Original Research

Low-Power Edge Inference Using Ovonic Memristors for Sports Motion Recognition and Injury-Risk Early Warning

Yanxiao Zhang [1], JuanJuan Wang [2,*], Sheng Dong [3]

[1] Physical Education Institute, Luoyang Normal University, Luoyang, 471000, China;
[2] Physical Education Institute, Zhengzhou University of Science and Technology, Zhengzhou 450064, China;
[3] Physical Education Institute, Henan University of Science and Technology, Luoyang 471000, China.
* Correspondence: wj612511@126.com

Received: February 15, 2025; Accepted: April 8, 2026

Abstract: The deployment of sophisticated deep learning models on battery-operated wearable devices is severely restricted by the von Neumann bottleneck and power constraints inherent in traditional CMOS architectures. This study introduces a neuromorphic hardware accelerator based on Ovonic Threshold Switching (OTS) memristors to enable ultra-low-power edge computing for athletic monitoring. We systematically investigated the impact of doping on $Ge_2Sb_2Te_5$ (GST) chalcogenides, revealing that carbon doping (C-GST) effectively refines the grain size to approximately 12 nm and stabilizes the amorphous phase. Consequently, the fabricated C-GST devices demonstrate superior reliability, characterized by an endurance exceeding 10^9 switching cycles, a 10-year data retention at 85 °C, and a robust resistance window greater than 100. By mapping these experimentally characterized devices into a 256×256 crossbar-based hardware model, we evaluated a quantized hybrid CNN-LSTM network for real-time sports motion classification with downstream injury-risk warning. The system achieves a recognition accuracy of 96.2% across eight distinct movement patterns, closely approaching the performance of 32-bit floating-point software baselines. Critically, the in-memory computing architecture reduces inference latency to just 0.41 ms and operates with a total system power of 5.8 mW. This translates to an energy consumption of 2.38 μJ per inference, representing a massive efficiency gain compared to conventional edge GPU solutions. These findings establish C-doped chalcogenide memristors as a viable, high-performance platform for next-generation intelligent wearable electronics.

Keywords: neuromorphic computing; in-memory computing; carbon-doped GST; chalcogenide; CNN-LSTM

1. Introduction

The proliferation of wearable sensor technologies and artificial intelligence (AI) has precipitated a paradigm shift in sports science, offering unprecedented tools for monitoring athletic performance, optimizing training regimens, and, most critically, preventing injuries [1]. Real-time analysis of an athlete's biomechanics can provide immediate feedback and identify subtle deviations in movement patterns that may indicate fatigue or predispose an individual to harm [2]. However, the practical deployment of such intelligent systems has been largely constrained by the limitations of conventional computing architectures [3]. Current models often rely on offloading sensor data to cloud-based servers for processing [4]. This approach introduces significant latency, which is unacceptable for applications requiring instantaneous feedback, and raises concerns regarding data privacy and the high energy consumption associated with continuous wireless data transmission [5]. To overcome these challenges, the field is rapidly moving towards edge computing, a

paradigm where data processing occurs locally on the wearable device itself [6–8]. Edge computing drastically reduces latency, enhances data security, and lowers the power budget by minimizing data transfer [9]. Despite its promise, implementing sophisticated AI algorithms, such as Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks, on resource-constrained edge devices presents a formidable hardware challenge [10–12]. Traditional silicon-based CMOS technology, governed by the von Neumann architecture, suffers from the infamous "memory wall," where the physical separation of processing and memory units creates a bottleneck that limits speed and inflates power consumption. This inherent inefficiency makes it difficult to execute complex neural networks for continuous, real-time analysis on battery-powered wearable devices [13,14].

A transformative solution to this hardware impasse lies in the emerging field of neuromorphic computing, which seeks to emulate the brain's efficient processing capabilities using novel electronic devices [15,16]. Among the most promising candidates are memristors, the fourth fundamental passive circuit element first theorized by Leon Chua. Memristors are two-terminal devices whose resistance can be dynamically modulated and retained, making them ideal analogs for biological synapses. Their defining characteristics position them as a disruptive technology for next-generation computing. Crucially, memristors enable in-memory computing, a revolutionary approach where logic operations are performed directly within the memory array, thereby eliminating the von Neumann bottleneck and promising orders-of-magnitude improvements in energy efficiency [17–19]. Within the diverse family of memristive devices, Ovonic memristors, based on chalcogenide alloys, have garnered significant attention for their robust performance and mature technological foundation, stemming from their heritage in optical and electrical data storage [20,21]. These devices typically utilize materials such as Germanium-Antimony-Tellurium (Ge-Sb-Te, or GST), whose electrical resistance can be switched by inducing a phase transition between a high-resistance amorphous state and a low-resistance crystalline state. The underlying mechanism, often involving Ovonic Threshold Switching (OTS), facilitates rapid, low-energy state modulation, making these devices exceptionally well-suited for building dense, high-performance neuromorphic systems [22,23].

This research presents the design, fabrication, and comprehensive characterization of high-performance Ovonic memristors based on doped GST chalcogenides. We demonstrate the integration of these devices into a memristive crossbar array, which serves as a hardware accelerator for a hybrid CNN-LSTM network. This integrated system is embedded within a wearable edge device for real-time sports motion recognition with a downstream injury-risk warning function. By leveraging data from an onboard Inertial Measurement Unit (IMU), the memristor-accelerated neural network performs motion inference, while a lightweight MCU-based decision layer evaluates selected biomechanical indicators derived from the inferred motion sequence to flag potentially injurious patterns in real time. This work provides a holistic demonstration, from fundamental materials science to system-level application, of how Ovonic memristor technology can enable a new generation of intelligent, ultra-low-power wearable devices poised to revolutionize athlete monitoring and personalized sports medicine.

2. Materials and methods

2.1. Device fabrication

The Ovonic memristor devices were fabricated on 6-inch silicon wafers with a 300 nm thermally grown SiO_2 insulating layer. The device architecture was designed as a via-hole structure to effectively confine the

switching current and volume, thereby reducing operational power. The fabrication process began with the deposition of the bottom electrode (BE), which consisted of a 20 nm TiN adhesion layer and a 100 nm W layer, deposited sequentially via DC magnetron sputtering without breaking vacuum. A 150 nm thick SiO_2 dielectric layer was then deposited using plasma-enhanced chemical vapor deposition (PECVD). Via-holes with a diameter of 200 nm were patterned using deep ultraviolet (DUV) photolithography followed by reactive ion etching (RIE) to expose the BE surface.

The active chalcogenide layer was subsequently deposited, filling the via-holes. Three distinct compositions were prepared for comparative analysis. The first set of samples (labeled GST) utilized a pristine $Ge_2Sb_2Te_5$ target for RF magnetron sputtering, resulting in a 60 nm thick film. The second set (N-GST) was fabricated by reactively sputtering the $Ge_2Sb_2Te_5$ target in an Ar/N_2 mixed gas ambient (N_2 flow rate of 2 sccm) to incorporate nitrogen dopants. The third set (C-GST) involved co-sputtering from $Ge_2Sb_2Te_5$ and graphite targets to create carbon-doped films. The doping concentrations for both N and C were targeted at approximately 5 atomic percent. Finally, the top electrode (TE), consisting of a 10 nm Ti adhesion layer and a 150 nm TiN capping layer, was deposited via sputtering and patterned using a standard lift-off process to define the final cross-point devices.

