Leveraging Intrinsic Physics of Spintronic Devices for Neuromorphic Computing

(Invited Special Session Paper)

Md Zesun Ahmed Mia and Abhronil Sengupta
School of Electrical Engineering and Computer Science
The Pennsylvania State University
University Park, PA 16802, USA
Email: {zesun.ahmed, sengupta}@psu.edu

***Abstract*—Neuromorphic computing has made significant strides over the past few years by mapping functional primitives of neurons and synapses to nanoelectronic devices, yet most implementations do not tap into the temporal non-linear dynamics of the devices. We argue that spintronic devices, specifically magnetic tunnel junctions (MTJs), offer a palette of intrinsic physical phenomena that can serve as native computational primitives. Spin-transfer torque (STT), spin-orbit torque (SOT), voltage-controlled magnetic anisotropy (VCMA), magnetoelectric effect (ME) and stochastic superparamagnetic dynamics each enable functionalities difficult to replicate in deterministic CMOS. We underscore this perspective through two specific system-level instances focusing on the synapse and neuron: selectorless synaptic crossbar arrays that exploit a VCMA-driven switching-mechanism to eliminate per-cell selectors, and temporal information encoding through superparamagnetic neuronal MTJs that achieves significant network-level spiking sparsity via stochastic dynamics. Building on these exemplars, we envision how these primitives compose with the broader spintronic ecosystem — domain-wall synapses, oscillatory primitives, and stochastic neurons — toward neuromorphic systems that are simultaneously sparser, denser, and more functionally diverse.**

***Index Terms*—Neuromorphic Computing, Spintronics, Hardware-Software Co-Design.**

I. Introduction

Neuromorphic computing has emerged as a compelling paradigm for energy-efficient intelligence, drawing inspiration from the brain's architecture to process information in ways fundamentally different from conventional von Neumann machines. Driven by advances in non-volatile memory technologies, the field has demonstrated in-memory computing with resistive, phase-change, ferroelectric and spintronic devices, enabling massively parallel multiply-accumulate operations directly at the data storage site [1]. Dedicated neuromorphic processors such as Intel's Loihi 2 have further shown that co-locating computation and memory at the architectural level delivers orders-of-magnitude efficiency gains for spiking neural network workloads [2]. Across these efforts, a common engineering pattern has emerged: device-level physics — noise, nonlinear switching dynamics, stochastic behavior — is typically treated as a limitation to be suppressed or compensated, rather than utilized as a computational resource.

Spintronic devices challenge this pattern. Magnetic tunnel junctions (MTJs) host a rich set of intrinsic physical phenomena — spin-transfer torque (STT), spin-orbit torque (SOT), voltage-controlled magnetic anisotropy (VCMA), magnetoelectric coupling (ME), and stochastic superparamagnetic dynamics — each arising naturally from the underlying material and device physics [3]–[5]. Rather than translating these phenomena solely into device-efficiency metrics, a growing body of work shows that they can act directly as computational primitives: SOT provides a decoupled write path for neuronal activation or synaptic plasticity [6], domain-wall MTJs support multilevel synaptic behavior [7], [8], VCMA provides voltage-tunable anisotropy for selective synaptic switching [9], and superparamagnetic switching maps naturally onto probabilistic spiking behavior [10].

In this perspective, we argue that the next generation of spintronic neuromorphic hardware should be designed *with* intrinsic device physics, not *around* it. We ground this argument in two canonical neuromorphic roles: at the synaptic level, intrinsic switching-mechanism transitions enable inherent write selectivity in dense crossbar arrays; at the neuronal level, stochastic magnetization dynamics provide a natural substrate for temporal spike encoding. Both functionalities are difficult or prohibitively expensive to replicate with deterministic CMOS.

Section II establishes the MTJ structures, operating regimes, and control mechanisms needed for the two neuromorphic primitives emphasized here, while Section III grounds them in applications related to selectorless synaptic arrays and temporal neurons. Section IV then distills the broader design perspective and the near-term challenges for spintronic neuromorphic hardware.

