nature nanotechnology

Article

https://doi.org/10.1038/s41565-026-02311-0

Integration of Mott neurons and spintronic memories for event-driven spiking neural networks

Received: 29 August 2025

Accepted: 23 September 2026

Published online: 25 November 2026

Check for updates

Guanmin Li[1,2], Renata Lei[1,2], Ran Zhang[1], Simon Arnold[1], Yuliang Chen[1], Lukas Fischer[1], Andrea Migliorini[1], Zihan Yin[1], Yung-Cheng Li[1], Jae-Chun Jeon[1] & Stuart S. P. Parkin[1]

Spiking neural networks provide energy-efficient, event-driven computing, but separating volatile neuronal dynamics from non-volatile memory imposes costly signal transfer and conversion. Here we demonstrate a hybrid compute-and-store core that directly integrates thermally coupled vanadium dioxide Mott leaky integrate-and-fire neurons to spintronic memory. The neurons require no external resistive or capacitive elements for temporal integration and write their output spikes into either domain-wall racetracks or spin–orbit torque magnetic tunnel junctions without amplifiers or analogue-to-digital converters. In racetracks, spike number and amplitude are converted into domain-wall displacement, spatially preserving temporal histories and enabling up to six spikes per processing cycle. Distinct spike trains generated from encoded input patterns produce separable non-volatile states, demonstrating event-based processing and read-out. A network model parameterized by experimental leaky integrate-and-fire neuron dynamics achieves 96.3% accuracy on MNIST digit recognition. Functional coupling of a 500-nm leaky integrate-and-fire neuron to a 100-nm spin–orbit torque magnetic tunnel junction further establishes nanoscale compatibility. The demonstrated compute-and-store core consumes 51.9 nJ per spike-to-memory conversion, with sub-200-nm vanadium dioxide devices indicating operation near 200 pJ, offering a scalable route towards dense, low-power edge neuromorphic systems for real-time sensory processing.

Spiking neural networks (SNNs) mimic biological neurons by communicating via discrete action potentials, offering energy-efficient, event-driven computation[1–4]. A typical structure of a spiking neuron that consists of dendrites, soma and axon is shown in Fig. 1a. As depicted in Fig. 1b, inputs arrive as spikes at the dendrites; the soma integrates these inputs into a membrane potential that leaks over time (Fig. 1b); when the potential exceeds a threshold, the neuron fires a spike along the axon and then enters a brief refractory period. These leaky integrate-and-fire (LIF) dynamics impart short-term memory, which distinguishes neurons in SNNs from conventional models that function merely as memoryless activation units.

Mott memristors, with their volatile memory and threshold switching characteristics, can emulate the LIF neuron behaviour[5,6]. So far, however, most implementations still rely on external passive components (for example, capacitors and resistors) to allow for temporal integration and leakage[7–11], which consequently increases the system footprint and

[1]Max Planck Institute of Microstructure Physics, Halle (Saale), Germany. [2]These authors contributed equally: Guanmin Li, Renata Lei. e-mail: jae-chun.jeon@mpi-halle.mpg.de; stuart.parkin@mpi-halle.mpg.de

https://doi.org/10.1038/s41565-026-02311-0

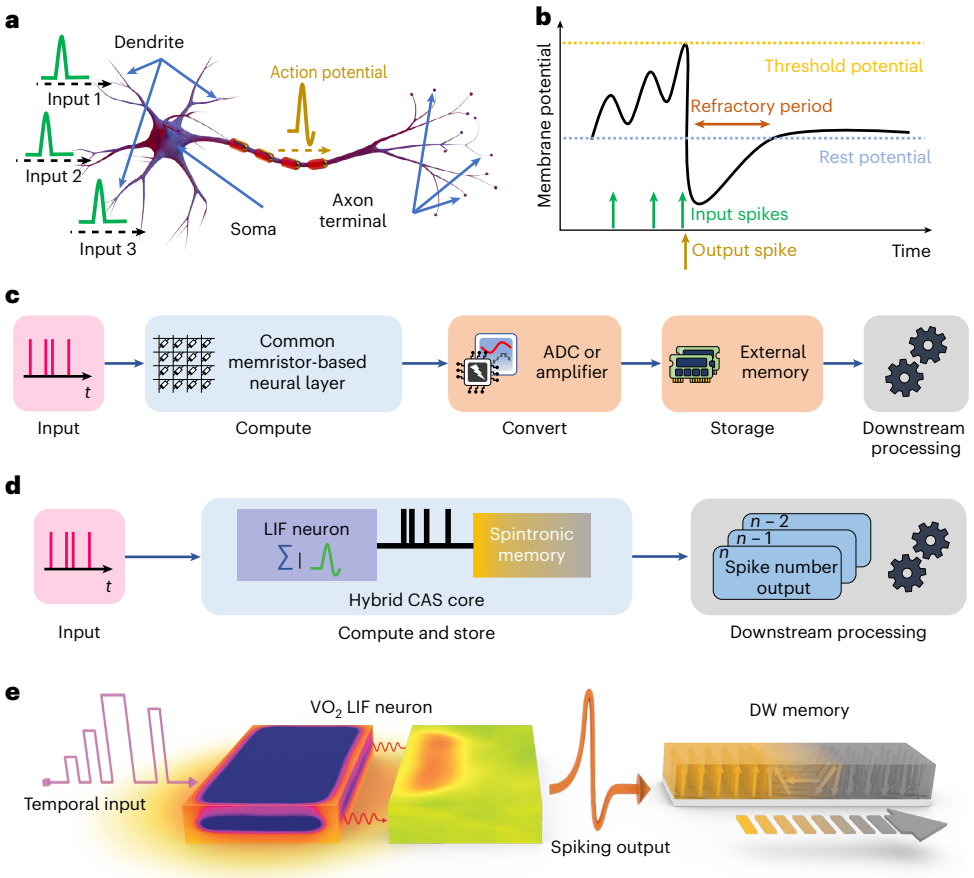


Fig. 1 | Hybrid CAS unit for spike-based computing architecture. a, Typical structure of a spiking neuron, including dendrites, soma and axon. **b**, A LIF neuron that integrates incoming spikes and builds up a membrane potential. Then, the neuron fires a spike and undergoes a refractory period until it returns to its rest state. **c**, Common memristor-based neural networks with separate compute units and memories that save the neuron outputs for downstream processing (that is, forward propagation/learning). ADC, analogue-to-digital converter. **d**, Spike-based computing architecture with a hybrid Mott-spintronic CAS unit. The CAS unit will compute and then directly store output spikes within the spintronic memories. The spike history of the firing neuron can be acquired after each process cycle *n*. **e**, A Mott-spintronic CAS unit consists of a VO_2 LIF neuron that processes the temporal input and a DW spintronic memory that acquires the spike history of the firing neuron.

impedes both scalability and tunability. Moreover, in memristor-based neural networks, amplifiers and analogue-to-digital converters are frequently required[12–14] to condition the neuron's analogue outputs before writing to off-chip memory for learning rules such as gradient descent or spike-timing-dependent plasticity[15–19], which adds additional energy overhead (Fig. 1c). Traditional complementary-metal–oxide–semiconductor-based SNNs further rely on asynchronous address-event representation, which serializes spike communication and limits throughput to one spike per cycle[20–23]. To overcome these bottlenecks, spintronic memory presents an ideal platform due to its low energy cost per operation and non-volatility[24,25]. Among various spintronic devices, racetrack memory, whose functionality relies on the synchronized motion of the domain walls (DWs) by nanosecond-long current pulses, offers a unique capability: temporal input, that is, a series of pulses, can be spatially stored as DW positions, owing to the near-linear response of DW displacement to either the input pulse number or amplitude[26–28].

In this work, we realize a hybrid compute-and-store (CAS) core by directly integrating Mott memristor (vanadium dioxide (VO_2))-based LIF neurons and spintronic memory (Fig. 1d) on the same substrate—a critical step towards all-in-memory SNN architectures. Our LIF neurons are implemented using two thermally coupled VO_2 memristors, without the need for external passive components and inherently supporting event-driven operation based on temporal dynamics. First, we show that VO_2 LIF neurons can directly write bits into DW-based racetracks without any analogue-to-digital converters or amplifiers. Unlike address-event representation, which handles only one spike per cycle, our Mott-spintronic CAS core can record multiple spikes per process cycle (six spikes per cycle). The serially generated spikes from the VO_2 LIF neuron shift the DW within the racetrack device (Fig. 1e), and the spatial evolution of DW position is proportional to the number of identical spikes—that is, the spike count is stored as the DW position. When the spike amplitudes differ (for example, because of differing neuron recovery times), distinct temporal sequences produce different net DW displacements. In this case, the memory implements a combinatorial read-out: varying amplitudes and timings map onto unique DW positions, thereby preserving temporal information of the input spike train. We demonstrate the reading of spike histories on two scales: first, via Kerr microscopy on a micrometre-scale racetrack device directly integrated with VO_2 neurons on the same chip; second, through the electrical read-out of VO_2 spikes using 100-nm spin–orbit torque (SOT) magnetic tunnel junctions (MTJs), thereby, showcasing our platform's scalability and flexibility.