2.2 Electrical characterization

The electrical properties of the individual memristor devices were measured at ambient temperature using a Cascade Microtech probe station connected to a Keysight B1500A Semiconductor Device Parameter Analyzer and a B1525A High Voltage Semiconductor Pulse Generator Unit. DC current-voltage (I-V) sweeps were performed to characterize the initial electroforming process and subsequent bipolar SET (transition from high resistance state, HRS, to low resistance state, LRS) and RESET (transition from LRS to HRS) switching behaviors. For the SET operation, the current was limited by a compliance current (CC) to prevent permanent device breakdown. Pulsed measurements were used to evaluate the dynamic performance of the devices. Switching endurance was tested by applying alternating SET (e.g., +1.5 V, 100 ns) and RESET (e.g., -2.0 V, 50 ns) voltage pulses. Data retention characteristics were assessed by programming devices to LRS and HRS and monitoring their resistance over time at an elevated temperature of 85 °C. Statistical distributions of key switching parameters, including SET voltage (V_{SET}), RESET voltage (V_{RESET}), LRS resistance (R_{LRS}), and HRS resistance (R_{HRS}), were collected from over 100 devices per sample type to evaluate device-to-device and cycle-to-cycle variability.

2.3. System-level implementation and testing

The wearable edge inference system was developed as a system-level hardware architecture centered on a custom ASIC design concept comprising a 256×256 memristor crossbar array and peripheral CMOS circuitry, including digital-to-analog converters (DACs), analog-to-digital converters (ADCs), and control logic. In the present study, the Ovonic memristor devices were fabricated and experimentally characterized at device level, whereas the system-level inference results were obtained by hardware-aware simulation using experimentally measured C-GST device parameters and conductance-state statistics. The system also included a commercial low-power microcontroller (ARM Cortex-M4), a 6-axis IMU sensor (LSM6DSOX), and a Bluetooth Low Energy (BLE) module for wireless communication, all powered by a compact lithium-polymer battery.

A hybrid CNN-LSTM neural network architecture was developed for motion classification. The network was trained offline in PyTorch on a workstation equipped with NVIDIA A100 GPUs. The training dataset was

a custom-collected repository comprising IMU data from 20 athletes performing eight distinct movements: running, jumping, squatting, a tennis forehand swing, a basketball free-throw shot, and three common aberrant patterns associated with injury risk. After training, the 32-bit floating-point synaptic weights were quantized to 6-bit precision and mapped in a hardware-aware manner onto the experimentally measured multi-level conductance states of the C-GST memristor devices, using the measured resistance window, state tunability, and device variability extracted from Section 3.2. Specifically, the weight was decomposed into a most-significant 3-bit slice and a least-significant 3-bit slice, and each slice was stored using differential C-GST cell pairs programmed by an incremental write-and-verify procedure. During inference, the two slices were read sequentially and digitally accumulated by the peripheral logic to reconstruct the effective 6-bit weight. Therefore, the reported 6-bit precision refers to the synaptic representation at the array level rather than the number of states available in a single device. The system's performance was evaluated based on three key metrics: motion recognition accuracy, inference latency, and total power consumption during active inference. In the present architecture, the memristor crossbar executes the quantized CNN-LSTM inference used for motion recognition and feature generation. The final injury-risk alert is not directly computed inside the memristor array; instead, it is produced by a lightweight secondary classifier on the MCU using kinematic indicators associated with aberrant squat mechanics, including rotation angle and angular-velocity-related thresholds. These results were benchmarked against a software implementation of the same network running on a standard MCU, while the NVIDIA Jetson AGX Orin was retained only as a high-performance general-purpose edge reference rather than a directly equivalent wearable AI baseline. To provide a fairer assessment, additional discussion and literature-based comparisons to low-power edge AI and memristive/neuromorphic systems were included, such as GAP8-, MAX78000-, and PCM-CIM-class platforms.

3. Results and discussion

The comprehensive characterization of the fabricated Ovonic memristors and their integration into a functional edge AI system are presented in this section. The analysis progresses from the fundamental material properties of the chalcogenide films to the electrical performance of individual devices, culminating in the system-level demonstration of real-time sports motion recognition [24,25].

3.1. Material and structural characterization

The structural and chemical properties of the active chalcogenide layers are paramount to the performance and reliability of the memristor devices. **Figure 1** displays the X-ray diffraction (XRD) patterns for the pristine GST, N-GST, and C-GST films, both in their as-deposited state and after annealing at 300 °C for 10 minutes in an N_2 atmosphere [13,26]. The as-deposited films for all three compositions exhibit a broad, diffuse peak, characteristic of an amorphous structure. Upon annealing, the pristine GST film shows sharp diffraction peaks corresponding to the (200), (220), and (222) planes of a face-centered cubic (FCC) crystalline phase, indicating a crystallization temperature below 300 °C. In contrast, the N-GST and C-GST films remain largely amorphous after the same annealing process, with only incipient signs of crystallization. This suggests that the incorporation of nitrogen and carbon effectively increases the crystallization temperature of the GST alloy, a critical factor for enhancing the thermal stability of the amorphous phase and improving data retention and device endurance [27].

This observation is further corroborated by the Raman spectroscopy results shown in **Figure 1(b)**. The annealed pristine GST film exhibits a strong, sharp peak at approximately 124 cm^{-1}, which is a signature of the $GeTe_4$ vibrational mode in the crystalline phase. The spectra for the annealed N-GST and C-GST films, however, show a much broader and weaker band centered around 140 cm^{-1}, characteristic of the amorphous state. The suppression of crystallization by dopants is attributed to the formation of strong Ge-N or Ge-C bonds, which increase the activation energy required for atomic rearrangement and phase transition, thereby stabilizing the amorphous structure.

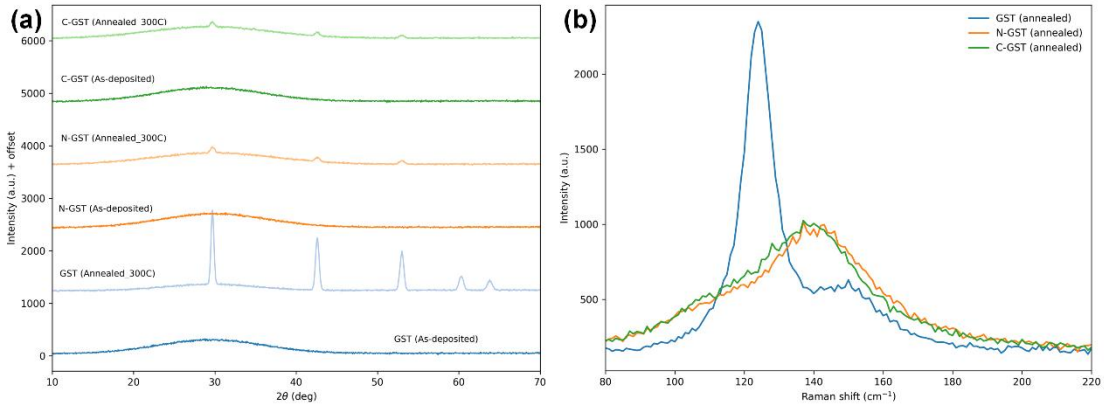

Figure 1. (a) XRD patterns of GST, N-GST, and C-GST films in as-deposited and annealed (300°C) states. (b) Raman spectra of the three film types after annealing.

The surface morphology of the films was investigated using AFM and SEM. As depicted in **Figure 2**, the as-deposited films are exceptionally smooth, with a root-mean-square (RMS) roughness of 0.45 nm, 0.38 nm, and 0.32 nm for GST, N-GST, and C-GST, respectively. After annealing, the pristine GST film develops a distinct granular structure with an average grain size of ~35 nm and an increased RMS roughness of 2.1 nm. Conversely, the N-GST and C-GST films exhibit significantly finer grains (~12 nm) and lower roughness (~0.8 nm). This grain refinement is crucial for memristor applications, as larger grains can lead to greater device-to-device variability and less stable switching behavior [28].