II. Spintronic Phenomena for Neuromorphic Primitives

Neuromorphic hardware must implement two core computational roles: weighted summation across synapses and conversion of the aggregated signal into a neuronal response. In crosspoint arrays, the first role appears naturally as dot-product or vector-matrix multiplication, with inputs applied along the rows and weighted currents summed along the columns. For spiking neuromorphic architectures, the second can be realized through rate- or time-based encoding of an output spike signal; this perspective emphasizes the stochastic temporal regime

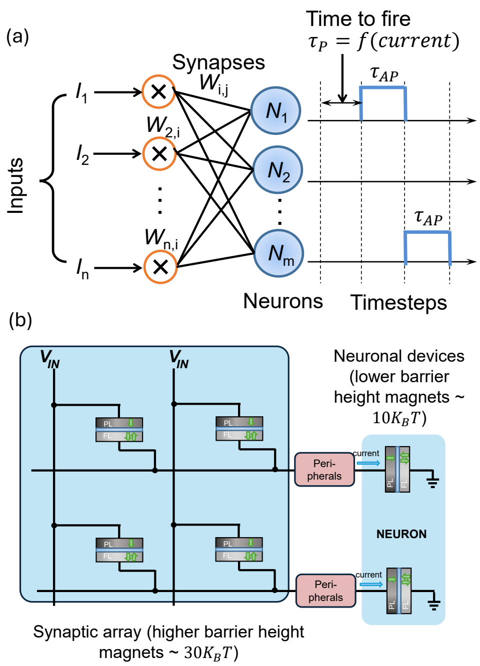

Fig. 1. From neuromorphic primitives to a spintronic hardware mapping: (a) Inputs are multiplied by synaptic weights and accumulated at each neuron; the resulting signal controls neuronal firing time through input-dependent switching statistics. (b) Weighted summation is realized in a crosspoint array of higher-barrier MTJ synapses, and the resulting column current is routed by peripheral circuitry to lower-barrier MTJ neurons, whose stochastic switching determines the output spike timing.

because it exposes functionality native to MTJ dynamics. Fig. 1 summarizes this mapping from neuromorphic primitives to a spintronic hardware realization, while Fig. 2 introduces the representative MTJ-derived structures and control routes discussed below.

A. MTJ Device Structures and Operating Regimes

A magnetic tunnel junction [Fig. 2(a)] consists of two ferromagnetic layers — a pinned layer (PL) with fixed magnetization and a free layer (FL) with switchable magnetization — separated by a thin oxide barrier, typically MgO. The device resistance depends on the relative magnetization orientation: low resistance in the parallel (P) state and high resistance in the anti-parallel (AP) state. The MTJ therefore provides a compact magnetic element whose internal state can be read electrically.

The magnetization dynamics of the FL is governed by the Landau-Lifshitz-Gilbert (LLG) equation:

$$\frac{d\hat{\boldsymbol{m}}}{dt} = -\gamma(\hat{\boldsymbol{m}}\times\boldsymbol{H}_{\text{eff}})+\alpha\left(\hat{\boldsymbol{m}}\times\frac{d\hat{\boldsymbol{m}}}{dt}\right)+\frac{1}{qN_s}(\hat{\boldsymbol{m}}\times\boldsymbol{I}_s\times\hat{\boldsymbol{m}}) \tag{1}$$

where $\hat{\boldsymbol{m}}$ is the unit magnetization vector of the FL, γ is the gyromagnetic ratio, α is the Gilbert damping constant, $\boldsymbol{H}_{\text{eff}}$ is the effective magnetic field (including anisotropy, demagnetization, and thermal noise contributions), q is the elementary charge, N_s is the number of spins in the FL, and $\boldsymbol{I}_s$ is the spin current. The three terms correspond to precession around $\boldsymbol{H}_{\text{eff}}$, damping toward equilibrium, and spin-torque-driven switching, respectively.

