Forty Tested Neuromorphic Capabilities on BrainChip Akida Hardware

Technical Paper
November 17, 2026
doi:10.5281/zenodo.24317943

Kevin D. Johnson, MBA, MAcc, MSGTD, MAT

PhD Candidate, Innovation in Global Development, Arizona State University
Strategic Research Systems LLC

***Abstract*—Forty brain-inspired capabilities were built and tested on BrainChip Akida AKD1500 neuromorphic processors, working alone and in combination with Tang Primer 20K FPGAs and host computers. The capabilities fall into four families. Memory includes fast weights, associative recall, sparse distributed memory, hippocampal indexing, multi-timescale synapses and a working memory held in the wiring of the network itself. Plasticity and learning include structural plasticity, metaplasticity, three-factor learning, predictive coding, dendritic segments, sleep, neurogenesis, neuroevolution, learning to learn, one-shot place fields and exact perturbation learning. Dynamics and computation include neuromodulation, neural fields, two-timescale control, oscillatory gating, event-driven dendrites, hyperdimensional binding, Boltzmann sampling, path integration, a dendritic XOR and an apical switch. A final family covers self-reference across chips: self-repair, a rewindable network, reflexive plasticity, introspection, exact unlearning, unihemispheric sleep, a split brain, dreaming by recombination, neural Darwinism and a causal self-map. Each capability was built and run on silicon, and every chip output used was checked against an exact integer reference of the same computation, with zero mismatches. All forty work, several within stated bounds. Each result is paired with a potential application, from verifiable data deletion and self-monitoring confidence signals to dead reckoning without GPS and fault-tolerant edge hardware.**

I. Introduction

Neuromorphic processors are usually evaluated as inference accelerators. A model is trained elsewhere, converted, mapped to the chip and run, and the figures of merit are accuracy, latency and energy per inference. BrainChip's Akida family fits the pattern well, and the AKD1000 and AKD1500 are deployed in exactly that role [1]. The author's own record of experimental demonstrations on Akida silicon, together with three co-authored IEEE NAECON 2026 papers, has used the chips at fleet scale as an orchestrated class of processor within a heterogeneous compute ontology [2], [3], [4], [5]. In such an ontology each class of processor is known for what it is, a resource with its own computational model, and is placed, sequenced and coupled with the others by orchestration [6], [7].

Brains do far more than map inputs to outputs. Synapses change on several timescales, neurons are modulated, memories are indexed and replayed, circuits are pruned and grown, and the brain monitors and repairs parts of itself. Computational neuroscience has a large literature of models for each of these functions, and most of them have been studied in software. The question taken up here is how many of them can run on commercially available neuromorphic silicon now, in combination with ordinary programmable logic and host computers, and what each would be good for.

Forty capabilities were selected, built and tested on AKD1500 hardware. The selection began from the brain functions in the literature and from what a system of exact, inspectable chips orchestrated together with FPGAs and hosts could offer each of them. Every capability is reported with its central measurement and a potential use case. All forty work. Several work within bounds stated with the result, and those bounds are reported as plainly as the successes.

Section II describes the platform. Section III gives the evaluation method. Section IV presents the forty results in four families. Section V states the limits of the work, Section VI discusses portability to the AKD1000, and Section VII concludes.

II. Platform

The work ran on ten AKD1500 processors, one per host computer, and four Sipeed Tang Primer 20K FPGA boards (Gowin GW2A-LV18), each with 128 MB of DDR3 memory and an Ethernet interface on the local network. The AKD1500 carries the Akida 1.0 neural processing architecture of the earlier AKD1000 [1]. Chips and boards were used singly, in pairs and, for the largest runs, all together: eight chips and four FPGAs operated as four pairs, each pair coupled through one FPGA.

The three kinds of device are combined within the heterogeneous compute ontology described in Section I [6], [7]. The chip contributes exact, parallel dendritic computation. The FPGA contributes state, time and the soma. The host contributes memory, logs and the library of stored networks. None of the forty capabilities belongs to one device alone. Each arises from the combination, and orchestration makes

the combination one system: it decides which device carries each step, in what order the steps run, and what state passes between them on every cycle. Without the ontology the chip, the FPGA and the host remain three separate tools. Within it they act as one brain-like system whose parts each do what they do best.