Temporal dynamics and computing capability of VO_2 neurons

To construct the LIF neuron, thermally coupled VO_2 Mott memristors are utilized (Supplementary Fig. 1 shows the details of film

https://doi.org/10.1038/s41565-026-02311-0

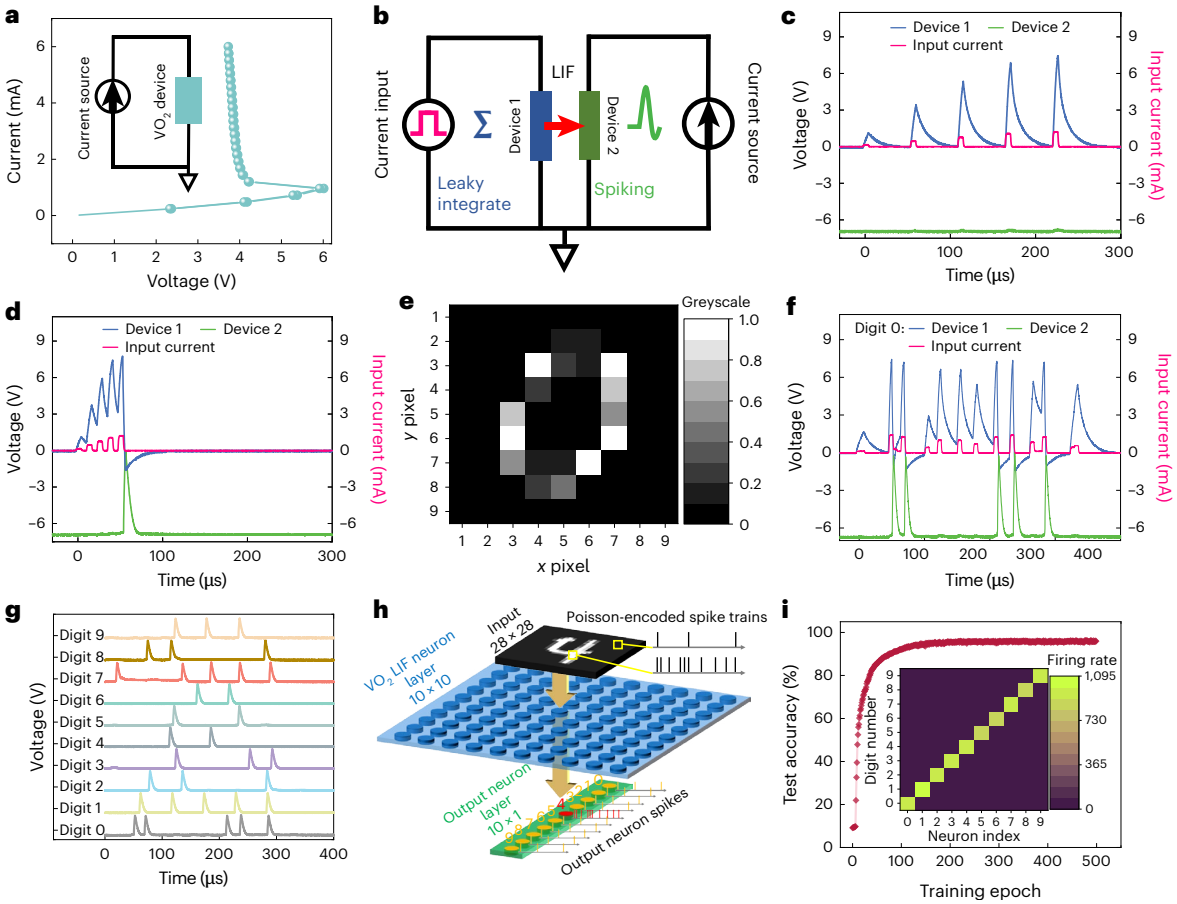


Fig. 2 | VO_2 LIF neuron with temporal dynamics and computing capability. a, *I*–*V* measurement (current is varied and voltage is measured) of a single VO_2 device ($5 \times 5\ \mu m^2$). Inset: diagram of a single device supplied with a constant current. **b**, Circuit diagram for a LIF neuron realized by two thermally coupled VO_2 Mott memristors. **c**, Experimental demonstration of the leaky integration (short-term memory) of the VO_2 LIF neuron. **d**, Experimental demonstration of the triggered spiking of the VO_2 LIF neuron, when the interval of the pulse train becomes shorter and the neuron reaches its threshold firing voltage. **e**, A handwritten digit 0 in a 9×9-pixel^2 array converted from the MNIST dataset. **f**, Neuron potential (blue line) and output spikes (green line) from a VO_2 LIF neuron with the input of digit 0. **g**, Different spiking patterns of a VO_2 LIF neuron for handwritten digits 0 to 9. **h**, SNN architecture based on 100 VO_2 LIF neurons for MNIST handwritten digit recognition. **i**, Test accuracy reaches 96.3% for the digit recognition after 300 epochs of training. Inset: the confusion matrix.

properties). VO_2 is a type of metal–insulator transition (MIT) material that shows substantial resistivity change at around 340 K due to the Mott transition[29–32]. In our VO_2 Mott memristor device ($5\ \mu m \times 5\ \mu m$), when the current (I) is gradually increased, the voltage starts to drop at $I \approx 1.5$ mA (threshold current), exhibiting negative differential resistance (Fig. 2a). In the current range between 1.5 mA and 5.0 mA, the voltage across the device oscillates with a d.c. current, where the MIT is stable beyond 10^6 switching cycles (Supplementary Fig. 2). At higher current values (>5 mA), the device becomes fully metallic and no longer oscillates. We confirm that the VO_2 Mott memristor device for current-driven self-oscillation can be scaled down to 100 nm (Supplementary Fig. 3). When two VO_2 devices are closely spaced, LIF neuron functionalities can be realized. As schematically illustrated in Fig. 2b, we fabricated two 5×5-μm^2 VO_2 memristors with a 2-μm spacing: device 2 (green) is set at a constant current at the subthreshold value ($I_2 = -0.75$ mA) and acts as a trigger for firing spikes. Device 1 (blue) receives input spike currents and serves as the leaky integrate component. In Fig. 2c, a series of gradually increasing current pulses (duration of 6 μs, from 0.16 to 1.22 mA, which is lower than the threshold switching value for device 1) with a 47-μs interval are injected into device 1. Note that between each input pulse, the potential of device 1 relaxes back to 0 V without reaching the switching voltage, which corresponds to leaky integration behaviour. The potential of device 2 stays at the rest level (−7 V). Next, a same series of pulses but with a smaller interval (6 μs) is sent to device 1 (Fig. 2d). The potential of device 1 gradually builds up and finally reaches the threshold switching value, undergoes a phase transition to a low resistance state, thermally triggers device 2 to 'fire' a spike (phase transition), and then both devices recover back to their rest level, that is, 0 V and −7 V for devices 1 and 2, respectively. The scalability of the device is further demonstrated by a nanoscopic VO_2 LIF neuron (500×500 nm^2) (Supplementary Fig. 4). Note that the thermal crosstalk becomes negligible at a centre-to-centre distance of twice the device's lateral size, as shown in our previous work[32]. Furthermore, the well-established '3W principle' in electronic design states that a spacing of three times the linewidth ensures minimal electric-field crosstalk between the lines. Therefore, a centre-to-centre pitch of three times the device width between neighbouring VO_2 LIF neuron units provides a safe, scalable design rule.

Next, we demonstrate the temporal dynamics of a VO_2 LIF neuron for its computing capability. First, each digit from 0 to 9 in the Modified National Institute of Standards and Technology (MNIST) handwritten dataset[33] is converted into a 9×9-pixel^2 array. Figure 2e shows an example of a handwritten digit 0. Each pixel is transformed into the current value proportional to its greyscale—that is, for pixel with greyscale value of 0.8, the current pulse amplitude is 0.8×1.5 mA (pulse duration is kept constant at 6 μs). The total encoding window for the pulse train is set to be 400 μs. After that, the pulse train is injected into the VO_2 LIF and triggers it to fire temporal spikes (Fig. 2f). The neuron

spike patterns for handwritten digit 0 to 9 are shown in Fig. 2g (detailed waveforms are shown in Supplementary Fig. 5). The experimental data are then fitted by the well-established RC model[34] to extract the leakage time constant τ_{leak} for the software reconstruction of a large-scale SNN (Supplementary Fig. 6). The network with 100 VO_2 LIF neurons (Fig. 2h) can reach 96.3% precision for the MNIST handwritten digit recognition task after 300 epochs of training (Fig. 2i). A 10 × 10 crossbar array for the VO_2 LIF neuron computing layer is fabricated to prove the feasibility of manufacturing these units at a technologically relevant scale (experimental details are shown in Supplementary Figs. 7 and 8).

Additionally, the rich temporal dynamics of multiple thermally coupled VO_2 neurons enable complex, timing-dependent computational functions (Supplementary Figs. 9–15). Specifically, the demonstrated capabilities for temporal integration and modulation (Supplementary Figs. 9 and 10) allow a neuron to act as a temporal filter and coincidence detector[35,36]. This empowers the network to discriminate spatiotemporal patterns based on the precise timing and rate of input spikes, which is crucial for processing dynamic sensory data like audio or video streams. Furthermore, the resonate-and-fire behaviour based on frequency selection and phase locking (Supplementary Figs. 11–15) enables neurons to function as feature-specific detectors for rhythmic patterns[37–39], without the need for complex digital circuits or algorithms. These are not merely device-level phenomena but constitute a set of fundamental building blocks for advanced system-level processing.