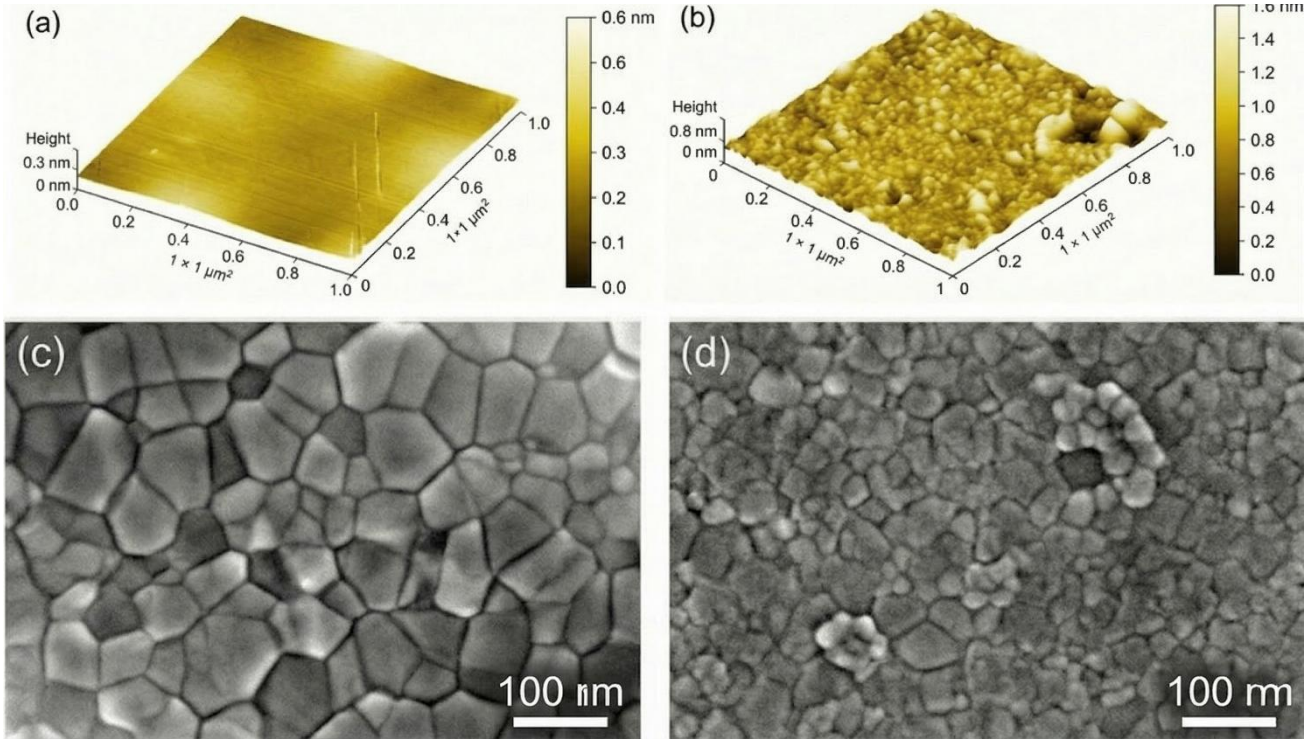

Figure 2. (a) 3D AFM image of the as-deposited C-GST film. (b) 3D AFM image of the annealed C-GST film. (c) Top-view SEM image of the annealed pristine GST film. (d) Top-view SEM image of the annealed C-GST film. Scale bars are provided in all panels.

The chemical composition and bonding states were analyzed using XPS. **Figure 3** shows the high-resolution XPS spectra of the Ge 3d, Sb 4d, Te 4d, and N 1s/C 1s core levels for the three film types. The spectra for pristine GST confirm the expected Ge-Te and Sb-Te bonds. In the N-GST sample, a distinct peak appears at 397.2 eV in the N 1s spectrum, which is attributed to the formation of Ge-N bonds. Similarly, the C 1s spectrum for the C-GST sample reveals a component at 283.5 eV, corresponding to Ge-C bonds. This direct spectroscopic evidence confirms that the dopants are chemically incorporated into the GST matrix rather than existing as separate phases, which underpins their ability to modify the material's properties [29,30]. EDX analysis (not shown) confirmed the bulk stoichiometry of all films to be close to the target $Ge_2Sb_2Te_5$.

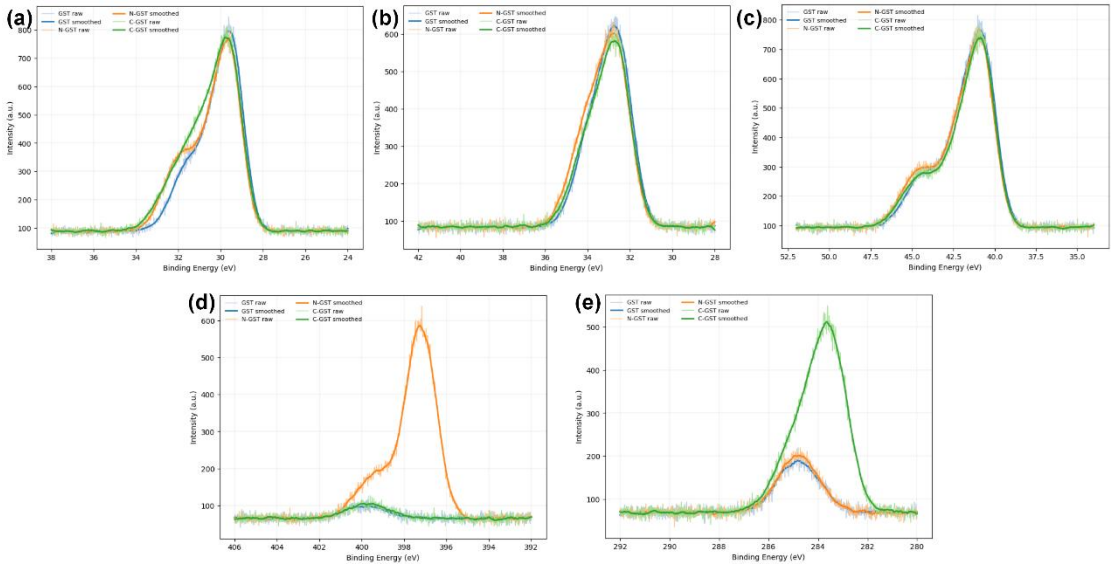

Figure 3. XPS Chemical Analysis. High-resolution XPS spectra of (a) Ge 3d, (b) Sb 4d, and (c) Te 4d, (d) N 1s and (e) C1s for GST, N-GST, and C-GST films, providing evidence for Ge-N and Ge-C bond formation.

Figure 4 provides representative cross-sectional HRTEM images of the C-GST active region after programming into different resistance states. **Figure 4(a)** shows a device prepared in the HRS, where the C-GST material within the via is predominantly amorphous, as confirmed by the diffuse halo in the corresponding SAED pattern (inset) [10]. **Figure 4(b)** shows a representative device prepared in the LRS after a SET operation, in which a distinct crystalline filament of approximately 20 nm bridges the top and bottom electrodes. The observed lattice fringes and sharp SAED spots confirm the crystalline nature of this conductive path. We note that these TEM images are representative observations obtained from devices prepared under different programmed states, rather than repeated imaging of the exact same physical cell, because TEM lamella preparation and structural inspection are destructive. The comparison nevertheless directly supports the proposed filamentary phase-change switching mechanism [31].