In the canonical STT geometry, a spin-polarized current flowing through the MTJ stack exerts torque on the FL magnetization, driving it from one stable state to the other when the current exceeds a critical threshold. At sufficiently large energy barrier, this behavior supports deterministic binary memory. As the FL is scaled and the barrier is reduced to a low-barrier regime, however, thermal fluctuations begin to dominate and the device crosses into a superparamagnetic regime in which spontaneous transitions between the P and AP states become frequent. The mean dwell times in the two states — τ_P and τ_{AP}, sketched for the neuron in Fig. 1(a) — then become the natural state variables of the device. Rather than viewing this stochasticity only as lost retention, it can be treated as a computational resource for spike-based computation [11], [12].

Beyond the basic monodomain MTJ, extending the free layer along one dimension allows the magnetic state to be represented by domain-wall position rather than by a single binary orientation [Fig. 2(b)]. Because the conductance depends on the fraction of the device in each magnetic configuration, domain-wall MTJs naturally support multilevel conductance states and thus map naturally onto synaptic weight storage [7], [8]. Other spin textures like skyrmions can also enable synaptic functionality [13]. Once the MTJ structure and its operating regimes are established, the remaining question is how to selectively access or manipulate them. SOT, VCMA, and ME control provide three such routes, among others.

B. Emerging MTJ Switching Mechanisms

Among these routes, SOT decouples write and read paths by placing the MTJ on a spin-orbit-coupled heavy-metal underlayer [Fig. 2(c)]. A charge current through the heavy metal generates a transverse spin current that switches the FL without directing the write current through the tunnel barrier [5], [14]. This geometry improves device endurance and allows the write path and read resistance to be optimized more independently than in STT switching.

VCMA provides a different handle: an applied electric field at the ferromagnet/oxide interface modulates the perpendicular magnetic anisotropy [Fig. 2(d)]. The effective interfacial anisotropy varies as:

$$K_i^{\text{eff}}(U) = K_i - \xi\frac{U}{t_{\text{OX}}} \tag{2}$$

where K_i is the intrinsic interfacial anisotropy energy density, ξ is the VCMA coefficient (typically $\sim$100 fJ/(V·m)), U is the applied voltage, and t_{OX} is the oxide thickness [9], [15]. When the applied voltage is sufficient to suppress the perpendicular anisotropy, the FL magnetization is no longer confined along the out-of-plane axis and instead undergoes precessional dynamics. As the FL precesses, it periodically passes through orientations that favor either the P or the AP state. The switching outcome is therefore set by the point in the precessional cycle at which the voltage pulse terminates, making the switching probability a sensitive function of pulse width. For neuromorphic applications, the critical consequence

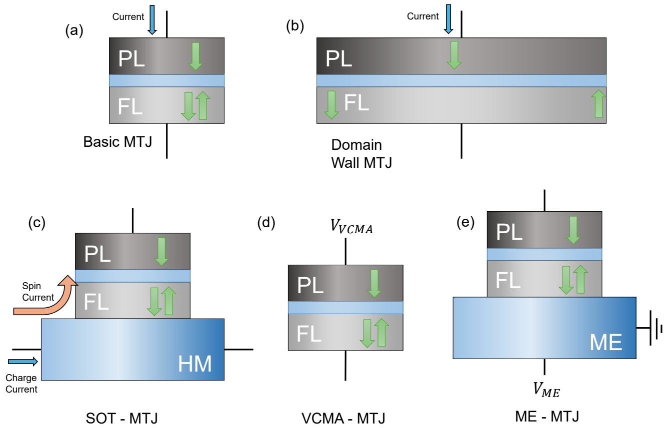

Fig. 2. Representative MTJ-derived device structures and control routes relevant to spintronic neuromorphic hardware. (a) Basic monodomain MTJ with pinned layer (PL), free layer (FL), and tunnel barrier. (b) Domain-wall MTJ, in which an extended free layer supports position-dependent conductance and thus multilevel synaptic behavior. (c) SOT-MTJ, where a heavy-metal (HM) underlayer converts charge current into a transverse spin current and decouples write and read paths. (d) VCMA-MTJ, where an applied voltage V_{VCMA} modulates interfacial anisotropy and enables pulse-width-sensitive precessional switching. (e) ME-MTJ, where a magnetoelectric layer, driven by V_{ME}, provides an additional voltage-controlled handle on the magnetic state.