Demonstration of a hybrid CAS core

Now, we demonstrate the operation of a Mott-spintronic CAS core that directly integrates VO_2 LIF neurons and racetrack devices. To realize such directly integrated hybrid device, we fabricated VO_2 neurons and racetrack memory devices on the same substrate (sapphire) using ion-beam etching and lift-off lithography techniques. Detailed information about the film, device and fabrication process is provided in Methods, Supplementary Figs. 16–23 and Extended Data Figs. 1–3. Furthermore, to perform computational tasks in a stand-alone environment, we integrate the hybrid CAS core with printed connector and control units. The photograph of the test system for the hybrid CAS core is shown in Fig. 3a. It consists of a temporal input (a programmable source unit, namely, a 20-bit digital-to-analogue converter (DAC)), a Mott-spintronic CAS core (VO_2 LIF neuron + racetrack memory), a sensor for detecting the DW position (Kerr microscopy), a reset unit for initializing spintronic memories and a clock that governs the execution order of the components (DAC unit). The operation of the CAS core (Fig. 3b) is explained as follows: at the beginning of a process cycle (n), the temporal data, encoded as a pulse train, is sent to the CAS core by the input unit. After the data have been processed by the VO_2 LIF neuron, the output spikes are sequentially transferred into the racetrack device in which a DW is positioned. The VO_2 spikes are converted into a spin current via the spin Hall effect within the heavy metal layer (Pt) of the racetrack, which then exerts an SOT on the DW, thereby inducing its motion[40–44]. The speed of the current-induced DW motion in our device is plotted as a function of pulse length (ranging from 20 ns to 1,500 ns) and amplitude (from 0.75 V to 3 V) (Supplementary Fig. 18). The DW position can be detected by many different methods (for example, Kerr microscopy, anomalous Hall signal or tunnelling magnetoresistance) to retrieve the spike history during the process. At the end, for the next process cycle (n + 1), the DW is reset to the same initial position by a series of nanosecond pulses. Besides, the same reset function could also be achieved using a self-reset mechanism, such as a magnetic inhibitor or a geometrically engineered device[45,46].

In Fig. 3c, we first demonstrate a VO_2 LIF neuron directly connected to the DW racetrack memory (2 × 15 μm^2) on the same substrate through a high-pass rectifier composed of a capacitor and diode, thereby eliminating additional power consumption from the analogue-to-digital converter and/or operational amplifier. The LIF neuron consists of two identical VO_2 devices (both 5 × 5 μm^2, spaced by 2 μm): one acts as the leaky integrator and the other one is the trigger for firing spikes (Fig. 2b). The unique characteristic of racetrack memory, where the total DW travelled distance depends linearly on the number of injected pulses, makes it ideal for storing the temporal spike information from the VO_2 LIF neuron. Figure 3d shows the pulse pattern of a single spike from the neuron to DW racetrack memory. The blue curve shows the potential of VO_2 device 1, whereas the green curve presents the potential of VO_2 device 2, thermally triggered by device 1. During this process, the current in device 2 (I_2) is fixed at −0.75 mA, a subthreshold current for MIT switching. The recovery time of VO_2 device 2 becomes longer after the phase transition due to the charging of the capacitor illustrated in Fig. 3c. The yellow line in Fig. 3d shows the spike received by the racetrack memory after the high-pass rectifier. The DW position within the racetrack linearly increases with the number of identical spikes, that is, ~1 μm per pulse (Fig. 3e). In this particular racetrack, up to six spikes generated from the VO_2 LIF neuron can be stored in each process cycle. The number of stored spikes can be increased by tuning the physical characteristics of the racetrack device, such as its length, anisotropy, edge roughness and so on. The differential Kerr microscopy images for the final DW position with each total spike number are shown in Fig. 3f. Note that our recent demonstration has revealed that racetrack memory devices only 50 nm wide exhibit similar functions as micrometre-sized devices[27]. This indicates that nanoscopic racetrack memory could readily be integrated into our Mott-spintronic CAS core. However, integrating read sensors, such as MTJs or Hall bars, with a high on–off ratio remains a key challenge.

Functional coupling in a nanoscale CAS core

To further demonstrate the scalability and versatility of the CAS architecture, we implemented magnetization switching induced by VO_2 LIF neuron spikes in a nanoscopic 100 × 400-nm^2 SOT-MTJ device. The device's magnetic properties, including the magnetization (M–H), resistance (R–H) hysteresis curves and current-induced switching characteristics are presented in Supplementary Figs. 20–22. The structural design and layer stack of the SOT-MTJ are illustrated in Fig. 3g. Repetitive switching of the SOT-MTJ by VO_2-generated spikes is shown in Fig. 3h, with detailed electrical connections provided in Supplementary Fig. 23. We demonstrate effective coupling between a VO_2 LIF neuron (500 × 500 nm^2) and an SOT-MTJ (100 × 400 nm^2) prepared on different substrates, thereby achieving functional operation at the nanoscale. The corresponding circuit configuration and operation are described in Extended Data Fig. 1. The successful switching of the nanoscale SOT-MTJ by VO_2 spikes underscores the versatility of our hybrid Mott-spintronic platform. Although the DW racetrack serves as a high-precision, analogue temporal buffer for encoding the spike history, the SOT-MTJ operates as a dense, binary event detector. This distinction reflects a hierarchical memory architecture suited for different computational roles within a full system: racetrack memory is ideal for tasks requiring fine-grained temporal history (multiple spikes read-out), whereas SOT-MTJs are directly applicable as binary memory elements for generating decision signals from an output neuron. This compatibility with diverse spintronic technologies highlights both scalability and architectural flexibility of our CAS core.

Demonstration of event-based processing and read-out

Next, we demonstrate the capability of the Mott-spintronic CAS core to process temporal inputs and generate a combinatorial read-out. Note that this CAS core directly integrates a VO_2 neuron and a racetrack memory on the same substrate (Extended Data Figs. 2 and 3). Due to the passive capacitor element, VO_2 device 2 undergoes a certain charging process after every time it is thermally triggered to fire a spike by device 1, as explained in the previous section. This charging process prevents device 2 from firing a second output spike right after the first spike