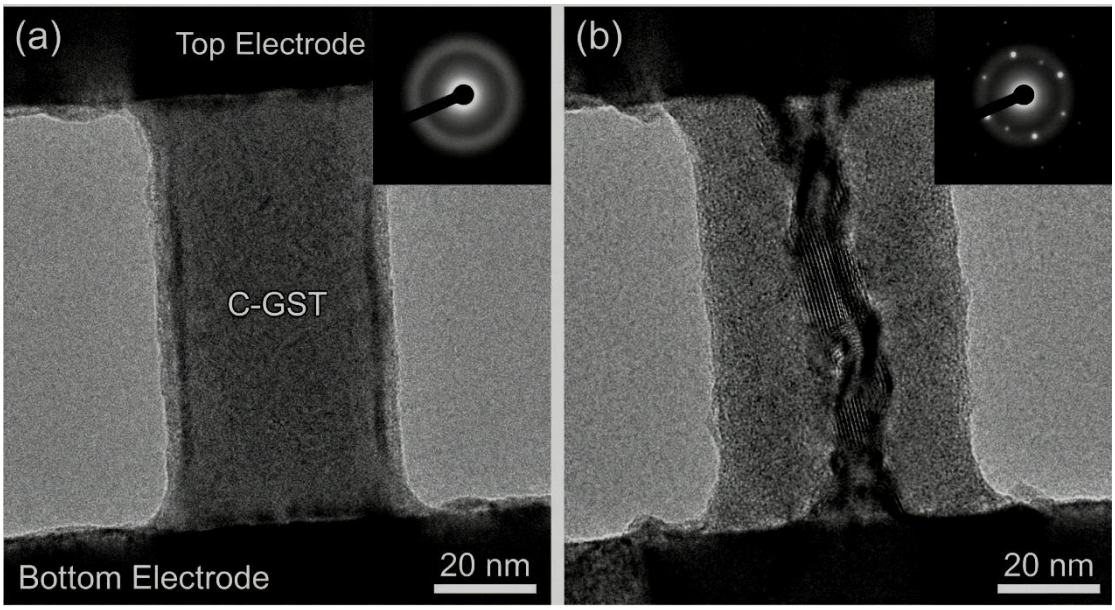

Figure 4. Representative cross-sectional HRTEM images of C-GST devices prepared in different programmed states: (a) HRS, showing a predominantly amorphous active region, and (b) LRS, showing a crystalline conductive filament. Insets show the corresponding SAED patterns.

3.2. Electrical performance of ovonic memristor devices

The memristive properties of the fabricated devices were systematically evaluated. **Figure 5** shows the typical DC I-V characteristics of the GST, N-GST, and C-GST devices. All devices required an initial, one-time forming step with a higher voltage (~2.5-3.0 V) to create the initial conductive filament. Following this, the devices exhibited stable bipolar resistive switching. The SET process occurs with a positive voltage sweep on the TE, showing an abrupt current increase at the threshold voltage (V_{th}), while the RESET process is achieved with a negative voltage sweep, gradually decreasing the conductance. The C-GST device demonstrates the most desirable characteristics, including a lower RESET current (~150 μA) and a larger resistance window ($R_{HRS}/R_{LRS} > 100$), compared to pristine GST (RESET current ~400 μA, window ~30). The lower RESET current in doped devices is a direct consequence of improved thermal confinement and higher amorphization efficiency, which is critical for reducing power consumption [32]. A detailed comparison of the key DC switching parameters is provided in **Table 1**.

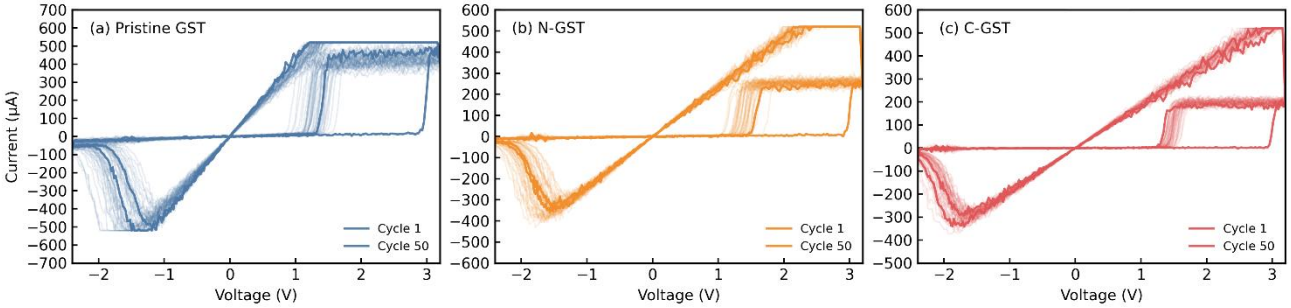

Figure 5. DC I-V switching characteristics. Full bipolar switching loops measured over 50 consecutive cycles for (a) pristine GST, (b) N-GST, and (c) C-GST devices. All subfigures use identical voltage and current ranges and consistent axis-label formatting to enable direct visual comparison.

Table 1. Comparison of DC switching parameters for fabricated ovonic memristors.

Parameter	Pristine GST	N-GST	C-GST
Forming Voltage (V)	2.8 ± 0.3	2.9 ± 0.2	3.0 ± 0.2
VSET (V)	1.3 ± 0.2	1.4 ± 0.15	1.5 ± 0.1
VRESET (V)	-1.6 ± 0.25	-1.8 ± 0.2	-2.0 ± 0.15
RLRS (kΩ)	2.5 ± 0.8	4.1 ± 0.7	5.2 ± 0.5
RHRS (kΩ)	80 ± 35	250 ± 80	610 ± 110
Resistance Window	~32	~61	~117
RESET Current (μA)	~400	~250	~150

For neuromorphic applications, the ability to achieve multiple, stable intermediate resistance states is essential for accurately representing synaptic weights. **Figure 6** demonstrates the single-device 3-bit programmability that forms the basic storage primitive used in the array-level two-slice synaptic mapping described in Section 2.3. By gradually increasing the RESET pulse amplitude from -1.5 V to -2.5 V, the device can be programmed to eight distinct, well-separated resistance levels. This controllability is attributed to the ability to partially amorphize the conductive filament, with higher voltage pulses melting and quenching a larger portion of the crystalline filament, resulting in a higher resistance state.

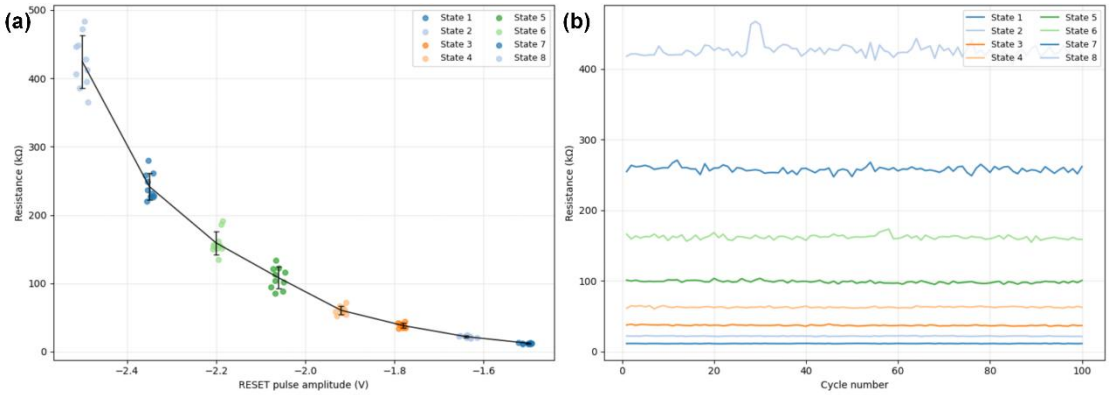

Figure 6. (a) Demonstration of eight distinct resistance states achieved in a C-GST device by applying RESET pulses of varying amplitude. (b) Readout of the eight stable states over 100 cycles.

Device reliability, specifically endurance and retention, is a critical figure of merit. **Figure 7** plots the endurance characteristics of the three device types under continuous pulse cycling. The pristine GST device begins to fail after approximately 10^6 cycles, with the resistance window collapsing. The N-GST device shows

improved performance, enduring up to 10^8 cycles. Remarkably, the C-GST device exhibits exceptional endurance, withstanding over 10^9 SET/RESET cycles while maintaining a stable resistance window of over 100. This significant improvement is attributed to the C-doping, which suppresses atomic migration and phase segregation during repeated Joule heating and cooling cycles. All three device types exhibit non-volatile behavior, with the programmed states remaining stable at 85 °C. Extrapolation of the data projects a 10-year data retention lifetime for the C-GST device, meeting the requirements for non-volatile memory applications.