Five properties of the combined system recur throughout the results. The chips are exact: a computation repeated on the same inputs gives bit-identical outputs, across chips as well as across runs. Synapses, thresholds and the wiring between neuron groups can be read back while the chip runs. The same quantities can be rewritten between two inferences, at microsecond to millisecond cost. A library of complete stored networks can be held outside the chip, any one of them active on the next call. Finally, state that a brain keeps inside its neurons, such as membrane potentials, traces, history and logs, is kept on the FPGAs and hosts, where it can be inspected, copied and restored.

Several capabilities use a two-compartment neuron modelled on the cortical pyramidal cell [8], [9]. Basal dendrites run on one chip and apical dendrites on a second chip, and the somas run on an FPGA that combines the two inputs and produces single events or bursts. In the cortical pyramidal cell, bursts signal coincidence between bottom-up input and top-down context. The neuron is exercised most directly in one-shot place fields (Capability 27) and the apical switch (Capability 29).

III. Evaluation Method

Each capability was built around one clear measurement. The measurement and its comparison condition were fixed before the run. Where the outcome could vary with initial conditions, runs were repeated over several random seeds and reported as counts of seeds, confidence intervals or p-values, as stated with each result.

Every chip output entering a result was compared against an exact integer reference of the same computation. A result with an unexplained mismatch was not accepted. Across all forty capabilities the mismatch count was zero. Soma computations on the FPGAs were checked in the same way against a bit-exact integer model. Quantities computed on the host, such as learning rules, statistics and controllers, are part of each design and are stated with the capability that uses them.

The exactness is more than a quality check. Several capabilities depend on it directly. A network can be rewound and replayed only if replay reproduces it, an experience can be removed exactly only if a run without it is reproducible, and a perturbation can be measured exactly only if the unperturbed run is.

IV. Results

The forty capabilities are grouped into four families of ten. Each table gives the central measurement and a potential use. The text following each table places the capabilities against the literature they draw on.

A. Memory and Representation

Fast weights, synapses that change quickly enough to hold the recent past, were proposed as a memory between activity and slow learning [10], [11]. In Capability 1 the output of one layer is written as the weights of another and used on the next call. Capability 2 is the content-addressable memory of Hopfield [12] in the form of a single feedforward recall. Capability 3 treats a timing race between two processors as an entropy source and holds it to the continuous health tests of NIST SP 800-90B [13]. The source is health-monitored but not certified.

Neuromodulation lets one anatomical circuit express several functions [14]. Capability 4 realises it as a threshold change alone, with the weights untouched. Structural plasticity, the growth and loss of synapses under a resource budget [15], appears in Capability 5. Capability 6 is engineering more than neuroscience: a processor whose memory has been corrupted is detected, its work moved to a spare and the network's wiring updated, all while the network keeps running.

Capability 7 holds hundreds of cortical columns as separate weight banks and gives each its turn on the chip. Capability 8 is a continuous attractor of the Amari type [16], with its lateral connections on the chip and its state on the FPGA. Capability 9 uses the chip's built-in learning layer as a continuously running novelty detector whose memory is read out without disturbing it. Capability 10 joins a fast inference loop to a slow homeostatic controller, the arrangement by which neurons hold their own activity near a set point [17].

B. Plasticity and Learning

Three-factor rules combine pre- and postsynaptic activity with a third, delayed signal such as reward or a neuromodulator [18]. Capability 11 shows the eligibility trace earning its place: the rule with traces beats the rule without them at every reward delay tested. Capability 12 implements the predictive coding scheme of Rao and Ballard [19]. A layer predicts its input through the exact transpose of its own weights, and both directions learn in place.

Capability 13 follows the account of active dendritic segments by Hawkins and Ahmad [20]: context selects a subset of neurons, so tasks that conflict use different parts of the network. Capability 14 is replay during sleep [21], run through the chip at about 8,000 trials per second. Capability 15 adds new classes in spare neurons while existing neurons stay bit-identical, a strict form of the stability that adult neurogenesis is thought to provide for new memories [22].