https://doi.org/10.1038/s41565-026-02311-0

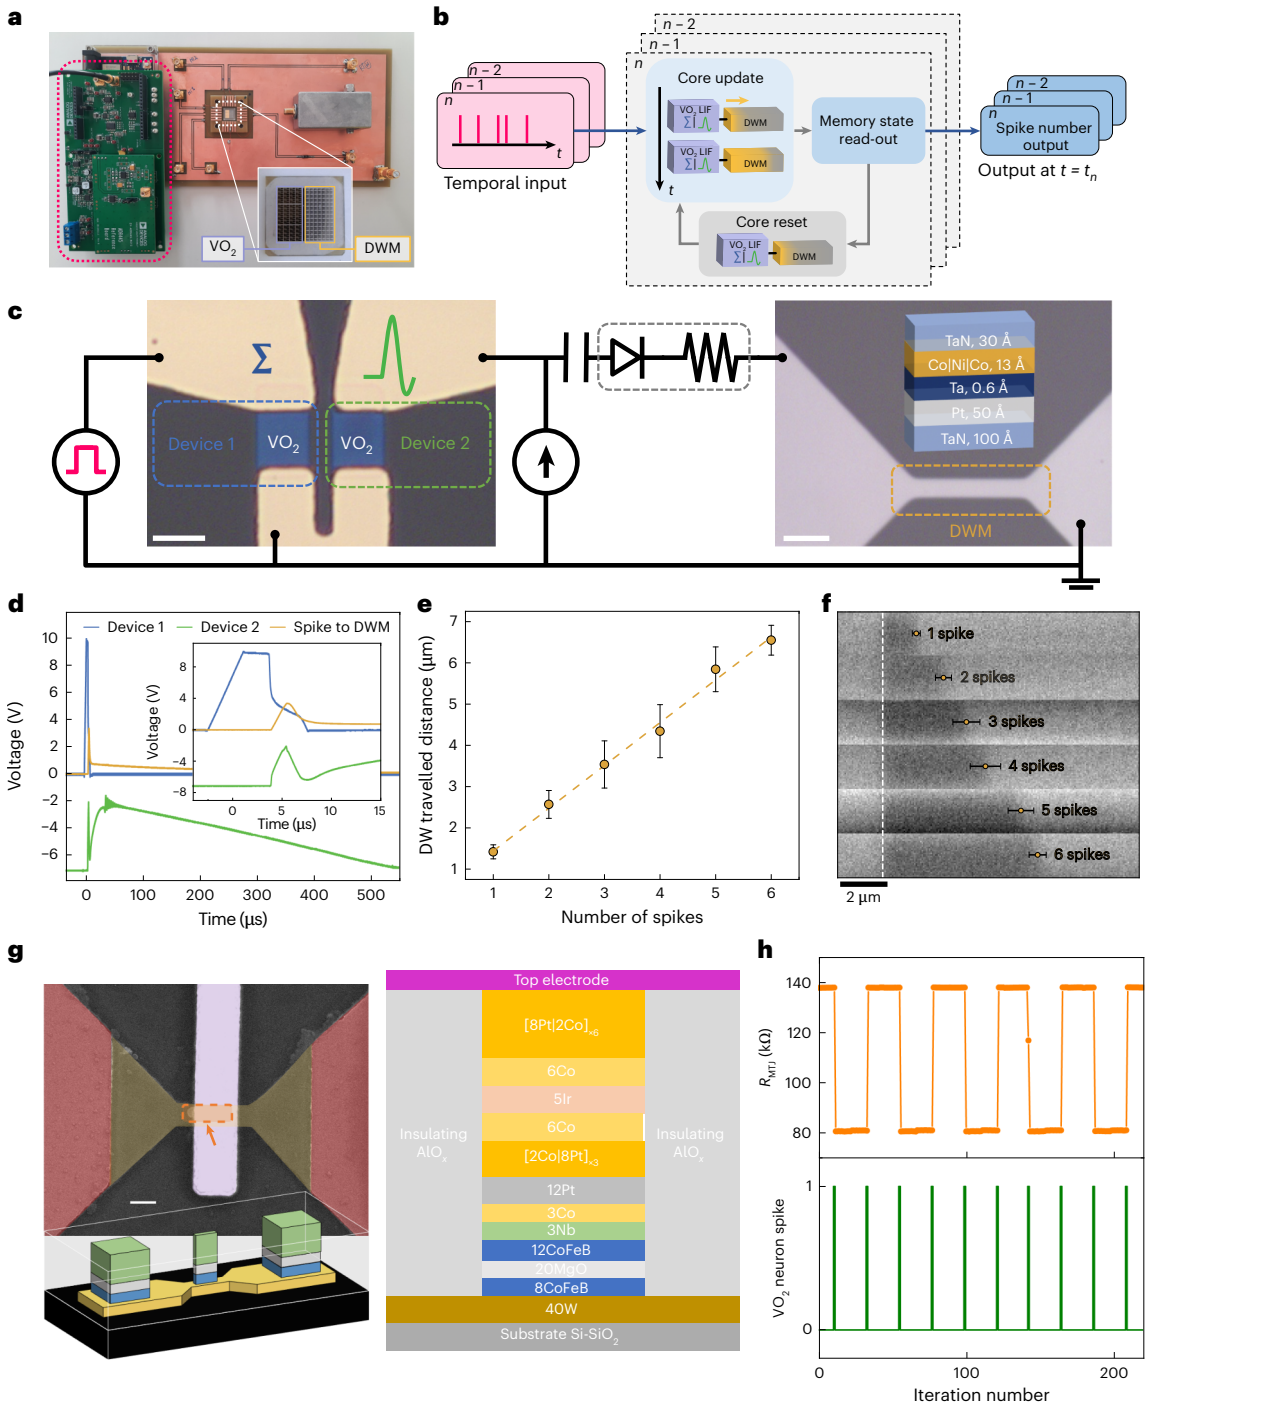


Fig. 3 | Operation of a hybrid CAS unit. a, Photograph of the integrated test setup. In the setup, a directly integrated CAS element, that is, a VO_2 + racetrack device fabricated on the same sapphire substrate (inset contains a magnified view of the chip), a high-precision 20-bit source unit (dotted pink outline) and passive circuit elements (capacitor and diode) are integrated onto a printed electrical connector board. **b**, Operation flowchart of the hybrid Mott-spintronic CAS unit. **c**, Optical microscopy image of a VO_2 LIF neuron and a DW memory (DWM) with the measurement circuit diagram. Scale bars, 5 μm. Note that in the setup, we added an external resistor in between the VO_2 neuron and the racetrack device since the spike amplitude was too large for fine DW motion. The additional resistor (3 kΩΩ) drops the spike potential. **d**, Potential of a VO_2 LIF neuron for generating a spike to move a DW. **e**, DW travelled distance as a function of the injected spikes generated from a VO_2 LIF neuron. Data show the mean and standard deviation across ten repetitions. **f**, Differential Kerr microscopy images of DW motion in response to 1–6 spikes. The white line indicates the starting position. Yellow dots and error bars indicate the average final DW position and standard deviation across ten repetitions, respectively. **g**, Left, top: SEM image of a 100 × 400-nm^2 SOT-MTJ device (yellow, SOT layer; red, bottom contact; purple, top contact; the dashed box shows the MTJ pillar). Scale bar (white), 200 nm. Left, bottom: schematic of the device structure (green, top reference layer; grey, MgO layer (tunnel barrier); blue, free layer; yellow, spin Hall layer (W); black, substrate). Right: device layer stack (thicknesses in ångströms). **h**, Preserving spike history and electrical read-out by SOT-MTJ. The green curve shows the output spikes from the VO_2 LIF neuron, whereas the orange curve corresponds to the resulting resistance (R_{MTJ}) switching in the SOT-MTJ by the spikes.

https://doi.org/10.1038/s41565-026-02311-0

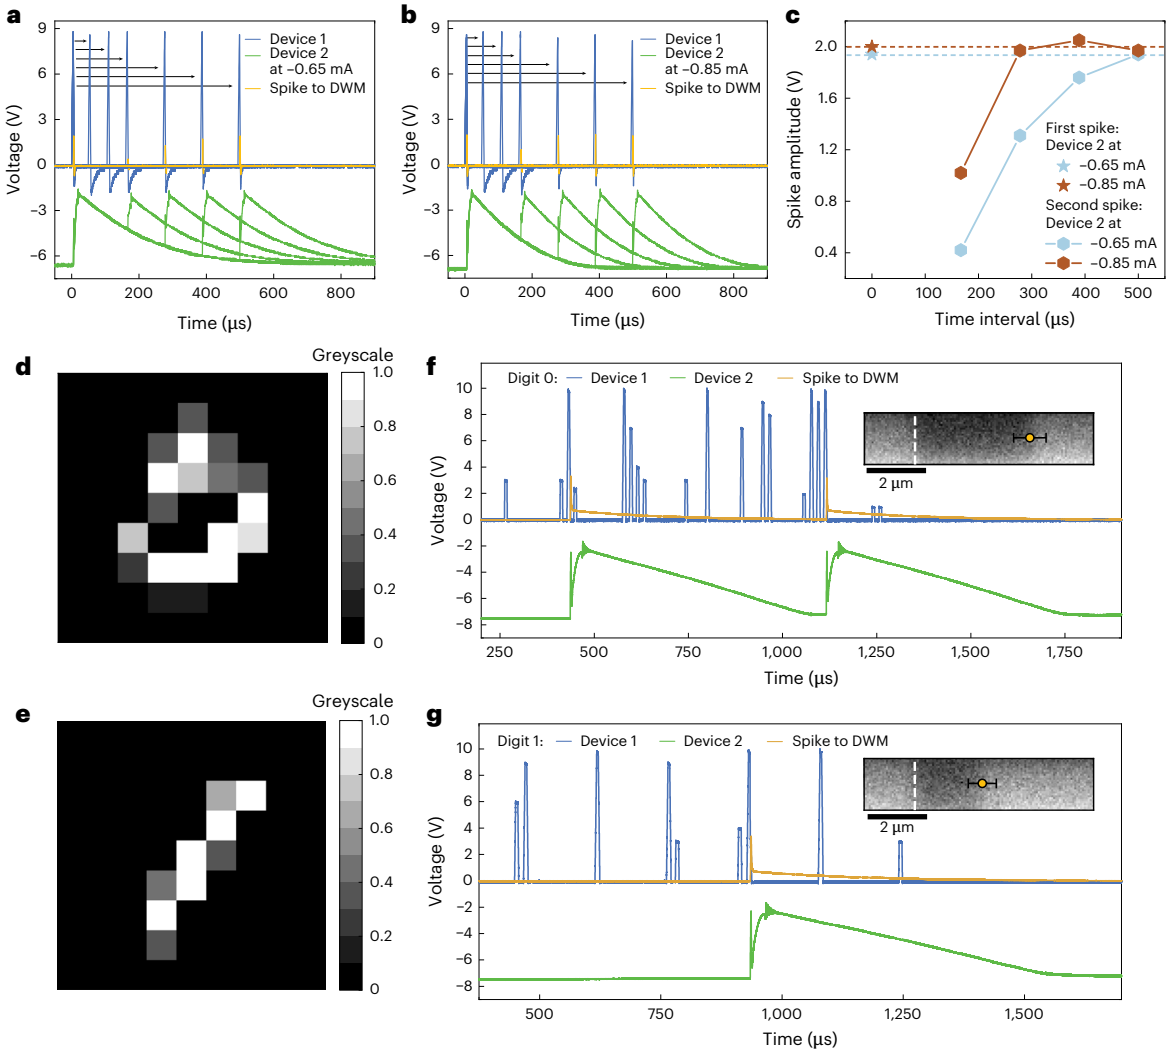


Fig. 4 | Spike-based information processing and store by CAS core. **a**,**b**, Refractory period and recovery effect of a VO_2 LIF neuron: device 2 at a bias of −0.65 mA (**a**) and −0.85 mA (**b**). Black arrows indicate different time intervals between the first and second stimuli. **c**, Comparison of the recovery speed of the VO_2 LIF neuron at different biases (device 2 at −0.65 mA and −0.85 mA). **d**, Digit 0 in a 9 × 9-pixel^2 array. **e**, Digit '1' in a 9 × 9-pixel^2 array. **f**,**g**, Potential of the VO_2 LIF neuron and the corresponding output spike to the DWM, in response to input digit 0 (**f**) and 1 (**g**). Top right inset: differential Kerr microscopy images for digit 0 and digit 1 (device 2 at −0.7 mA). The white dashed line indicates the starting position. The yellow dot and error bar show the average DW position and standard deviation across 20 repetitions, respectively.

(Fig. 4a). The duration of the refractory period—defined as the minimal interval required for a second spike to recover its full amplitude—can be directly and reliably controlled by adjusting the bias current (I_2) on VO_2 device 2. With increasing interval between the first and second stimuli, the amplitude of the second output spike gradually increases to the same level as the first output spike (fired at 0 μs; Fig. 4c). When VO_2 device 2 is set at a higher bias level (−0.85 mA), the neuron is more sensitive and recovers faster (Fig. 4b). The amplitude of the second output spike increases back to the level of the first spike more rapidly (it takes 500 μs for −0.65-mA bias, and 277 μs for −0.85-mA bias). Thus, the output spikes emitted from the VO_2 LIF neurons at different times with varying amplitudes incorporate a temporal correlation between the input stimuli. Additionally, we demonstrate a synapse-free modulation of the VO_2 LIF neurons to stimuli by simply changing the bias of device 2 (Supplementary Fig. 24).