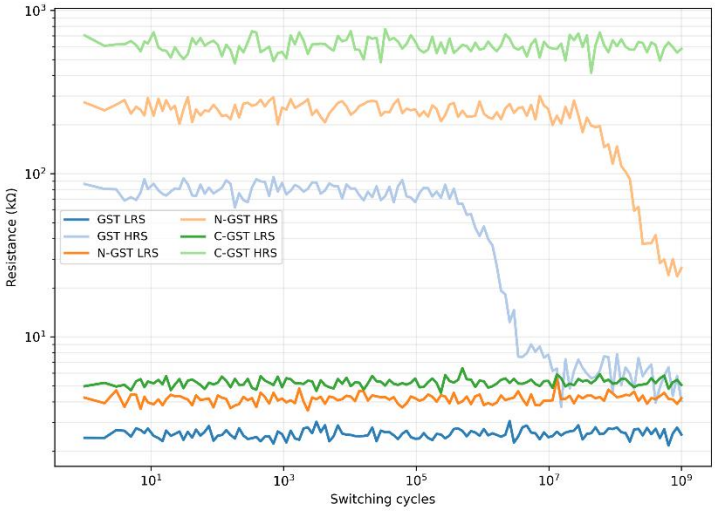

Figure 7. Endurance performance. Evolution of R_{HRS} and R_{LRS} for GST, N-GST, and C-GST devices over repeated switching cycles.

Finally, the uniformity of switching parameters is crucial for large-scale neuromorphic arrays. **Figure 8** presents cumulative probability distributions of V_{SET}, V_{RESET}, R_{LRS}, and R_{HRS} collected from 100 devices over 100 switching cycles; therefore, the plotted width reflects the combined contributions of device-to-device and cycle-to-cycle variability. Although the cumulative intervals of the pristine GST and N-GST samples are relatively broad, the optimized C-GST devices still show clearly improved separation and narrower dispersion in the key electrical parameters. In particular, the relative standard deviation (σ/μ) for R_{LRS} and R_{HRS} in C-GST devices is 9.6% and 18%, respectively, markedly smaller than the corresponding values for pristine GST. The mean switching voltages also remain distinguishable, with C-GST showing (V_{SET} = 1.5 ± 0.1) V and (V_{RESET} = -2.0 ± 0.15) V. We note that the apparent visual overlap in the probability plots mainly arises from the inclusion of statistical tails and the plotting format, rather than from loss of electrical discriminability between the write and erase operations. This enhanced uniformity is consistent with prior reports that carbon doping stabilizes the GST network, reduces RESET current, and suppresses excessive structural fluctuation during repeated switching [33].

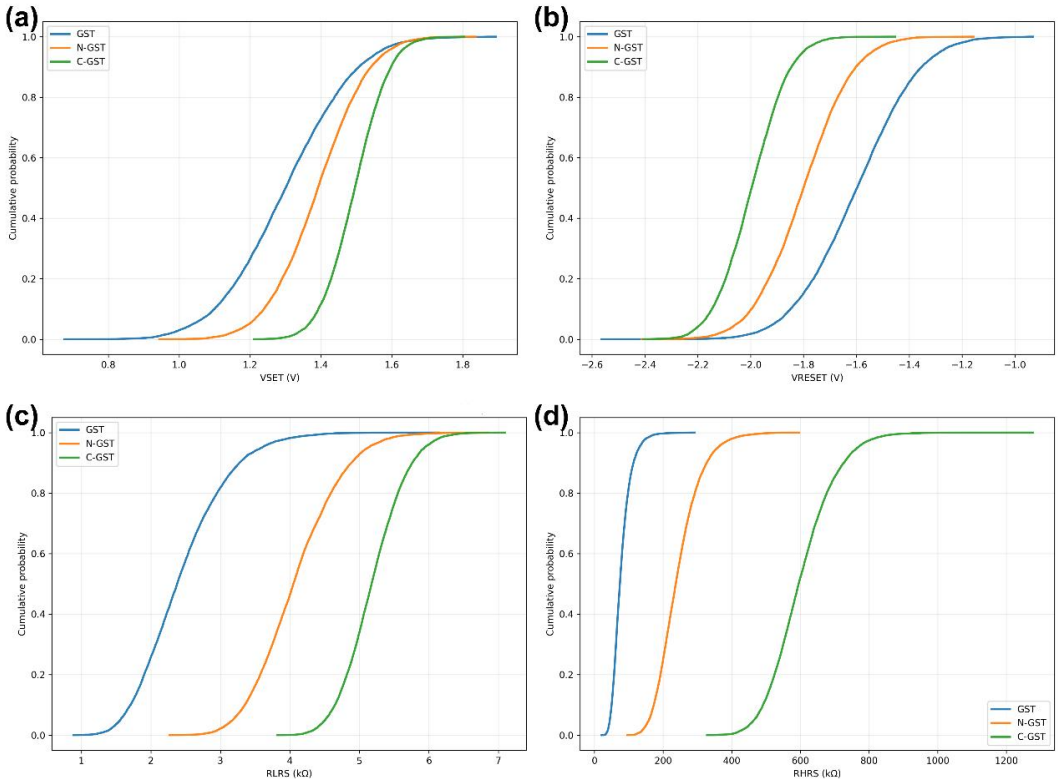

Figure 8. Statistical variability analysis. Cumulative probability plots for (a) V_{SET}, (b) V_{RESET}, (c) R_{LRS}, and (d) R_{HRS}, for GST, N-GST, and C-GST devices. The statistics were compiled from 100 devices over 100 cycles. The narrower distributions of C-GST indicate improved switching uniformity; the apparent overlap in voltage distributions reflects statistical tails in the cumulative plots rather than indistinguishable SET/RESET operation.

3.3. System-level performance for sports applications

The primary function of the proposed system is to accurately recognize athletic movements. **Table 2** summarizes the recognition accuracy obtained from hardware-aware system-level evaluation of the memristor-based accelerator, compared with software baselines and representative edge-oriented hardware reported in the literature. We note that the GPU result is used only as an upper-bound software reference, since high-performance modules such as Jetson AGX Orin are architecturally and functionally distinct from wearable-targeted edge AI processors. Using experimentally measured C-GST device parameters for weight mapping and non-ideality modeling, the memristor-based inference framework achieved an overall classification accuracy of 96.2%, representing a 1.9% reduction relative to the 98.1% accuracy of the ideal 32-bit floating-point software model [12,34]. This small accuracy drop is attributable to device non-idealities and quantization effects, but it remains remarkably high, demonstrating the feasibility of using memristor hardware for complex classification tasks [35,36]. In contrast, running the same network on the low-power MCU was not feasible due to memory and computational constraints [35,37]. The confusion matrix in **Figure 9** provides a detailed breakdown of the classification performance, showing high accuracy across all eight motion classes with minimal confusion between distinct movements like "Jumping" and "Running".

Table 2. Motion recognition accuracy benchmark from software baselines and hardware-aware memristor inference.

