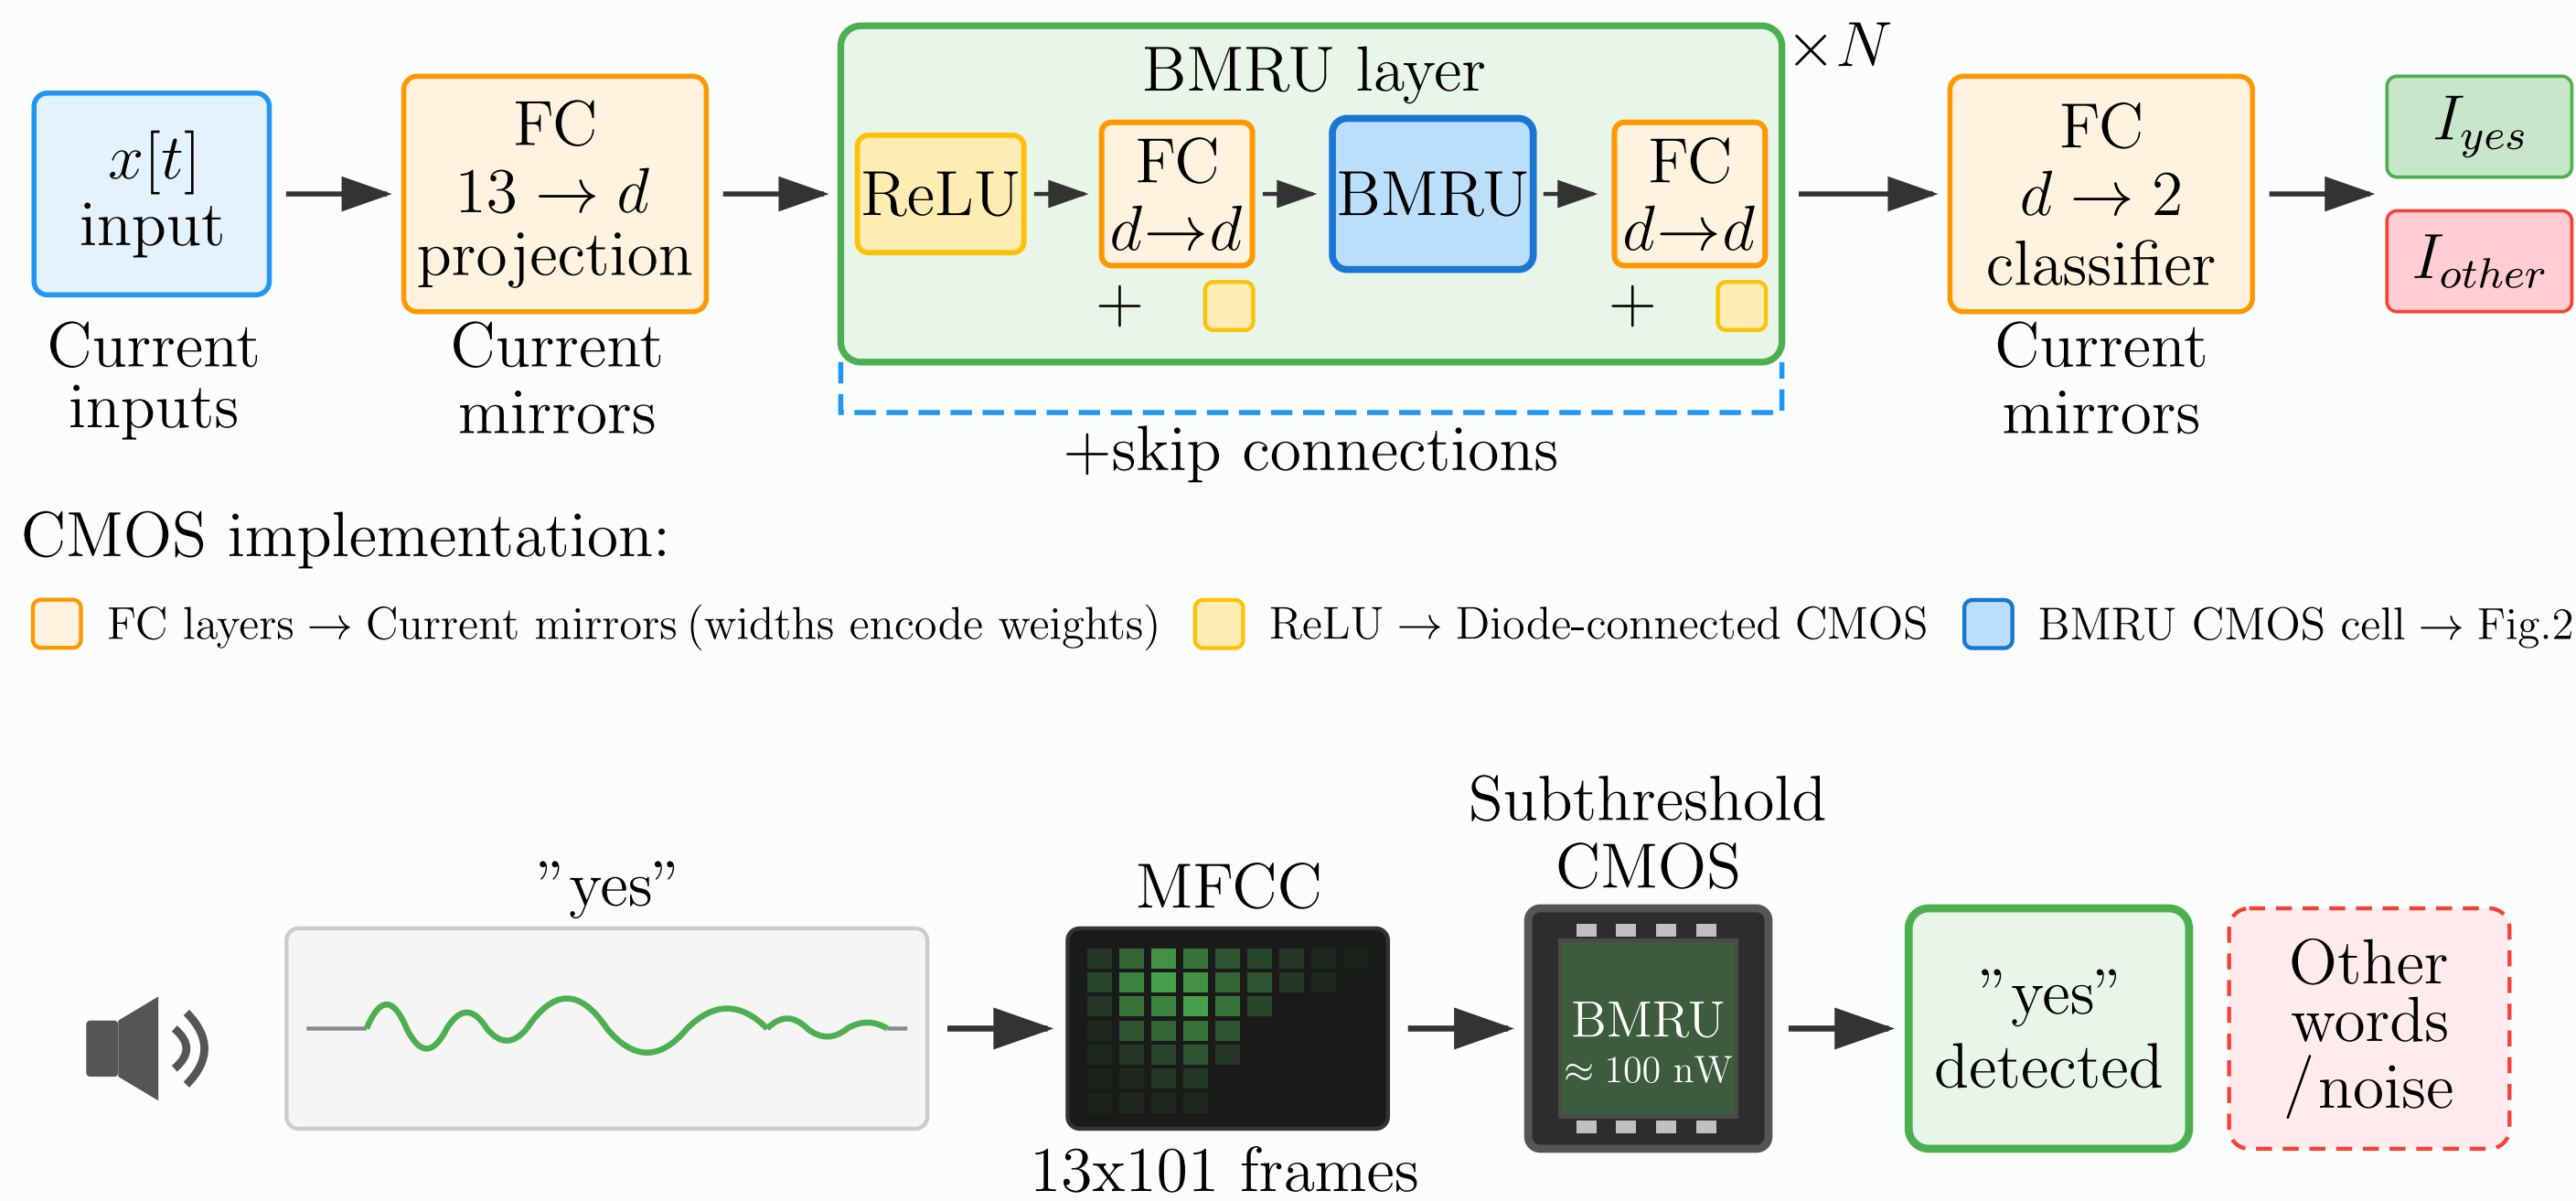
What we build !



The recurrent cell is a **hysteretic comparator**: its discrete state holds inside the learned band $[\beta_{lo}, \beta_{hi}]$ and updates to 0 or α once input pushes it outside the band.

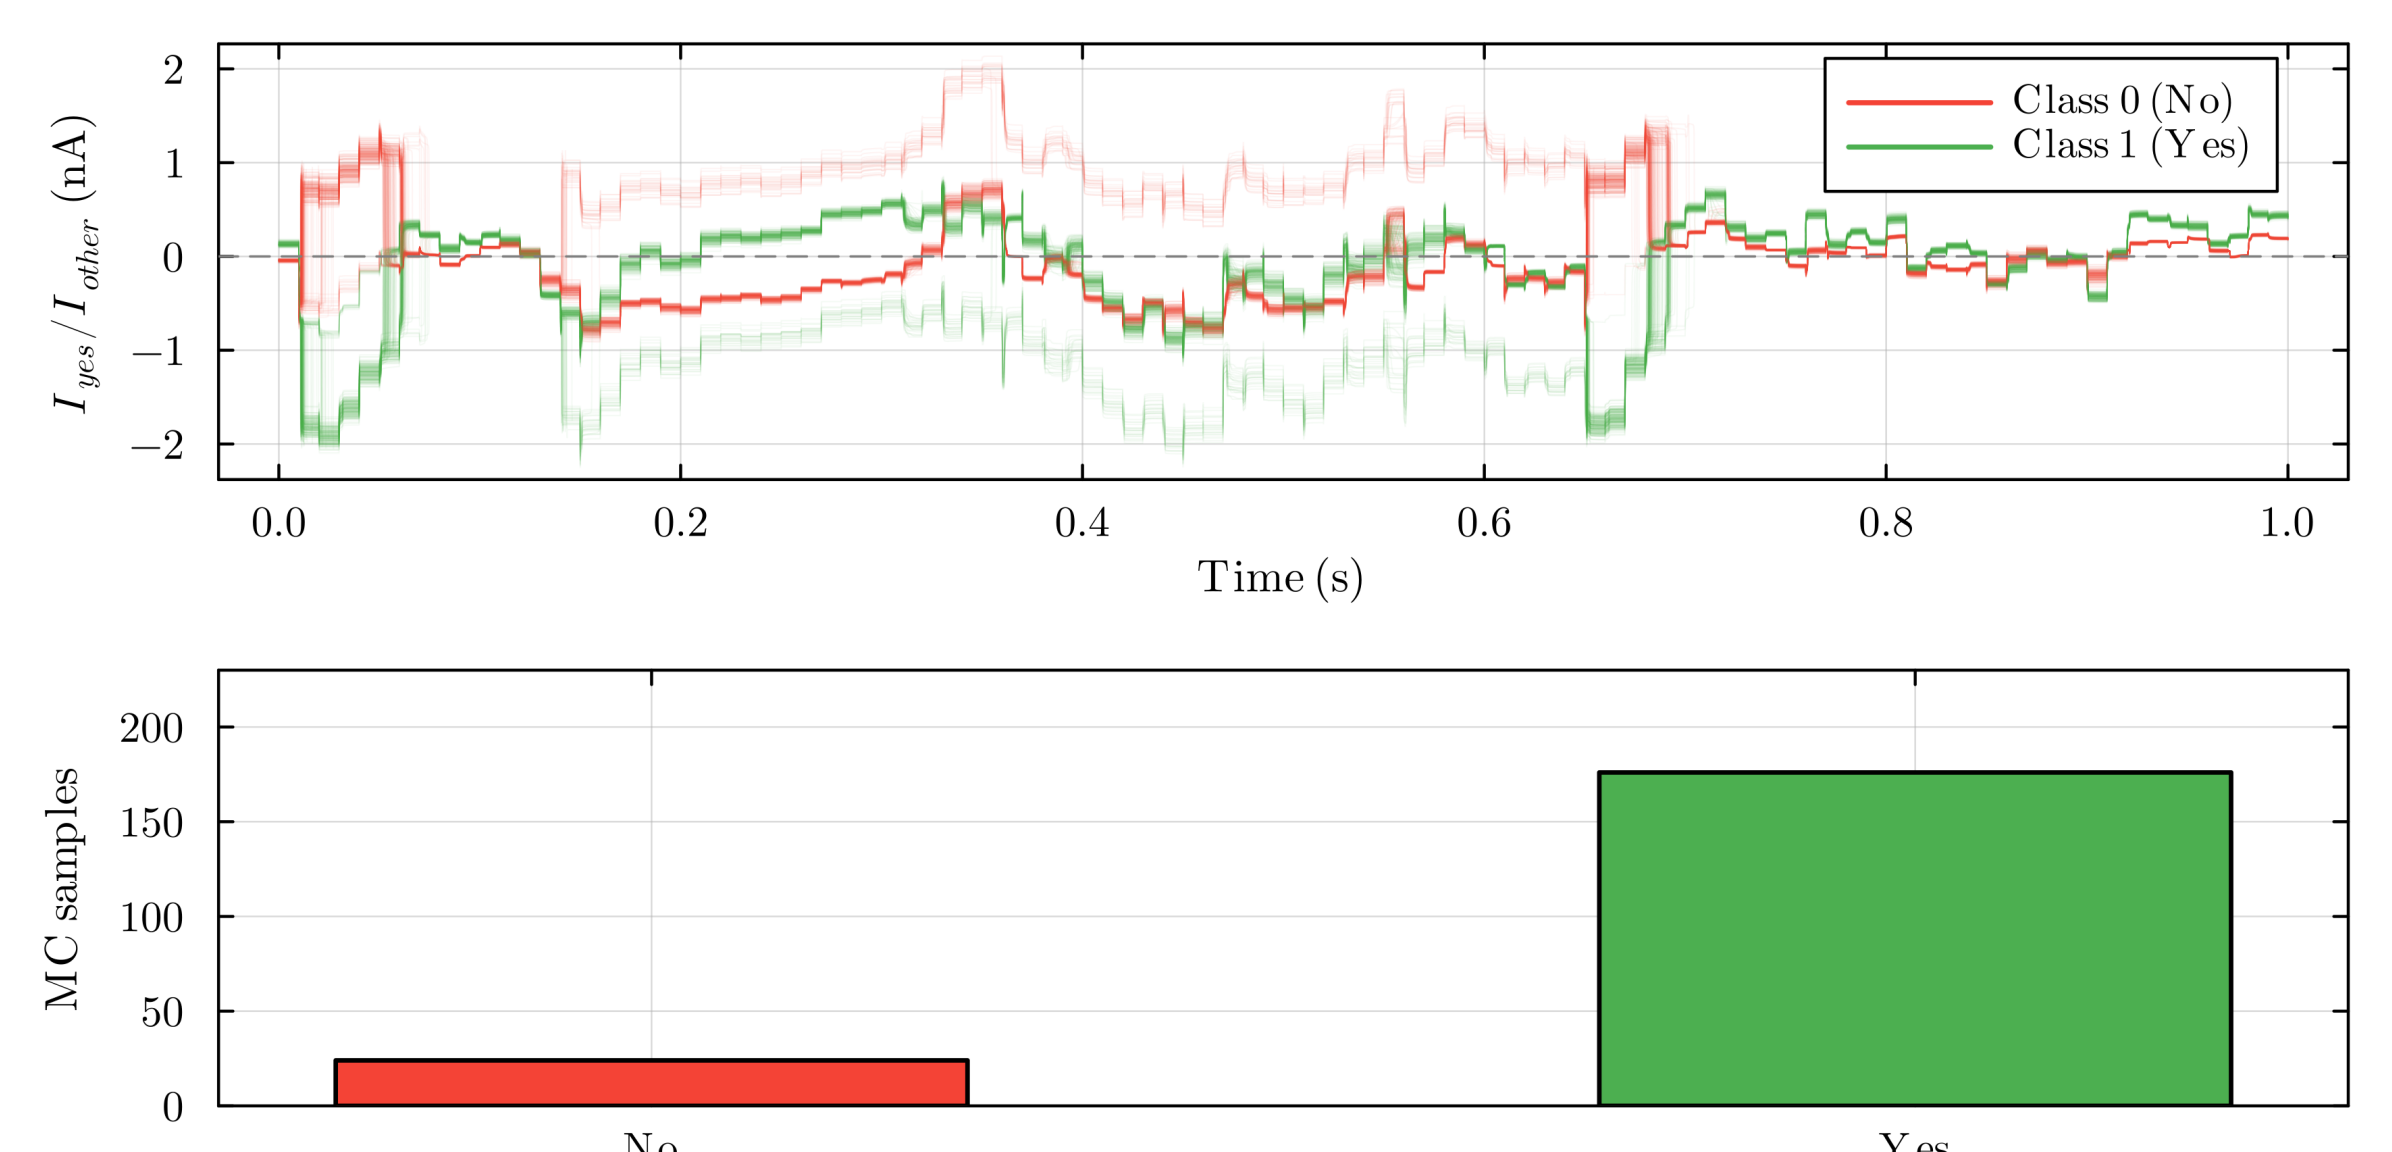
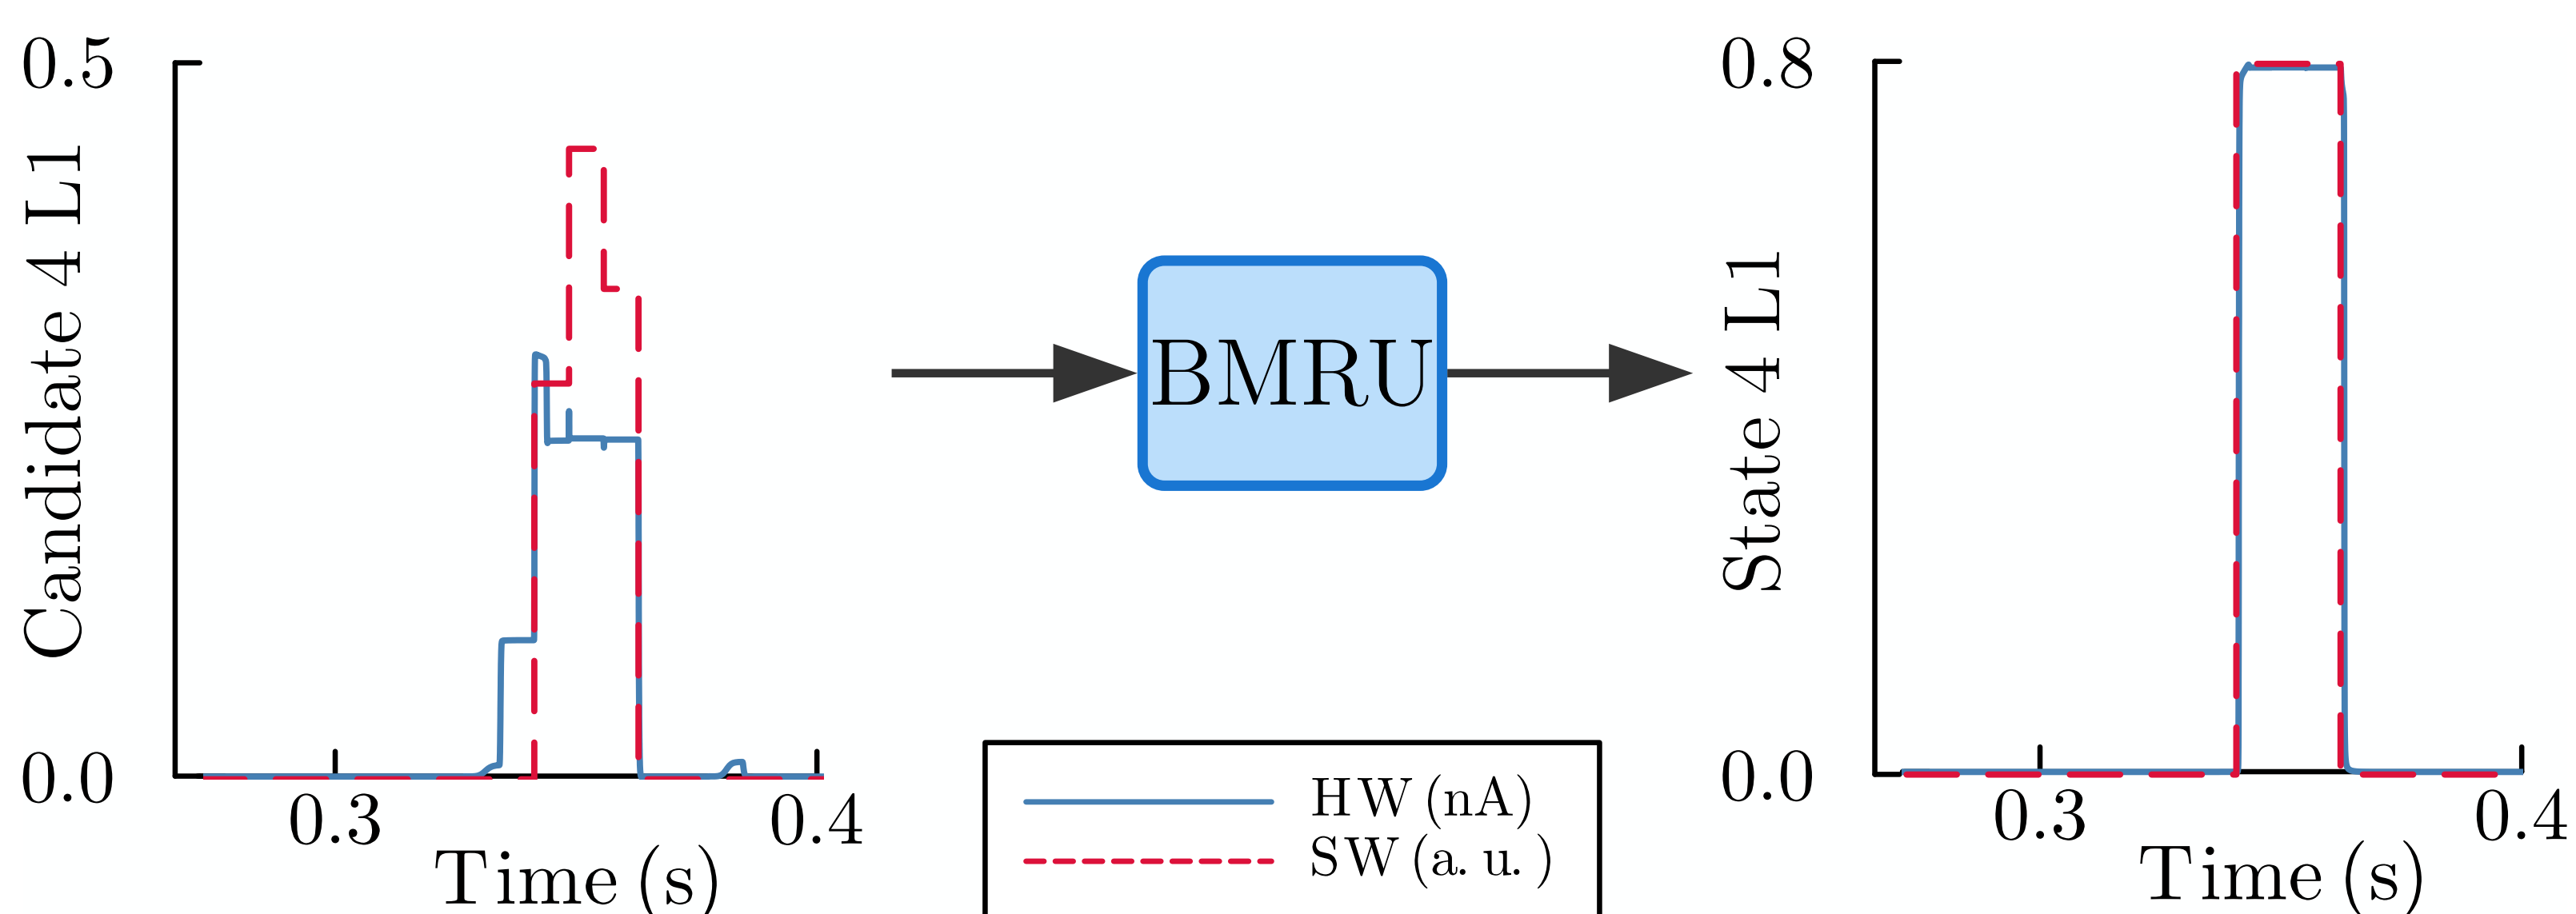
Two Heaviside stages in feedback form a **9-transistor tunable Schmitt trigger** whose threshold (I_{thresh}), output level (I_{gain}) and hysteresis width (I_{width}) are set independently by picoamp biases.

A CO-DESIGN ECOSYSTEM





Every block is a CMOS primitive: weights become current mirrors, ReLU a diode, each cell the tunable Schmitt trigger, so the trained model is the chip.

Result



Transition-level simulations confirm that the design is **highly robust to analog mismatch**. The software and hardware behaviors show **very strong agreement**, and the **quantization** prevents noise from accumulating across recurrent steps.

What this opens up



Reinforcement learning. Multistability, the very property this cell is built on, lets memory-based agents trained on short horizons stay optimal at far longer ones.

Cumulative memory. The same idea extends from two states to many: richer on-chip memory that rivals parallelizable RNNs and SSMs. (ICML 2026 Spotlight).