is that VCMA-dominated switching can occur on nanosecond-scale pulse widths, whereas pure STT switching at the same short pulse width has negligible probability [9], [15]. Section III-A shows how this timing asymmetry enables inherent write selectivity in dense crossbar arrays.

Beyond torque-driven and anisotropy-mediated switching, the MTJ platform also admits an additional field-like control path through magnetoelectric coupling. An ME layer integrated beneath the MTJ provides an additional voltage-controlled handle on the magnetic state [Fig. 2(e)]. The voltage across the ME layer generates an effective magnetic field-like control on the FL. In the present context, the importance of this added terminal is that it introduces a control axis distinct from the spin-current path through the PL; Section III-B shows how that extra degree of freedom enables temporal encoding.

III. Translating Physics to Primitives: Two System-Level Implementations

This section translates the phenomena introduced in Section II into functional hardware through two case studies — one targeting the synaptic primitive of weighted summation and the other the neuronal primitive of temporal encoding introduced conceptually in Fig. 1.

A. Sneak-Path Mitigation via VCMA

The synaptic array, sketched in Fig. 1(b), is naturally realized as a crossbar, enabling single-step vector–matrix multiplication by exploiting Ohm's law and Kirchhoff's current law along shared bitlines [1]. However, writing to a selected cell inevitably applies partial voltages to unselected cells sharing the same row or column, producing sneak-path disturbances that corrupt stored synaptic weights. Conventional mitigation — adding a selector transistor (one transistor–one resistor, 1T-1R), diode (1D-1R), or self-rectifying element per cell — sacrifice the crossbar's density advantage.

Yang and Sengupta [9] showed that in spintronic crossbars, this problem can be solved at the device physics level by exploiting the mechanism described in Section II-B. The key insight is that VCMA-assisted and STT switching have fundamentally different pulse-width requirements. At the full programming voltage, the applied field suppresses PMA sufficiently to initiate in-plane precessional dynamics, and the simultaneous STT assists the switching — completing within $\sim$1.8 ns [9]. At half voltage, PMA suppression is insufficient to trigger the VCMA effect, leaving only the STT torque, which at the same short pulse width has negligible switching probability. This is not merely a quantitative voltage margin but a qualitative change in the underlying physics: the selected cell switches through a fast precessional mechanism while half-selected cells remain in a slow, STT-dominated regime that cannot complete within the pulse window.

When scaled to an array level, this pulse-width selectivity remains robust to device-to-device free-layer thickness variations, provided the operating voltage preserves deterministic precessional trajectories [9]. Although the selected-cell switching probability is not unity, network-level simulations show that on-chip learning compensates for such write errors with negligible accuracy loss [9]. The key lesson is that selectivity in spintronic crossbars need not rely on voltage-amplitude margins alone; a change in the underlying switching mechanism provides an intrinsically sharper distinction between target and non-target cells.

B. Temporal Encoding via Superparamagnetic MTJs

The weighted summation in Fig. 1(a) and (b) addresses the synaptic side of the computation; the remaining question is how the accumulated signal is converted into an output spike. In standard rate-coded spiking networks, information is carried by spike counts over a time window — a mapping that scales poorly with network depth, as spiking activity grows with each layer [16]. Temporal coding, where information resides in precise spike timing, offers a sparser alternative but demands neurons whose first-spike latency is a controllable function of the input [17].