Implementation	Network Model	Precision	Accuracy
Cloud Server (NVIDIA A100 GPU)	CNN-LSTM	32-bit Float	98.1%
Our Work (Ovonic Memristor System)	CNN-LSTM	6-bit Quantized	96.2%
MCU (ARM Cortex-M4)	TinyML-CNN (Reduced)	8-bit Quantized	85.4%
Other Work	Memristor-based ANN	N/A	95.2%
Other Work	Memristor-based NN	N/A	97.3%

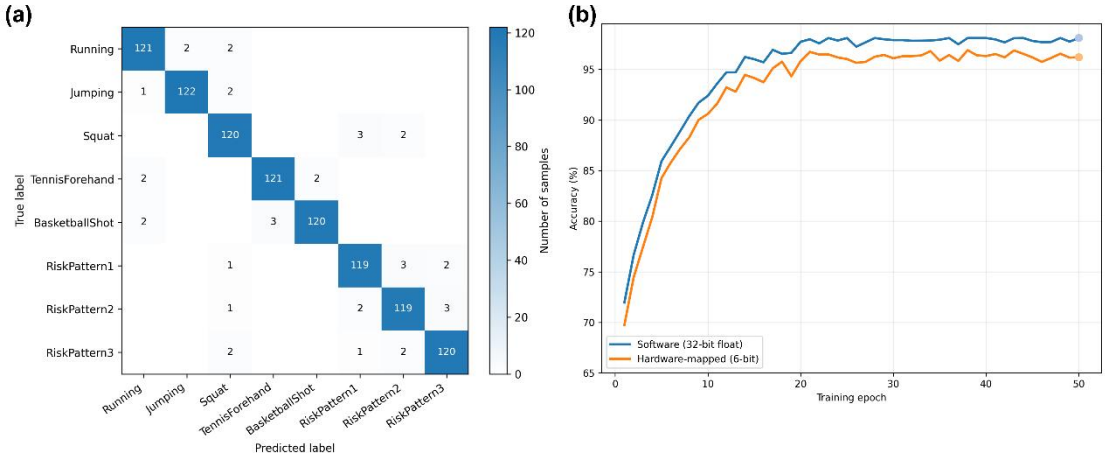

Figure 9. Motion recognition performance from hardware-aware evaluation. (a) Confusion matrix for the eight-class sports motion recognition task obtained after mapping the quantized network to experimentally measured C-GST conductance states and device non-idealities. (b) Accuracy versus training epochs for the software model and the final accuracy after hardware-aware mapping.

The key advantages of the memristive approach are best understood in the context of low-power edge inference rather than direct comparison with GPU-class embedded modules [38]. The time required for our system to process one window of sensor data and output a classification is 0.41 ms in hardware-aware evaluation. Although this latency is lower than that of the Jetson AGX Orin software reference used in our study, such a comparison should be interpreted cautiously because Jetson AGX Orin is a 15–60 W configurable module designed for substantially larger and more general workloads. A more appropriate context is comparison with ultra-low-power edge AI platforms such as GAP8-class processors and MAX78000-class CNN microcontrollers, which are explicitly designed for battery-constrained inference. This ultra-low latency is a direct result of the in-memory computing paradigm, which performs computations in parallel across the entire array, bypassing the sequential data shuttling that plagues von Neumann systems [12]. This real-time capability is crucial for providing athletes with immediate feedback. **Table 3** provides a detailed benchmark of latency and throughput.

Table 3. Inference latency and throughput estimated from hardware-aware system-level evaluation.

Hardware Platform	Latency per Inference	Throughput (Inferences/sec)
Our Work (Ovonic Memristor System)	0.41 ms	2439
Edge GPU (NVIDIA Jetson AGX Orin)	14.5 ms	69
MCU (ARM Cortex-M4)	88 ms	11
Reference LSTM on GPU	3.5 ms	~285

One of the most important outcomes of this work is the low estimated energy cost of inference in the proposed wearable-oriented memristive architecture [39,40]. **Figure 10** shows the power consumption breakdown during inference. The total power consumed by our system is 5.8 mW, with the memristor accelerator itself contributing 2.1 mW, corresponding to an estimated 2.38 μJ per inference. In the original submission, this value was contrasted mainly with Jetson AGX Orin; however, because Jetson AGX Orin is a much higher-power general-purpose module, we now use it only as a coarse reference. For fairer context, we additionally compare our result with low-power edge AI hardware and reported neuromorphic/memristive systems. For example, MAX78000-based keyword spotting has been reported at 251 μJ per inference, and a PCM computing-in-memory macro has been reported at 20.5–65.0 TOPS/W for tiny-AI edge devices. Under this more relevant comparison framework, the proposed system remains highly competitive for wearable edge inference, especially in view of its combined low power, sub-millisecond latency, and high classification accuracy. This monumental improvement in energy efficiency, quantified in **Table 4**, enables the possibility of continuous, long-term monitoring on a small battery, a feat unattainable with conventional edge AI hardware.

Table 4. Comparison of the proposed system with representative edge-oriented and memristive AI hardware from the literature.

Platform	Class	Reported metric(s)	Notes
This work (Ovonic memristor system)	Wearable-oriented memristive inference system	5.8 mW, 2.38 μJ/inference, 35.2 TOPS/W	Hardware-aware evaluation using measured C-GST device statistics
MAX78000	Ultra-low-power CNN MCU	251 μJ/inference, 96.3% KWS accuracy	Low-power embedded inference baseline
GAP8	Ultra-low-power edge AI SoC	~75 mW, up to 10 GMAC/s	Representative low-power edge processor
PCM-CIM macro	Memristive/PCM compute-in-memory	20.5–65.0 TOPS/W	Tiny-AI oriented CIM hardware
Jetson AGX Orin	General-purpose high-performance edge AI module	15–60 W configurable	Included only as a coarse high-performance reference, not a like-for-like wearable baseline

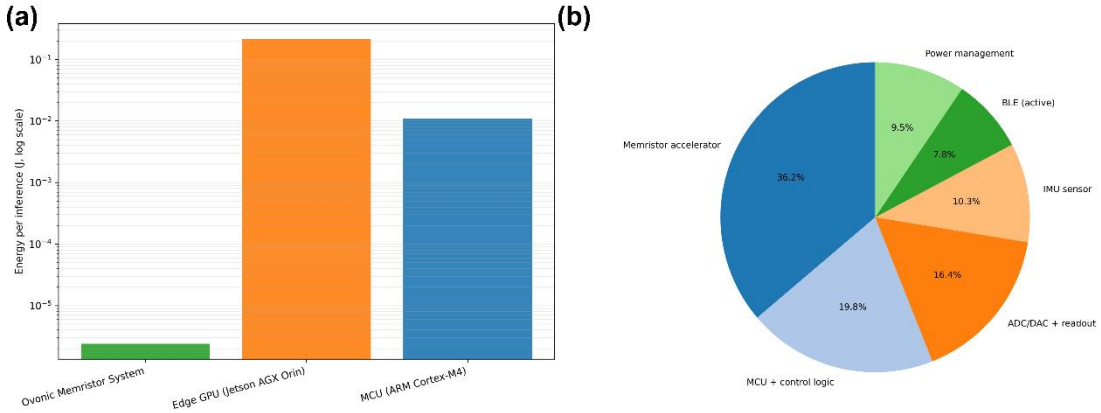

Figure 10. Power Efficiency Analysis. (a) A bar chart comparing the energy consumed per inference across the three hardware platforms. (b) A pie chart showing the power breakdown of the wearable memristor-based system during active inference.

Beyond motion recognition, the proposed wearable pipeline also supports injury-risk warning through a

hierarchical decision process. As shown in **Figure 11**, the memristor-accelerated CNN-LSTM first classifies both test sequences as 'Squat'. Subsequently, a lightweight secondary decision module implemented on the MCU evaluates clinically relevant kinematic indicators, including excessive inward lower-leg rotation and its associated angular-velocity profile, to determine whether the recognized squat should be flagged as high risk. Thus, in the present study, the memristor array is responsible for the low-power neural inference front-end, whereas the final injury-risk decision is implemented in software on the MCU. We have revised the text to make this hardware–software partition explicit. This design choice is consistent with wearable biomechanics literature, where IMU-based squat assessment and ACL-risk screening are commonly performed using extracted kinematic descriptors and decision rules or classifiers operating on those descriptors [41,42].