As shown in Fig. 4d,e, digits 0 and 1 represented in 9 × 9-pixel^2 arrays in the greyscale are encoded with the same method introduced in the previous section: each pixel is transformed into a current pulse with the amplitude proportional to the pixel greyscale. The pulse width is set to 9 μs and the period, to 19.5 μs. The current pulse trains converted from the image are injected into the directly integrated CAS core. Figure 4f,g shows the potential of the VO_2 LIF neuron (device 2 fixed at −0.7 mA) and output spikes for different input digits. The LIF neuron emits a spike with a 3.3-V amplitude at 936 μs for digit 1. For digit 0, the neuron emits two spikes with an amplitude of 3.3 V and 3.1 V at 438 μs and 1,118 μs, respectively. Each spike train is repeated two times and the final positions of the DW are then read out by the Kerr microscope (top right inset). As shown, the distinct spike patterns for 0 and 1 result in different DW displacements, enabling hardware-level distinction between the input patterns and validating the hybrid CAS core as a platform for event-based processing and read-out. Testing the same task across three distinct device sets produced qualitatively consistent results (Extended Data Fig. 4). For more complex letters such as 'M', 'P' and 'G' (Max-Planck-Gesellschaft), the system is also able to generate different outputs (Supplementary Fig. 25 shows the different number of spikes in different process cycles). Furthermore, we perform system-level software evaluation of a 16 × 16 1T-CAS array to identify M, P and G (each 1T-CAS unit cell consists of a metal–oxide–semiconductor field-effect transistor, VO_2 neuron and racetrack memory (Extended Data Fig. 5 and Supplementary Video 1 and code of ref. 47)).

Conclusion

In this work, we present a hybrid CAS platform that integrates VO_2-based LIF neurons with spintronic memories for event-driven neuromorphic computing. Our thermally coupled Mott neurons interface effectively with DW-based racetrack memory, enabling direct spike writing. This integration allows the analogue temporal dynamics of neurons to be spatially mapped into the non-volatile states of spintronic devices. As a result, our Mott-spintronic architecture achieves a substantial reduction in energy per spike-to-memory event for cognitive computing tasks compared with existing systems (Supplementary Table 1), and supporting multispike processing beyond the throughput limits of address-event representation schemes. Furthermore, we have shown nanoscale compatibility by functional coupling between VO_2 neurons and SOT-MTJs, highlighting the flexibility and scalability of the platform. Our CAS core promises dense, low-power neuromorphic hardware, providing a critical foundation for real-time sensory processing and energy-efficient edge intelligence.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41565-026-02311-0.

References

1. Yamazaki, K., Vo-Ho, V.-K., Bulsara, D. & Le, N. Spiking neural networks and their applications: a review. *Brain Sci.* **12**, 863 (2022).
2. Kim, R., Li, Y. & Sejnowski, T. J. Simple framework for constructing functional spiking recurrent neural networks. *Proc. Natl Acad. Sci. USA* **116**, 22811–22820 (2019).
3. Yin, B., Corradi, F. & Bohté, S. M. Accurate and efficient time-domain classification with adaptive spiking recurrent neural networks. *Nat. Mach. Intell.* **3**, 905–913 (2021).
4. Maass, W. Networks of spiking neurons: the third generation of neural network models. *Neural Netw.* **10**, 1659–1671 (1997).
5. Pickett, M., Medeiros-Ribeiro, G. & Williams, R. A scalable neuristor built with Mott memristors. *Nat. Mater.* **12**, 114–117 (2013).
6. del Valle, J., Salev, P., Kalcheim, Y. & Schuller, I. K. A caloritronics-based Mott neuristor. *Sci. Rep.* **10**, 4292 (2020).
7. Zhao, S. et al. A 1T1M programmable artificial spiking neuron via the integration of FeFET and NbO_x Mott memristor. *IEEE Electron Device Lett.* **45**, 1169–1172 (2024).
8. Liu, G. et al. Experimental demonstration of coplanar NbO_x Mott memristors for spiking neurons. *IEEE Electron Device Lett.* **45**, 708–711 (2024).
9. Xu, Z., Bernussi, A. A. & Fan, Z. Voltage pulse driven VO_2 volatile resistive transition devices as leaky integrate-and-fire artificial neurons. *Electronics* **11**, 516 (2022).
10. Yuan, R. et al. A neuromorphic physiological signal processing system based on VO_2 memristor for next-generation human–machine interface. *Nat. Commun.* **14**, 3695 (2023).
11. Qiu, E., Zhang, Y.-H., Di Ventra, M. & Schuller, I. K. Reconfigurable cascaded thermal neuristors for neuromorphic computing. *Adv. Mater.* **36**, 2306818 (2024).
12. Merrikh-Bayat, F. et al. High-performance mixed-signal neurocomputing with nanoscale floating-gate memory cell arrays. *IEEE Trans. Neural Netw. Learn. Syst.* **29**, 4782–4790 (2017).
13. Wang, Z. et al. Fully memristive neural networks for pattern classification with unsupervised learning. *Nat. Electron.* **1**, 137–145 (2018).
14. Duan, Q. et al. Spiking neurons with spatiotemporal dynamics and gain modulation for monolithically integrated memristive neural networks. *Nat. Commun.* **11**, 3399 (2020).
15. Diehl, P. U. & Cook, M. Unsupervised learning of digit recognition using spike-timing-dependent plasticity. *Front. Comput. Neurosci.* **9**, 99 (2015).
16. Cramer, B. Surrogate gradients for analog neuromorphic computing. *Proc. Natl Acad. Sci. USA* **119**, e2109194119 (2022).
17. Eshraghian, J. K. et al. Training spiking neural networks using lessons from deep learning. *Proc. IEEE* **111**, 1016–1054 (2023).
18. Masquelier, T., Guyonneau, R. & Thorpe, S. J. Competitive STDP-based spike pattern learning. *Neural Comput.* **21**, 1259–1276 (2009).
19. Aguirre, F. et al. Hardware implementation of memristor-based artificial neural networks. *Nat. Commun.* **15**, 1974 (2024).
20. Boahen, K. A. Point-to-point connectivity between neuromorphic chips using address events. *IEEE Trans. Circuits Syst. II* **47**, 416–434 (2000).
21. Davies, M. et al. Loihi: a neuromorphic manycore processor with on-chip learning. *IEEE Micro* **38**, 82–99 (2018).
22. Merolla, P. A. et al. A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* **345**, 668–673 (2014).
23. Roy, K., Jaiswal, A. & Panda, P. Towards spike-based machine intelligence with neuromorphic computing. *Nature* **575**, 607–617 (2019).
24. Liu, L. et al. Domain wall magnetic tunnel junction-based artificial synapses and neurons for all-spin neuromorphic hardware. *Nat. Commun.* **15**, 4534 (2024).
25. Wang, D. et al. Spintronic leaky-integrate-fire spiking neurons with self-reset and winner-takes-all for neuromorphic computing. *Nat. Commun.* **14**, 1068 (2023).
26. Hayashi, M. et al. Current-controlled magnetic domain-wall nanowire shift register. *Science* **320**, 209–211 (2008).
27. Jeon, J.-C. et al. Multicore memristor from electrically readable nanoscopic racetracks. *Science* **386**, 315–322 (2024).
28. Parkin, S. S. P. et al. Magnetic domain-wall racetrack memory. *Science* **320**, 190–194 (2008).
29. Zhao, Z. et al. Mott memristors for neuromorphics. *Adv. Phys. Res.* **4**, 2400129 (2025).
30. Imada, M., Fujimori, A. & Tokura, Y. Metal–insulator transitions. *Rev. Mod. Phys.* **70**, 1039 (1998).
31. Mott, N. F. The basis of the electron theory of metals, with special reference to the transition metals. *Proc. Phys. Soc. Lond. Sect. A* **62**, 416–422 (1949).
32. Li, G. et al. Computational elements based on coupled VO_2 oscillators via tunable thermal triggering. *Nat. Commun.* **15**, 5820 (2024).
33. Deng, L. The MNIST database of handwritten digit images for machine learning research. *IEEE Signal Process. Mag.* **29**, 141–142 (2012).
34. Brunel, N. & van Rossum, M. C. W. Quantitative investigations of electrical nerve excitation treated as polarization. *Biol. Cybern.* **97**, 341–349 (2007).
35. Bostner, Ž, Knoll, G. & Lindner, B. Information filtering by coincidence detection of synchronous population output: analytical approaches to the coherence function of a two-stage neural system. *Biol. Cybern.* **114**, 403–418 (2020).
36. Ashida, G., Kretzberg, J. & Tollin, D. J. Roles for coincidence detection in coding amplitude-modulated sounds. *PLoS Comput. Biol.* **12**, e1004997 (2016).
37. Hutcheon, B. & Yarom, Y. Resonance, oscillation and the intrinsic frequency preferences of neurons. *Trends Neurosci.* **23**, 216–222 (2000).
38. Izhikevich, E. M. Resonate-and-fire neurons. *Neural Netw.* **14**, 883–894 (2001).
39. Guth, T. A. et al. Theta-phase locking of single neurons during human spatial memory. *Nat. Commun.* **16**, 7402 (2025).