Communication through coherence holds that rhythms gate which signals pass between areas [23]. In Capability 16 the gate is a rhythmic change of thresholds, and information passes only at the matching phase. Capability 17 evolves networks in the manner of neuroevolution [24], with every candidate evaluated on the chip rather than in a software model. Capability 18 runs a dendritic layer as an event program whose cost scales with the number of input events. The result supports cheaper processing when inputs are very sparse, below about 12% activity in the measured layer.

TABLE I
CAPABILITIES 1 TO 10

No.	Capability	Result on silicon	Potential use
1	Fast weights	One layer's output becomes another layer's weights on the very next call, in a 1.66 ms loop.	Short-term context memory for edge agents, without retraining.
2	Associative memory	1,024 patterns stored as weights across 4 banks, recalled perfectly from queries with up to 30% noise, in one call.	On-device lookup of signatures, fingerprints or anomalies from partial cues.
3	Chip-internal randomness	A race between two on-chip processors routed to the same output, characterised as a noise source; NIST health tests pass over 30 minutes.	Hardware noise for stochastic inference and sampling.
4	Neuromodulation	The same weights compute two different mappings, switched by live threshold writes in 39 us.	Task or attention switching on one deployed model, without reloading.
5	Structural plasticity	Synapses pruned and grown live under a fixed budget; selectivity rose from 0.48 to 0.95.	Sparse networks that rewire themselves to the data they see in the field.
6	Self-repair	A corrupted processor is found by scrubbing its memory, migrated to a spare and rerouted in 5.7 ms.	Fault-tolerant edge and space hardware that heals without a reboot.
7	Virtual cortex	64 to 512 columns, each its own weight bank, time-shared on one chip at 1,350 column-steps per second.	Many small models on one chip, per user, per sensor or per context.
8	Neural fields	A ring attractor with its connections on the chip and its state on the FPGA; a stable bump, with diffusion and survival under noise measured.	Persistent tracking state, such as a target bearing, held by the network itself.
9	Native learning, repurposed	Akida's own on-chip learning run continuously and read back without disturbing it; novelty detection AUC 1.0.	Always-on anomaly detection with an inspectable memory.
10	Two timescales	A 6.7 kHz chip loop steered by a 100 Hz FPGA loop through threshold writes; 128 of 128 neurons held at target.	Self-calibrating sensors with a slow homeostatic controller.

TABLE II
CAPABILITIES 11 TO 20

No.	Capability	Result on silicon	Potential use
11	Three-factor learning	Delayed reward turned into weight writes through eligibility traces; beats the no-trace rule at every delay, over 6 seeds.	On-device reinforcement learning with delayed feedback.
12	Predictive coding	A prediction and its exact transpose in one layer, both learned in place; 16 of 16 bars learned.	Unsupervised feature learning and reconstruction-error anomaly detection.
13	Dendritic segments	Context segments gate neurons during continual learning; +0.24 to +0.29 accuracy over 10 to 20 conflicting tasks.	Learning task after task without forgetting the old ones.
14	Sleep	Logged experience replayed through the chip at about 8,000 trials per second; retention rose from 0.917 to 0.997.	Consolidation during idle time for devices that learn in the field.
15	Neurogenesis	New classes born into silent spare neurons while the old neurons stay bit-identical.	Adding classes after deployment without touching what already works.
16	Oscillations	Rhythmic threshold rewrites gate communication by phase: 0.43 bits get through at the right phase, zero out of phase.	Routing information between modules by timing.
17	Neuroevolution	Every genome evaluated on the chip itself, straight from the bank store; accuracy rose from 22.7% to 74.2%.	Gradient-free model search on the target hardware.
18	Event-driven dendrites	The dendrite layer run as a sparse-input event program, costing 3,688 cycles plus 35.6 per event.	Cheaper processing of very sparse sensor streams.
19	Metaplasticity	A slow controller adjusts thresholds to stop runaway Hebbian learning; selectivity 0.91 to 0.99.	Stable unsupervised learning left running unattended.
20	Hippocampal indexing	Episodes stored as pointers and reinstated bit for bit from partial cues.	Episodic recall for agents: "what happened last time I saw this".

Metaplasticity, the plasticity of plasticity [25], keeps Hebbian learning from running away in Capability 19. Capability 20 stores episodes as the hippocampal indexing theory proposes [26], as pointers to cortical patterns, and reinstates the full pattern bit for bit from a partial cue.