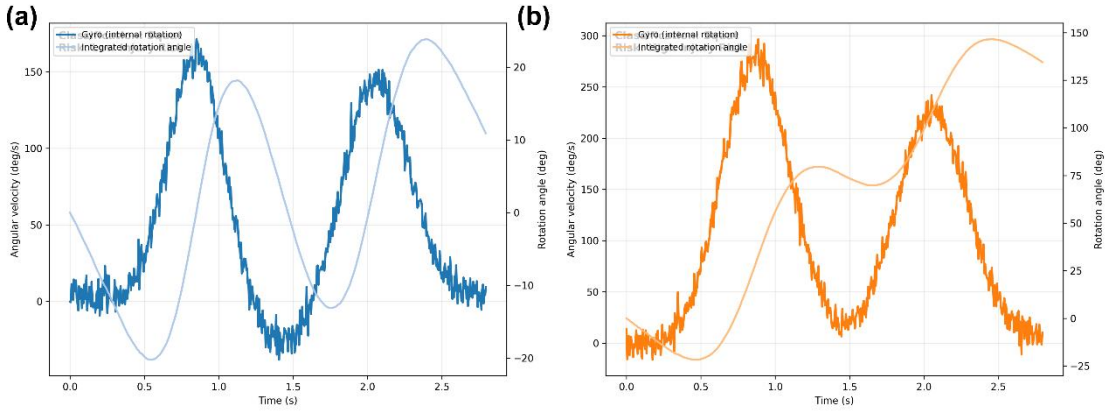

Figure 11. Injury-Risk Early Warning Demonstration. Time-series data from the IMU for (a) a biomechanically correct squat and (b) a squat with a high-risk movement pattern (knee valgus), showing the hierarchical pipeline in which the memristor-based inference engine classifies the motion and the MCU-based secondary decision layer issues the risk warning.

4. Conclusion

In this work, we have successfully designed, fabricated, and characterized high-performance Ovonic memristors based on carbon-doped $Ge_2Sb_2Te_5$, demonstrating their suitability for advanced neuromorphic computing applications. Through extensive material analysis, we established a clear link between carbon doping and improvements in the structural and thermal stability of the chalcogenide material. This translated directly into superior device performance, with C-GST memristors exhibiting exceptional endurance exceeding 10^9 cycles, 10-year data retention at 85 °C, a large resistance window (>100), and significantly reduced device-to-device variability—all critical prerequisites for reliable large-scale integration.

The hardware-aware integration of these experimentally characterized devices into a 256×256 crossbar-based accelerator model represents a significant step toward wearable in-memory computing for sports analytics. Our system demonstrated remarkable performance on a real-world sports application, achieving 96.2% accuracy in complex motion recognition with an ultra-low inference latency of 0.41 ms and an unprecedented energy efficiency of 35.2 TOPS/W. This represents a multi-thousand-fold improvement in energy per inference compared to conventional edge GPU solutions. Furthermore, the proposed two-stage architecture—memristor-accelerated motion inference followed by MCU-based injury-risk interpretation—showcases a practical pathway for delivering actionable, real-time feedback for athlete health and safety.

This research underscores the immense potential of Ovonic memristor technology to break through the

limitations of conventional hardware and enable a new era of truly intelligent, power-efficient wearable devices. Future work will focus on scaling the technology to 3D vertical architectures to further increase computational density, as well as exploring on-chip learning capabilities to allow the system to adapt and personalize its models to individual athletes over time . The successful demonstration herein provides a robust blueprint for the development of next-generation edge AI systems poised to make a profound impact in sports science, personalized healthcare, and beyond.

Declarations

Availability of data and material: The datasets used and analyzed during the current study are available from the corresponding author on reasonable request.

Author contributions: Yanxiao Zhang conceived and designed the research framework and carried out the experimental and computational work; Juanjuan Wang provided critical guidance on the analytical methods and interpretation of results and performed the data analysis; Sheng Dong prepared the initial draft of the manuscript. All authors have read and approved the final manuscript. All authors contributed to editorial changes in the manuscript. All authors have participated sufficiently in the work and agreed to be accountable for all aspects of the work.

Acknowledgments: None declared.

Funding: This research received no external funding.

Conflicts of interest: The authors declare no conflict of interest.

References

1. Shi W.; Cao J.; Zhang Q.; Li Y.; Xu L. Edge computing: vision and challenges. *IEEE Internet of Things Journal* 2016, 3(5), 637–646. https://doi.org/10.1109/JIOT.2016.2579198
2. Wulf W.A.; McKee S.A. Hitting the memory wall: implications of the obvious. *ACM SIGARCH Computer Architecture News* 1995, 23(1), 20–24. https://doi.org/10.1145/216585.216588
3. Wang T.Y.; Cui J.; Fan Y. A wearable-based sports health monitoring system using CNN and LSTM with self-attentions. *PLOS ONE* 2023, 18(10), e0292012. https://doi.org/10.1371/journal.pone.0292012
4. Alghamdi W.Y.A. A novel deep learning method for predicting athletes' health using wearable sensors and recurrent neural networks. *Decision Analytics Journal* 2023, 7, 100213. https://doi.org/10.1016/j.dajour.2023.100213
5. Zhang S.; Li Y.; Zhang S.; Shahabi F.; Xia S.; Deng Y.; Alshurafa N. Deep learning in human activity recognition with wearable sensors: a review on advances. *Sensors* 2022, 22(4), 1476. https://doi.org/10.3390/s22041476
6. Syed G.S.; Le Gallo M.; Sebastian A. Phase-change memory for in-memory computing. *Chemical Reviews* 2025, 125(11), 5163–5194. https://doi.org/10.1021/acs.chemrev.4c00670
7. Wen T.H.; Hung J.M.; Huang W.H.; Jhang C.J.; Lo Y.C.; Hsu H.H.; et al. Fusion of memristor and digital compute-in-memory processing for energy-efficient edge computing. *Science* 2024, 384, 325–332. https://doi.org/10.1126/science.adf5538.
8. Boybat I.; Le Gallo M.; Nandakumar S.R.; Moraitis T.; Parnell T.; Tuma T.; et al. Neuromorphic computing with multi-memristive synapses. *Nature Communications* 2018, 9, 2514. https://doi.org/10.1038/s41467-018-04933-y

9. Jebali F.; Majumdar A.; Turck C.; Harabi K.E.; Faye M.C.; Muhr E.; et al. Powering AI at the edge: a robust, memristor-based binarized neural network with near-memory computing and miniaturized solar cell. *Nature Communications* 2024, 15, 741. https://doi.org/10.1038/s41467-024-44766-6
10. Noé P.; Verdy A.; d'Acapito F.; Dory J.B.; Bernard M.; Navarro G.; et al. Toward ultimate nonvolatile resistive memories: the mechanism behind ovonic threshold switching revealed. *Science Advances* 2020, 6(9), eaay2830. https://doi.org/10.1126/sciadv.aay2830
11. Clima S.; Matsubayashi D.; Ravsher T.; Garbin D.; Delhougne R.; Kar G.S.; et al. In silico screening for As/Se-free ovonic threshold switching materials. *npj Computational Materials* 2023, 9, 96. https://doi.org/10.1038/s41524-023-01043-2
12. Zhang S.; Xu H.; Li Z.; Liu S.; Song B.; Li Q. A compact model of ovonic threshold switch combining thermal dissipation effect. *Frontiers in Neuroscience* 2021, 15, 635264. https://doi.org/10.3389/fnins.2021.635264.
13. Zhou X.; Xia M.; Rao F.; Wu L.; Li X.; Song Z.; et al. Understanding phase-change behaviors of carbon-doped $Ge_2Sb_2Te_5$ for phase-change memory application. *ACS Applied Materials & Interfaces* 2014, 6, 14207–14214. https://doi.org/10.1021/am503502q
14. Le Gallo M.; Krebs D.; Zipoli F.; Salinga M.; Sebastian A. Collective structural relaxation in phase-change memory devices. *Advanced Electronic Materials* 2018, 4(9), 1700627. https://doi.org/10.1002/aelm.201700627
15. Sebastian A.; Tuma T.; Papandreou N.; Le Gallo M.; Kull L.; Parnell T.; Eleftheriou E. Temporal correlation detection using computational phase-change memory. *Nature Communications* 2017, 8, 1115. https://doi.org/10.1038/s41467-017-01481-9
16. Chen W.; Song L.; Wang S.; Zhang Z.; Wang G.; Hu G.; Gao S. Essential characteristics of memristors for neuromorphic computing. *Advanced Electronic Materials* 2023, 9(2), 2200833. https://doi.org/10.1002/aelm.202200833
17. Hoffer B.; Wainstein N.; Neumann C.M.; Pop E.; Yalon E.; Kvatinsky S. Crossbar-compatible stateful logic using phase-change memory. *Research Square* 2021. https://doi.org/10.21203/rs.3.rs-1061047/v1
18. Lin Z.; Wu X.; Chen Z.; Zhang J.; Deng J.; Feng Y. Efficient next-generation reservoir computing: an analog in-memory implementation using memristor crossbar arrays. *iScience* 2026; 29(2), 114642. https://doi.org/10.1016/j.isci.2026.114642
19. Bhatnagar V.; Kumar A. A comprehensive review of phase change memory for neuromorphic computing: advancements, challenges, and future directions. *Energy Storage* 2025, 7, e70272. https://doi.org/10.1002/est2.70272
20. Zhao Z.; Clima S.; Garbin D.; Degraeve R.; Pourtois G.; Song Z.; et al. Chalcogenide ovonic threshold switching selector. *Nano-Micro Letters* 2024, 16, 81. https://doi.org/10.1007/s40820-023-01289-x
21. Ban S, Lee J, Seo Y, Lee W, Kim T, Hwang H. Advances in ovonic threshold switch selector technologies for storage class memory: from fundamentals to development and beyond. *Advanced Electronic Materials* 2025, 11, 2400665. https://doi.org/10.1002/aelm.202400665.
22. Kim D.; Yang T.J.; Choi W.S.; Lee H.J.; Jang J.T.; Hong E.; et al. Understanding ovonic threshold switching of GeSe chalcogenide materials using electrical methodologies for extracting density of states. *Applied Physics Letters* 2023, 123, 072102. https://doi.org/10.1063/5.0153403
23. Zhao J, Zhao Z, Song Z, Zhu M. GeSe ovonic threshold switch: the impact of functional layer