40. Dzyaloshinsky, I. A thermodynamic theory of 'weak' ferromagnetism of antiferromagnetics. *J. Phys. Chem. Solids* **4**, 241 (1958).
41. Moriya, T. Anisotropic superexchange interaction and weak ferromagnetism. *Phys. Rev.* **120**, 91 (1960).
42. Miron, I. M. et al. Fast current-induced domain-wall motion controlled by the Rashba effect. *Nat. Mater.* **10**, 419 (2011).
43. Ryu, K.-S., Thomas, L., Yang, S.-H. & Parkin, S. S. P. Chiral spin torque at magnetic domain walls. *Nat. Nanotechnol.* **8**, 527 (2013).
44. Liu, L., Lee, O. J., Gudmundsen, T. J., Ralph, D. C. & Buhrman, R. A. Current-induced switching of perpendicularly magnetized magnetic layers using spin torque from the spin Hall effect. *Phys. Rev. Lett.* **109**, 096602 (2012).
45. Durner, C. A., Migliorini, A., Jeon, J.-C. & Parkin, S. S. P. Reconfigurable magnetic inhibitor for domain-wall logic and neuronal devices. *ACS Nano* **19**, 5316–5325 (2025).
46. Leonard, T. et al. Shape anisotropy-dependent leaking in magnetic neurons for bio-mimetic neuromorphic computing. *ACS Nano* **19**, 5370–3477 (2025).
47. Zhang, R. Code of 16×16 VO2-MTJ crossbar array simulation. *Zenodo* https://doi.org/10.5281/zenodo.18800785 (2026).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

© The Author(s) 2026

Methods

VO_2 thin films and devices

VO_2 films with a thickness of about 23 nm (determined by X-ray reflectivity measurement) were deposited on a *c*-plane sapphire [0001] substrate by the pulsed laser deposition technique. A KrF excimer laser (Coherent LPX Pro) beam with a frequency of 3 Hz was focused onto the VO_2 target (99.9% purity, Plasmaterials) in a chamber under ambient O_2 of 0.020 mbar at 450 °C. Energy and fluence of the laser beam on the target surface were 50.3 mJ and 340 mJ cm^{-2}, respectively. After deposition, the sample was cooled down under ambient O_2 of 0.020 mbar. Then, the prepared thin films were patterned into devices by conventional optical lithography techniques—negative-tone photoresist (ARN-4340, Allresist), maskless aligner (Heidelberg MLA 150) and ion-beam etching (scia Coat 200) were used. For the electron-beam lithography, a JEOL EBL machine (JBX-8100FS, 100 kV) with an ARN-7520-18 resist was used. Then, electrodes were prepared by a lift-off method: ~77 nm of Ti/Au via ion-beam deposition (scia Coat 200).

Racetrack and MTJ films and devices

All spintronic films were deposited using a multisource, atomically engineered, next-generation alloys and compounds deposition system with a base pressure of $<10^{-9}$ torr. Films were grown at room temperature using magnetron sputtering in an Ar environment (3 mtorr). The film structures are as follows: (substrate) Si-SiO_2 or sapphire||100TaN|50Pt|0.6Ta|3Co|7Ni|3Co|30TaN for racetrack devices and (substrate) Si-SiO_2||40W|8CoFeB|20(10)MgO|12CoFeB|3Nb|3Co|12Pt|$[2Co|8Pt]_{\times 3}$|6Co|5Ir|6Co|$[8Pt|2Co]_{\times 6}$|30Ta for the SOT-MTJ device. Note that the MTJ stack was in situ annealed at 350 °C for 20 min after the Nb spacer layer. Magnetic characteristics of the samples are shown in Supplementary Figs. 17 (racetrack) and 20 (MTJ), respectively. Electrical characteristics of the samples are shown in Supplementary Fig. 18 (racetrack) and Supplementary Figs. 21 and 22 (MTJ). The prepared films were then patterned using conventional optical and electron-beam lithography techniques (as described above). The detailed device fabrication processes for racetracks, SOT-MTJs and directly integrated CAS cores are illustrated in Supplementary Figs. 16 and 19 and Extended Data Fig. 2, respectively.

Hybrid CAS core operation

For directly integrated devices, where VO_2 and racetrack devices are fabricated on the same substrate, a custom-built compact arbitrary waveform generator and circuit elements (bias tee and diode) with a printed connector are integrated. Device 1 was driven by a custom-built compact arbitrary waveform generator (see the next section) programmed to output current pulses converted from pixel values. Concurrently, device 2 was biased using a Keithley 6221 a.c./d.c. current source configured in a constant d.c. mode. The setup is shown in Extended Data Fig. 3.

For functionally integrated devices in which VO_2 and spintronic devices are prepared on different substrates, two independent Keithley 6221 a.c./d.c. current sources were used to apply current to VO_2 devices 1 and 2. Current source 1 to device 1 was programmed by Python script to send out current pulses converted from the pixel value (arbitrary waveforms). Current source 2 to device 2 was set in the constant-current mode.

Kerr microscopy measurements were conducted at room temperature using a modified system from Evico Magnetics. Differential Kerr imaging, which is sensitive to the out-of-plane component of magnetization, was used to track DW motion within the racetrack device. The entire experimental sequence, from driving VO_2 to DW read-out and reset, was automated using ScopeFoundry[48], an open-source Python library for controlling custom laboratory experiments. The DW position was extracted by averaging the Kerr contrast across the racetrack, followed by Gaussian filtering and first-derivative analysis to locate the edge. For repeated measurements, the DW position was reset by using nanosecond pulses generated by a PSPL 10300B pulse generator (20 ns, 2.25 V). On the basis of the displacement measured from the Kerr images after VO_2-induced motion, the automated routine estimates the number of reset pulses required to return the DW to its initial position and verifies successful DW repositioning. A radio-frequency switch (single-pole double-throw, d.c. to 40 GHz, RF Lambda), controlled by a National Instruments data acquisition USB-6363, toggles the racetrack input between the VO_2 spike output and the pulse generator, depending on the operation mode.

The minimum energy required for a single spike-to-memory conversion in the CAS core is calculated to be 51.9 nJ (two 5×5-μm^2 VO_2 memristors; Supplementary Fig. 26). The racetrack memory/SOT-MTJ is passively connected to the VO_2 LIF neuron, meaning that no additional power supply is required for the spintronic memories. As a result, the energy consumption of the integrated CAS core is entirely attributed to the VO_2 LIF neurons. Both racetrack memory and VO_2 device exhibit excellent scalability down to nanoscopic dimensions. For the VO_2 device below 200 nm, the energy required for generating a single spike operation can be greatly reduced to the ~200-pJ level (Supplementary Fig. 27). The energy consumption for single spike-to-memory conversion of the CAS core in this work is compared with other VO_2 as well as spintronic-memristor-based neural networks (Supplementary Table 1). A more detailed system-level energy consumption analysis (accounting for standby power, duty cycle and sparsity) is provided in Supplementary Section VIII and Supplementary Fig. 28.

Custom-built compact arbitrary waveform generator module

To demonstrate the feasibility of the hybrid CAS core operation, a compact arbitrary waveform generator was engineered using a high-precision 20-bit DAC (Analog Devices AD5791). The hardware comprises a tightly integrated three-tier stacked architecture: an AD5791 evaluation board (EVAL-AD5791ARDZ) conforming to the standard Arduino pinout, an Arduino Uno R4 microcontroller and an ultralow-noise precision voltage reference board (EV-ADR445-REFZ, equipped with an ADR445 reference). This compact configuration operates entirely via a single USB power supply and facilitates high-speed arbitrary waveform generation in the output range from −10 V to +10 V.

Custom C++ firmware was developed for the microcontroller to manage high-speed serial peripheral interface communication with the DAC at a 24-MHz clock frequency. To guarantee deterministic, nanosecond-level precision for the waveform output, the firmware utilizes register-level direct port manipulation for chip select and load DAC toggling, deliberately by-passing standard, high-latency software libraries. Furthermore, global hardware interrupts are temporarily disabled during the active firing sequence, and hardware-level no-operation assembly loops are used to strictly control transient pulse widths and interpulse gap timings. Waveform execution is dynamically orchestrated via a serial interface (115,200 baud) with a host PC, allowing custom Python scripts to seamlessly write sequenced voltage amplitudes and timing parameters into the microcontroller's memory buffer for immediate conversion.

Electrical transport measurements

For the electrical characterization of VO_2 memristors, the current–voltage (*I*–*V*) and current–frequency (*I*–*f*) measurements were carried out in a multiprobe room temperature probe station (Lakeshore) with 25-μm-diameter tungsten tips. Two current sources (Keithley 6221 a.c./d.c. current source) were used to drive the two VO_2 memristors. The output voltage was detected by an oscilloscope (DSO5052A, InfiniiVision) with custom-built LabView codes. The tests performed on VO_2 LIF neurons and arrays were carried out in a custom-built printed circuit board with multiple contact ports.

The VO_2-MTJ CAS core was measured using a custom-built automatic probe station. VO_2 spikes were injected into the channel of the

SOT-MTJ via a high-pass rectifier and a radio-frequency Be–Cu probe (40 GHz). Subsequently, the MTJ resistance was measured with a 100-μA d.c. current (Keithley 6221 a.c./d.c. current source and Keithley 2182a nanovoltmeter). A 25-μm-diameter tungsten tip (d.c.) and a Be–Cu high-frequency probe tip were used for contacts to the MTJ.

Data availability

The data that support the findings of this study are available from the corresponding authors upon request. Source data are provided with this paper.

Code availability

The source code for the 16 × 16 VO_2-MTJ crossbar array simulation is available via Zenodo at https://doi.org/10.5281/zenodo.18800785 (ref. 47). Code for training a software reconstruction of a SNN with 100 VO_2 LIF neurons in the MNIST handwritten dataset is provided in Supplementary Section IX.

References

48. Barnard, E. ScopeFoundry: a Python platform for controlling custom laboratory experiments and visualizing scientific data. https://scopefoundry.org/ (2017).