C. Dynamics and Computation

Kanerva's sparse distributed memory stores data at hard locations near an address and reads it from every location within a Hamming radius [27]. In Capability 21 that radius is a neuron's threshold, and recall agrees with the theory's predictions. Benna and Fusi showed that synapses built as cascades of coupled variables retain memories with power-law decay and far longer than single-variable synapses [28]. Capability 22 builds such a cascade from six levels of stored weights and measures a retention slope of -0.503, close to the -0.5 the theory predicts.

Capability 23 performs the binding operation of vector-symbolic architectures [29], [30], circular convolution, as multiplication by a circulant weight bank that the chip itself writes. Exact unbinding doubles the number of role-filler pairs that can be superposed and recovered. Capability 24 samples a restricted Boltzmann machine [31] on the chip, using an exact transpose for the reverse pass. The error between the sampled and exact distributions reaches the floor set by finite sample size.

Head-direction and grid cells integrate self-motion through a continuous attractor whose connectivity is shifted by velocity [32]. Capability 25 selects a stored connection bank for each velocity bin, holding heading error to 0.41 degrees. Gidon and colleagues found dendritic action potentials in human cortex whose amplitude falls for stronger input, which lets a single dendrite compute XOR [33]. Capability 26 reproduces the non-monotonic response with two thresholds and switches the same dendrite to OR or AND in place.

Behavioural timescale plasticity forms a hippocampal place field in a single trial, triggered by a dendritic plateau [34]. Capability 27 uses the two-compartment neuron: apical drive from a second chip produces plateaus on the FPGA, and each plateau triggers a rewrite of that neuron's weights. All 16 fields formed in one lap. Capability 28 tunes a plasticity rule by an outer evolutionary loop [35]. Every inner episode is exactly repeatable, so the outer loop sees no hardware noise.

The dendritic integration theory of Aru, Suzuki and Larkum links conscious processing to coupling between the apical and somatic compartments of layer-5 pyramidal neurons, a coupling that anaesthetics disrupt [36]. Capability 29 opens and closes that coupling. Feedforward activity survives, while bursts and the learning they drive disappear, as the theory predicts. Capability 30 runs a network forward, restores it to an earlier state, changes one thing and runs it forward again. The two branches are bit-identical except for the consequences of the change, which turns the lesion and perturbation logic of neuroscience [37] into exact measurement.

D. Self-Reference and Multiple Chips

The final family uses the system's ability to observe and act on its own substrate. A self-referential weight matrix, a network able to inspect and modify its own weights, was proposed as a route to learning without an external optimiser [38]. In Capability 31 the network's outputs are themselves the words that rewrite its synapses and its plasticity, and what the network writes matches, word for word, what an external host would have written. Capability 32 gives a second chip a sense organ pointed at the first chip's synapses. From the synapses alone it predicts the first chip's competence better than the first chip's own confidence does. The task was constructed to contain confident false memories, so the comparison measures exactly the case where self-reported confidence fails. Metacognition, knowing what one knows, is otherwise inferred from behaviour [39].

Machine unlearning seeks to remove the influence of particular training data from a model [40], [41]. Capability 33 does it exactly. A network restored to a snapshot and replayed without one experience is identical to the network that never saw it, and the difference from the original run measures the full footprint of that experience. Capability 34 borrows unihemispheric sleep from dolphins and birds [42]: one chip serves while the other replays and consolidates, then they swap, and retention improves with no downtime. Capability 35 is a split brain [43] in which two chip hemispheres share information through a narrow FPGA callosum. The same starting state is run with the callosum intact and cut, so the effect of the cut is measured with a perfect control.

Weight and node perturbation estimate credit from a global scalar by measuring how a small change alters performance [44], [45]. On a noisy substrate each measurement is a sample. In Capability 36 each measurement is the true effect, and exact perturbation beats a noise-injected copy on every seed. Capability 37 tests the idea that dreams recombine fragments of experience [46]. New networks are recombined from stored ones and kept if they help, and they give a head start on situations that blend familiar ones. They give none on material the stored networks do not cover, and the result is reported with that bound.