As discussed in Section II-A, lowering the MTJ barrier drives the device into a stochastic regime characterized by the dwell times τ_P and τ_{AP}. In a standard two-terminal superparamagnetic MTJ, a single spin current mainly modulates the switching statistics, a mapping better aligned with rate-based stochastic neurons [18]. The three-terminal ME-MTJ, introduced in Section II-B, provides a richer temporal primitive because the ME terminal and the spin-current terminal can be used to tune τ_{AP} and τ_P separately. In the operating scheme proposed by Yang *et al.* [19], a shared bias establishes a common reference timescale through τ_{AP} (duration of each timestep), while each neuron's weighted input independently modulates τ_P; the first neuron to fire determines the classification output. Stronger inputs reduce τ_P, causing the device to leave the P state, and hence fire sooner. The result is a direct mapping from input strength to first-spike latency: a time-to-first-spike neuron realized through stochastic switching dynamics.

This approach has three noteworthy properties. First, a binary two-state MTJ can encode multibit information through *timing*: the exponential lifetime distribution maps a continuous-valued latency onto a nominally single-bit device, eliminating the need for explicit multilevel cells. Second, in the networks evaluated [19], temporal coding preserved sparse activity across depth, with only 2–3 spikes per neuron required for inference [19]. Third, Yang *et al.* derived the training loss directly from the exponential waiting-time statistics of the stochastic device, so the learning rule followed the hardware model rather than imposing an external abstraction [19].

Taken together, these properties make low-barrier spintronic neurons a natural substrate for sparse temporal coding.

IV. Perspective

The two case studies discussed above point to a broader lesson: spintronic neuromorphic hardware should be viewed not as a single device proposal, but as a wider design space in which distinct magnetic phenomena give rise to different computational primitives. Beyond the VCMA- and ME-based examples emphasized here, prior work has explored multi-level synaptic behavior in domain-wall devices [7], oscillatory dynamics for associative computing [20], and stochastic p-bits for unconventional computing [12], [21]. Other emergent materials like antiferromagnets (AFMs), ferrimagnets and spin textures like skyrmions, spin waves are being actively investigated [22]. Taken together, these studies delineate a broader spintronic design space in which distinct magnetic mechanisms can be harnessed for complementary neuromorphic functions (Fig. 3).

Realizing that broader repertoire will require a corresponding shift in algorithm design. Rather than forcing idealized deterministic abstractions onto the hardware, learning and inference strategies should be formulated around native device behavior: lifetime-statistics-based objectives for temporal encoding (Section III-B), learning rules that tolerate the intrinsic write uncertainty of dense selectorless crossbars (Section III-A), and probabilistic frameworks that exploit thermal fluctuations as a computational resource [10], [23].

Several challenges remain. First, the case studies emphasized here still rely heavily on idealized device-level modeling, and validation under experimentally realistic operating conditions is needed. Second, device-physics-aware training has thus far been demonstrated only on relatively small benchmarks [9], [19], leaving its scalability to larger workloads unresolved. Third, the field needs evaluation metrics that capture what spintronic hardware uniquely offers — for example, temporal sparsity, selectorless write fidelity, and stochastic inference capability — rather than relying solely on conventional static-task accuracy.