Author contributions

S.S.P.P. directed the project. G.L., J.-C.J. and S.S.P.P. conceived of the project and designed the experiments. G.L. and Y.C. prepared the VO_2 films. Z.Y. performed the transmission electron microscopy cross-section. J.-C.J., L.F. and A.M. prepared and characterized the magnetic thin films. G.L. characterized the prepared VO_2 films. G.L. fabricated the VO_2 devices and arrays. G.L. performed the large-scale network simulation for MNIST recognition. J.-C.J. and Y.-C.L. prepared the MTJ devices. J.-C.J. and G.L. prepared the racetrack devices. R.L. performed the Kerr microscopy measurements with the help of L.F. and A.M. G.L. and J.-C.J. conducted the electrical transport measurements. G.L., R.L. and S.A. prepared the automated Python codes for electrical transport and Kerr microscopy measurements. J.-C.J. and Y.-C.L. prepared the electrical transport measurement setups and LabVIEW programs. G.L., R.L., R.Z. and J.-C.J. prepared directly integrated device with circuit integration. G.L., R.L., J.-C.J. and S.S.P.P. prepared the paper. All authors participated in the discussion of the results.

Funding

S.S.P.P. acknowledges the International Max Planck Research School for Science and Technology of Nano-Systems. This research was also supported by the Samsung Electronics R&D program 'Material and Device research on Racetrack Memory'. Funding by the Max Planck Graduate Center for Quantum Materials and by the Fraunhofer-Max Planck cooperation programme in the 'Joint Initiative for Research and Innovation' ('Pakt für Forschung und Innovation'/project CONDOR) is gratefully acknowledged. We also acknowledge funding from the Deutsche Forschungsgemeinschaft within TRR404 Active-3D (project number 528378584), project A04. Open access funding provided by Max Planck Society.

Competing interests

The authors declare no competing interests.

Additional information

Extended data is available for this paper at https://doi.org/10.1038/s41565-026-02311-0.

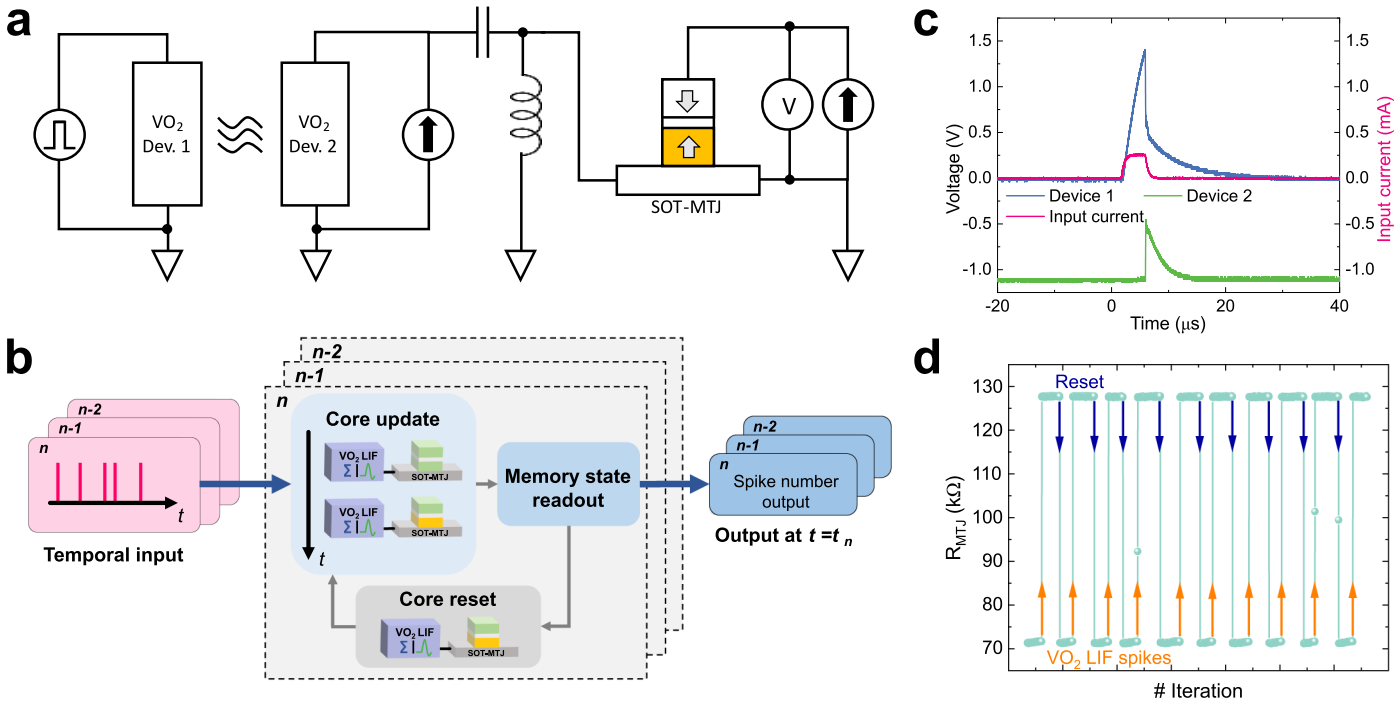
Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41565-026-02311-0.

Correspondence and requests for materials should be addressed to Jae-Chun Jeon or Stuart S. P. Parkin.

Peer review information *Nature Nanotechnology* thanks Zhongrui Wang and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

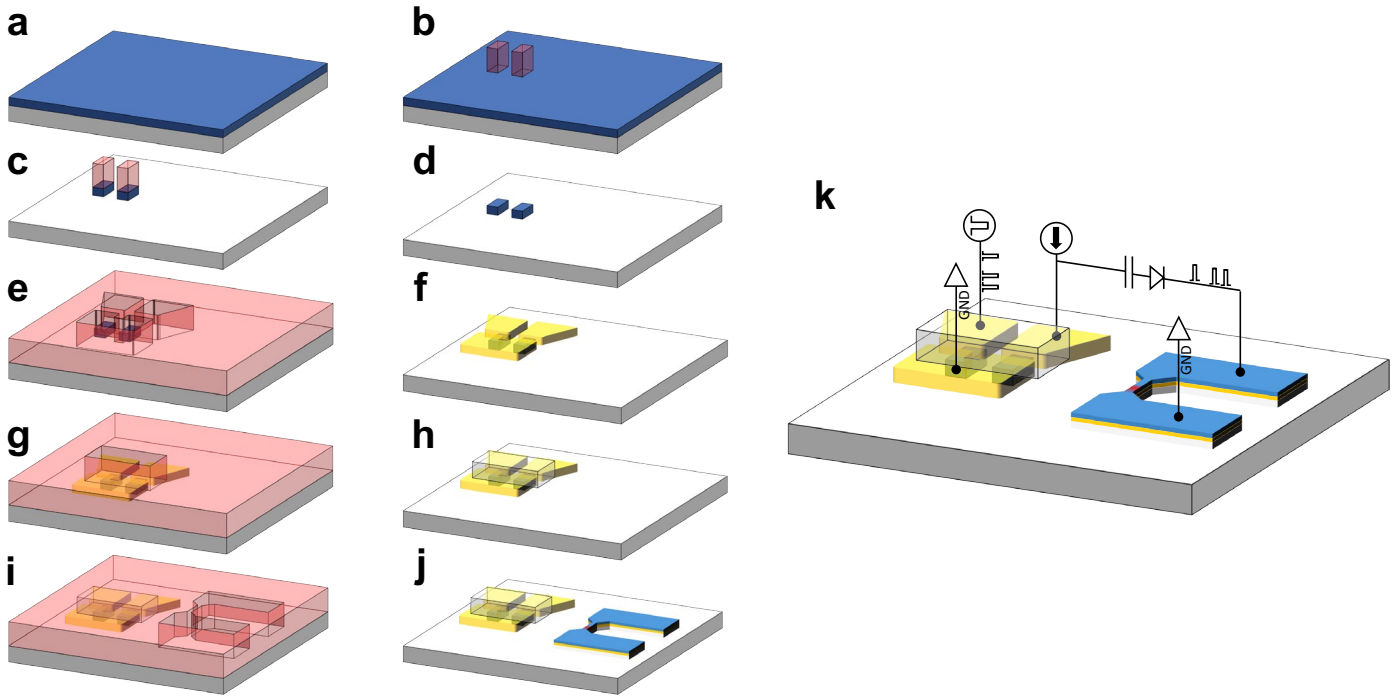
Reprints and permissions information is available at www.nature.com/reprints.