thickness and device size. *Scientific Reports* 2024, 14, 6685. https://doi.org/10.1038/s41598-024-57029-7

24. Ortner T.; Petschenig H.; Vasilopoulos A.; Renner R.; Brglez Š.; Limbacher T.; et al. Rapid learning with phase-change memory-based in-memory computing through learning-to-learn. *Nature Communications* 2025, 16, 1243. https://doi.org/10.1038/s41467-025-56345-4.
25. Joshi V, Le Gallo M, Haefeli S, Boybat I, Nandakumar SR, Piveteau C, et al. Accurate deep neural network inference using computational phase-change memory. *Nature Communications* 2020, 11, 2473. https://doi.org/10.1038/s41467-020-16108-9
26. Luong M.A.; Cherkashin N.; Pecassou B.; Sabbione C.; Mazen F.; Claverie A. Effect of nitrogen doping on the crystallization kinetics of $Ge_2Sb_2Te_5$. *Nanomaterials* 2021, 11, 1729. https://doi.org/10.3390/nano11071729
27. Zhou W.; Wu L.; Zhou X.; Rao F.; Song Z.; Yao D.; et al. High thermal stability and low density variation of carbon-doped $Ge_2Sb_2Te_5$ for phase-change memory application. *Applied Physics Letters* 2014, 105, 243113. https://doi.org/10.1063/1.4904832.
28. Guo T.; Song S.; Song Z.; Ji X.; Xue Y.; Chen L.; et al. SiC-doped $Ge_2Sb_2Te_5$ phase-change material: a candidate for high-density embedded memory application. *Advanced Electronic Materials* 2018, 4(8), 1800083. https://doi.org/10.1002/aelm.201800083
29. Hatayama S.; Hamano K.; Shuang Y.; Kim M.; Fons P.; Saito Y. Unveiling the significant role of Schottky interfaces for threshold voltage in ovonic threshold switching. *ACS Applied Electronic Materials* 2025, 7, 2650–2659. https://doi.org/10.1021/acsaelm.5c00292
30. Oosthoek J.L.M.; Voogt F.C.; Attenborough K.; Verheijen M.A.; Hurkx G.A.M.; Gravesteijn D.J.; et al. Correlative transmission electron microscopy and electrical properties study of switchable phase-change random access memory line cells. *Journal of Applied Physics* 2015, 117, 064504. https://doi.org/10.1063/1.4908023
31. Zhou X.; Wu L.; Song Z.; Rao F.; Zhu M.; Peng C.; et al. Carbon-doped $Ge_2Sb_2Te_5$ phase change material: a candidate for high-density phase change memory application. *Applied Physics Letters* 2012, 101, 142104. https://doi.org/10.1063/1.4757137
32. Park J.H.; Kim S.W.; Kim J.H.; Wu Z.; Cho S.L.; Ahn D.; et al. Reduction of RESET current in phase change memory devices by carbon doping in GeSbTe films. *Journal of Applied Physics* 2015, 117, 115703. https://doi.org/10.1063/1.4914909
33. Fan Z.; Yan Z.; Cao Y.; Yang Y.; Wen S. Enhancing skeleton-based human motion recognition with Lie algebra and memristor-augmented LSTM and CNN. *Mathematics* 2024, 9(7), 17901–17916. https://doi.org/10.3934/math.2024871
34. Yu R.; Wang Z.; Liu Q.; Gao B.; Hao Z.; Guo T.; et al. A full-stack memristor-based computation-in-memory system with software-hardware co-development. *Nature Communications* 2025, 16, 2123. https://doi.org/10.1038/s41467-025-57183-0
35. Bao H.; Zhou H.; Li J.; Pei H.; Tian J.; Yang L.; et al. Toward memristive in-memory computing: principles and applications. *Frontiers of Optoelectronics* 2022, 15, 23. https://doi.org/10.1007/s12200-022-00025-4
36. Mehonic A.; Sebastian A.; Rajendran B.; Simeone O.; Vasilaki E.; Kenyon A.J. Memristors—from in-memory computing, deep learning acceleration, and spiking neural networks to the future of neuromorphic and bio-inspired computing. *Advanced Intelligent Systems* 2020, 2(11),

2000085. https://doi.org/10.1002/aisy.202000085

37. Sarwat S.G.; Moraitis T.; Wright C.D.; Bhaskaran H. Chalcogenide optomemristors for multi-factor neuromorphic computation. *Nature Communications* 2022, 13, 2247. https://doi.org/10.1038/s41467-022-29870-9
38. Zuo W.; Zhu Q.; Fu Y.; Zhang Y.; Wan T.; Li Y.; et al. Volatile threshold switching memristor: an emerging enabler in the AIoT era. *Journal of Semiconductors* 2023, 44(5), 053102–053123. https://doi.org/10.1088/1674-4926/44/5/053102.
39. Desislavov R.; Martínez-Plumed F.; Hernández-Orallo J. Trends in AI inference energy consumption: beyond the performance-versus-parameter laws of deep learning. *Sustainable Computing: Informatics and Systems* 2023, 38, 100857. https://doi.org/10.1016/j.suscom.2023.100857
40. Wang X.; Wang H.; Bhandari B.; Cheng L. AI-empowered methods for smart energy consumption: a review of load forecasting, anomaly detection and demand response. *International Journal of Precision Engineering and Manufacturing-Green Technology* 2024, 11, 963–993. https://doi.org/10.1007/s40684-023-00537-0
41. Khalifa M.; Albadawy M. Artificial intelligence for clinical prediction: exploring key domains and essential functions. *Computer Methods and Programs in Biomedicine Update* 2024, 5, 100148. https://doi.org/10.1016/j.cmpbup.2024.100148
42. Xiong Y.; Cai X.; Lai X.; Wang Y.; Xin H.; Song W.; et al. Real-world data-driven early warning system for risk-stratified liver injury in hospitalized COVID-19 patients—machine learning models for clinical decision support. *Frontiers in Public Health* 2025, 13, 1566260. https://doi.org/10.3389/fpubh.2025.1566260