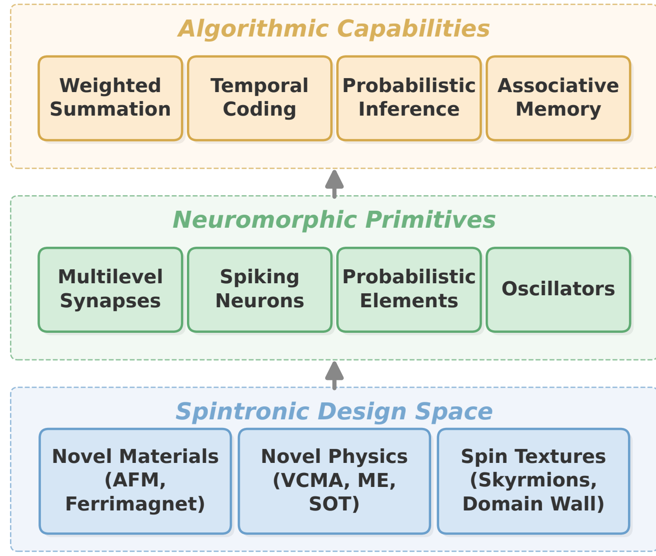

Fig. 3. Hierarchical view of the spintronic neuromorphic ecosystem. Three broad spintronic directions — novel materials (AFMs, ferrimagnets), novel switching physics (VCMA, ME, SOT), and spin textures (skyrmions, domain walls) — form the foundation (bottom) from which reusable neuromorphic building blocks (middle) emerge. These primitives collectively support a range of algorithmic capabilities (top). Algorithms designed to exploit these native primitives — rather than abstract them into generic models — can access functionalities that deterministic CMOS cannot efficiently replicate.

V. Conclusions

The two case studies examined here point to a common design lesson: in spintronic neuromorphic hardware, intrinsic device physics need not be suppressed to emulate computation, but can instead be used directly as the computational primitive. VCMA-enabled write selectivity in dense crossbars and ME-assisted control of stochastic switching in low-barrier MTJs illustrate how mechanism-level magnetic dynamics can be translated into synaptic and neuronal functionality. The next step is therefore not merely to substitute spintronic devices for CMOS analogues, but to align algorithms, training objectives, metrics, and experimental validation with the native selectivity, temporal statistics, and stochasticity of the hardware itself.

Acknowledgment

The authors acknowledge the use of generative AI tools in the preparation of this manuscript. Specifically, GPT-5.4 (OpenAI) and Claude 4.6 (Anthropic) were used for summarization and wording refinement, limited figure-script edits, and identifying potentially relevant prior work. Any generated content was reviewed and verified by the authors; all technical claims, cited references, and final text remain the authors' responsibility.

This material is based upon work supported in part by the U.S. National Science Foundation under award No. CCSS #2333881, ECCS #2028213 and CCF #1955815 and by Oracle Cloud credits and related resources provided by the Oracle for Research program.

References

[1] I. Chakraborty, A. Jaiswal, A. Saha, S. K. Gupta, and K. Roy, "Pathways to efficient neuromorphic computing with non-volatile memory technologies," *Appl. Phys. Rev.*, vol. 7, p. 021308, 2020.

[2] M. Davies, A. Wild, G. Orchard, Y. Sandamirskaya, G. A. F. Guerra, P. Joshi, P. Plank, and S. R. Risbud, "Advancing neuromorphic computing with Loihi: A survey of results and outlook," *Proc. IEEE*, vol. 109, no. 5, pp. 911–934, 2021.

[3] J. Grollier, D. Querlioz, K. Y. Camsari, K. Everschor-Sitte, S. Fukami, and M. D. Stiles, "Neuromorphic spintronics," *Nat. Electron.*, vol. 3, pp. 360–370, 2020.

[4] C. H. Marrows, J. Barker, T. A. Moore, and T. Moorsom, "Neuromorphic computing with spintronics," *npj Spintronics*, vol. 2, p. 12, 2024.

[5] K. Roy, C. Wang, S. Roy, A. Raghunathan, K. Yang, and A. Sengupta, "Spintronic neural systems," *Nat. Rev. Electr. Eng.*, vol. 1, no. 11, pp. 714–729, 2024.

[6] A. Sengupta, P. Panda, P. Wijesinghe, Y. Kim, and K. Roy, "Magnetic tunnel junction mimics stochastic cortical spiking neurons," *Sci. Rep.*, vol. 6, p. 30039, 2016.

[7] A. Sengupta, Z. A. Azim, X. Fong, and K. Roy, "Spin-orbit torque induced spike-timing dependent plasticity," *Appl. Phys. Lett.*, vol. 106, no. 9, p. 093704, 2015.

[8] S. Liu, T. P. Xiao, C. Cui, J. A. C. Incorvia, C. H. Bennett, and M. J. Marinella, "A domain wall-magnetic tunnel junction artificial synapse with notched geometry for accurate and efficient training of deep neural networks," *Appl. Phys. Lett.*, vol. 118, no. 20, p. 202405, 2021.