https://doi.org/10.1038/s41565-026-02311-0



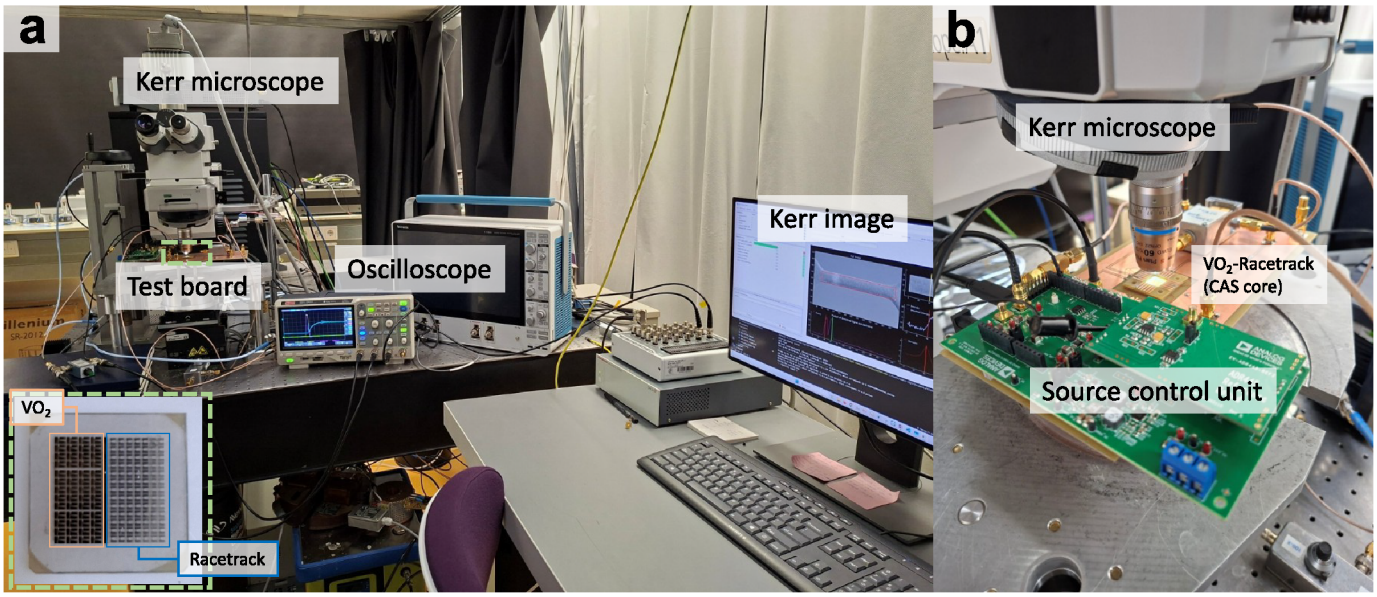
Extended Data Fig. 1 | Nanoscopic VO_2 LIF neuron spike to MTJ. a, Schematic illustration of the circuit diagram for VO_2 LIF neuron (500×500 nm^2) and SOT-MTJ (100×400 nm^2). **b**, Architecture for a hybrid Mott-spintronic CAS unit and its operation procedure. **c**, A single input current pulse triggers VO_2 LIF to fire a spike. **d**, Preserving spike history and electrical readout by SOT-MTJ. Blue curve represents the switching of a SOT-MTJ by the VO_2 LIF neuron spikes (orange arrows), and reset by external field (purple arrows).

https://doi.org/10.1038/s41565-026-02311-0



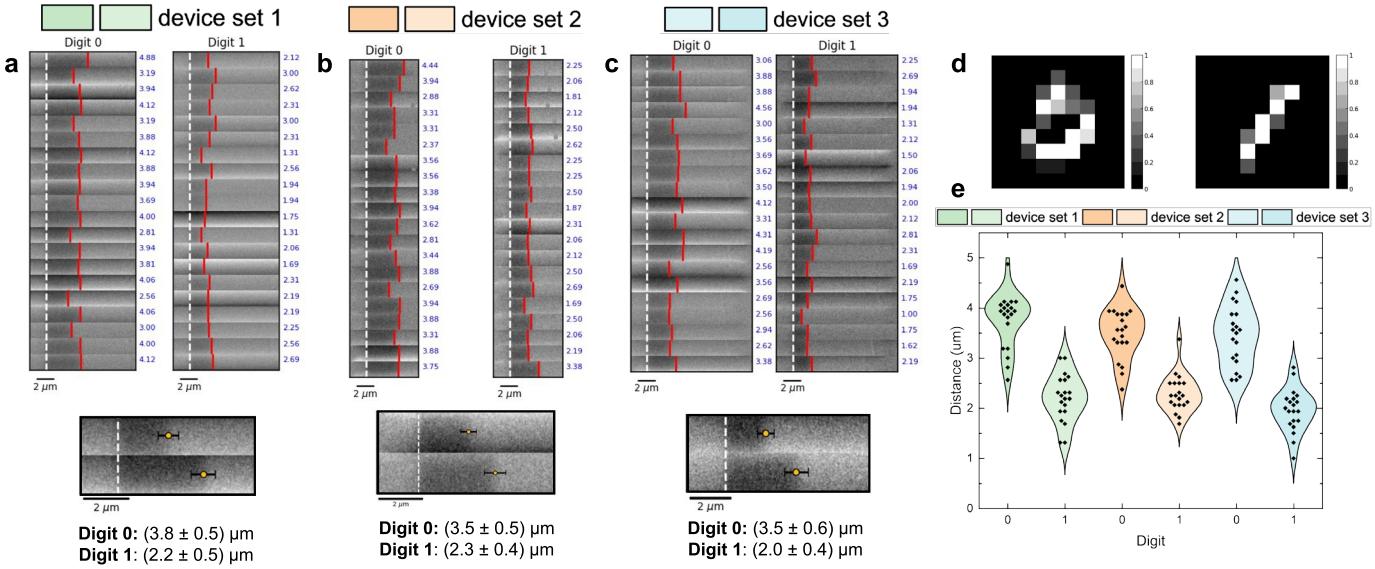
Extended Data Fig. 2 | Device fabrication process for directly integrated CAS element on the same substrate. a, Preparation VO_2 film on a sapphire substrate. **b**, Patterning of thermally coupled VO_2 device, that is, VO_2 neuron using a negative tone resist (ARN 4340). **c**, Ion beam etching of the pattern. At this point, we also completely etch half of the substrate for racetrack integration in later process. **d**, Cleaning of resist. **e**, Patterning of contact for VO_2 neuron using a positive tone resist (ARP 3540 T). **f**, Cleaning of the resist after the lift-off process. **g**, Patterning of VO_2 neuron protection layer. **h**, Deposition of Al_2O_3 (150 nm using e-beam evaporator) and cleaning of the sample. **i**, Patterning of racetrack devices using a positive tone resist (ARP 3540 T) with undercut (ARBR 5480). **j**, Deposition of racetrack stack (100TaN – 50Pt – 0.6Ta – 3Co – 7Ni – 3Co – 30TaN; units in Å – nominal thickness). **k**, Circuit diagram for the measurement.

https://doi.org/10.1038/s41565-026-02311-0



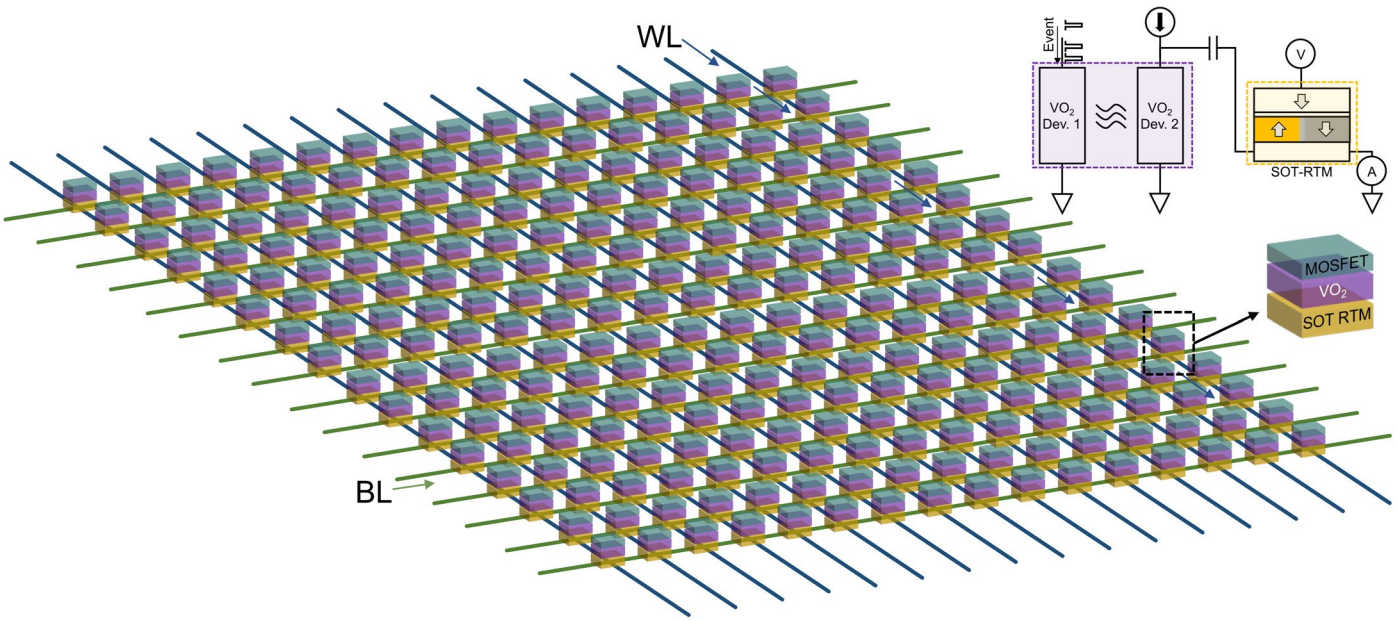
Extended Data Fig. 3 | Photo of test setup. a, Test setup with Kerr microscope. **b**, Photo of the test board, taken during the measurement.

https://doi.org/10.1038/s41565-026-02311-0



Extended Data Fig. 4 | Pattern recognition tasks using different sets of CAS core. a, Task performance from device-set 1; **b**, device-set 2; **c**, device-set 3. Note that device-sets 1 and 2 were tested with inductor connected – while device-set 3 was tested without inductor. The results are qualitatively the same among different measurements. **d**, Grey scale pattern '0' and '1'. **e**, summary of results.

https://doi.org/10.1038/s41565-026-02311-0



Extended Data Fig. 5 | Proposed 16 × 16 array device for 1T-CAS simulation. Each 1T-CAS unit cell is composed of a MOSFET, VO_2 neuron, and racetrack memory – where WL (word-line) controls MOSFET while BL (bit-line) reads MTJ signal out. Top right inset shows simplified circuit diagram for CAS. See supplementary video for real-time simulation results to identify 'M', 'P', and 'G' using the proposed 16 × 16 array device.