[9] K. Yang and A. Sengupta, "Leveraging voltage-controlled magnetic anisotropy to solve sneak path issues in crossbar arrays," *IEEE Trans. Electron Devices*, vol. 70, no. 4, pp. 2021–2027, 2023.

[10] C. M. Liyanagedera, A. Sengupta, A. Jaiswal, and K. Roy, "Stochastic spiking neural networks enabled by magnetic tunnel junctions: From nontelegraphic to telegraphic switching regimes," *Phys. Rev. Applied*, vol. 8, p. 064017, 2017.

[11] A. A. Faisal, L. P. J. Selen, and D. M. Wolpert, "Noise in the nervous system," *Nat. Rev. Neurosci.*, vol. 9, no. 4, pp. 292–303, 2008.

[12] K. Y. Camsari, R. Faria, B. M. Sutton, and S. Datta, "Stochastic p-bits for invertible logic," *Phys. Rev. X*, vol. 7, p. 031014, 2017.

[13] K. M. Song, J.-S. Jeong, B. Pan, X. Zhang, J. Xia, S. Cha, T.-E. Park, K. Kim, S. Finizio, J. Raabe *et al.*, "Skyrmion-based artificial synapses for neuromorphic computing," *Nature Electronics*, vol. 3, no. 3, pp. 148–155, 2020.

[14] L. Liu, O. J. Lee, T. J. Gudmundsen, D. C. Ralph, and R. A. Buhrman, "Current-induced switching of perpendicularly magnetized magnetic layers using spin torque from the spin hall effect," *Phys. Rev. Lett.*, vol. 109, p. 096602, 2012.

[15] T. Nozaki, T. Yamamoto, S. Miwa, M. Tsujikawa, M. Shirai, S. Yuasa, and Y. Suzuki, "Recent progress in the voltage-controlled magnetic anisotropy effect and the challenges faced in developing voltage-torque mram," *Micromachines*, vol. 10, no. 5, p. 327, 2019.

[16] A. Borst and F. E. Theunissen, "Information theory and neural coding," *Nat. Neurosci.*, vol. 2, no. 11, pp. 947–957, 1999.

[17] R. VanRullen, R. Guyonneau, and S. J. Thorpe, "Spike times make sense," *Trends Neurosci.*, vol. 28, no. 1, pp. 1–4, 2005.

[18] K. Yang and A. Sengupta, "Stochastic magnetoelectric neuron for temporal information encoding," *Appl. Phys. Lett.*, vol. 116, no. 4, p. 043701, 2020.

[19] K. Yang, D. P. G. M, and A. Sengupta, "Leveraging probabilistic switching in superparamagnets for temporal information encoding in neuromorphic systems," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 42, no. 10, pp. 3464–3468, 2023.

[20] W. A. Borders, H. Akima, S. Fukami, S. Moriya, S. Kurihara, Y. Horio, S. Sato, and H. Ohno, "Analogue spin–orbit torque device for artificial-neural-network-based associative memory operation," *Appl. Phys. Express*, vol. 10, no. 1, p. 013007, 2016.

[21] W. A. Borders, A. Z. Pervaiz, S. Fukami, K. Y. Camsari, H. Ohno, and S. Datta, "Integer factorization using stochastic magnetic tunnel junctions," *Nature*, vol. 573, pp. 390–393, 2019.

[22] G. Finocchio, M. Di Ventra, K. Y. Camsari, K. Everschor-Sitte, P. K. Amiri, and Z. Zeng, "The promise of spintronics for unconventional computing," *Journal of Magnetism and Magnetic Materials*, vol. 521, p. 167506, 2021.

[23] A. Sengupta, M. Parsa, B. Han, and K. Roy, "Probabilistic deep spiking neural systems enabled by magnetic tunnel junction," *IEEE Trans. Electron Devices*, vol. 63, no. 7, pp. 2963–2970, 2016.