Edelman's neural Darwinism holds that a repertoire of variant circuits is selected by value during behaviour, with reentry between repertoires [47]. Capability 38 keeps such a repertoire live and selects one circuit per call, and reentry between two chips raises reward. Working memory is usually attributed to persistent activity or short-term synaptic change [48]. Capability 39 stores items in a third medium, the routing of neuron groups to outputs, and recalls them exactly. Capability 40 combines the rewindable network with introspection. The system lesions its own parts in a sandbox copy, measures what each contributes and protects the parts that matter during further learning, in the spirit of importance-weighted consolidation [49]. With a protection budget that covers the population at risk, it matches an oracle that knows the true importance.

TABLE III
CAPABILITIES 21 TO 30

No.	Capability	Result on silicon	Potential use
21	Sparse distributed memory	Each neuron's threshold serves as a Hamming radius over Kanerva's hard locations; recall matches his theory.	Robust approximate-match memory and nearest-neighbour search on device.
22	Multi-timescale synapses	A coupled six-level ladder of weight banks gives power-law retention, slope -0.503, as Benna and Fusi predict.	Streaming learners that keep recent detail and old knowledge.
23	Hyperdimensional binding	Circulant weight banks written by the chip bind role-filler pairs; exact unbinding doubles capacity.	Structured, symbolic reasoning on neuromorphic hardware.
24	Boltzmann sampler	A restricted Boltzmann machine sampled on the chip with an exact transpose; error at the finite-sample floor.	Probabilistic inference, generative sampling and optimisation at the edge.
25	Path integration	Velocity selects a connection bank each step; 0.41 degree heading error.	Dead reckoning for drones and robots without GPS.
26	Dendritic XOR	Two thresholds make one non-monotonic dendrite: XOR at 100%, switched live to OR or AND.	More computation per neuron; reconfigurable logic inside a network.
27	One-shot place fields	Apical drive from a second chip makes plateaus on the FPGA, which trigger in-place weight rewrites; 16 of 16 place fields formed in one lap.	One-shot learning of places or events.
28	Learning to learn	Evolution tunes the plasticity rule, with every on-chip episode exactly repeatable; the tuned rule learned 126 of 256 against 30.	On-device learning rules tuned automatically for each deployment.
29	The apical switch	Cutting apical coupling, a silicon "anaesthesia", keeps feedforward activity but removes bursts and learning.	Switching learning safely on and off; a test bed for theories of awareness.
30	The rewindable brain	Snapshot, restore, change one thing, run forward: branches are bit-identical except for the change.	Debugging and auditing learning systems with exact what-if analysis.

TABLE IV
CAPABILITIES 31 TO 40

No.	Capability	Result on silicon	Potential use
31	Reflexive plasticity	The network's own outputs are the configuration packets that rewrite its synapses and its plasticity; self-written equals host-written, word for word.	Self-modifying controllers with an exact log of every self-change.
32	Introspection	A second chip sees only the first chip's synapses and predicts its competence better than the first chip's own confidence (p = 8e-9, 6 seeds).	Self-monitoring AI that catches confident false memories.
33	Exact unlearning	Restoring and replaying with one experience left out gives exactly the network that never saw it, and that experience's footprint is measured.	Verifiable deletion of data from a learning device.
34	Unihemispheric sleep	Two chips take turns: one serves while the other replays and consolidates. Retention improves, with no downtime.	Systems that learn around the clock and never go offline.
35	Split brain	Two chips joined by a narrow FPGA "corpus callosum", run intact and cut from the identical starting state.	Designing distributed AI over limited links.
36	Exact perturbation learning	Perturb, measure the exact effect, restore; beats a noisy copy on 8 of 8 seeds.	Gradient-free learning where backpropagation is impossible.
37	Dreaming	New networks recombined from stored ones and kept if they help: a +0.26 to +0.32 head start on new situations that blend familiar ones, none on truly novel material.	Preparing for new mixes of familiar situations during idle time.
38	Neural Darwinism	A value-driven selector picks one variant circuit per call from a live repertoire; reentry across two chips raises reward.	Adaptive ensembles that keep many solutions ready.
39	The connectome as working memory	Items held as live rewirings of which neurons feed which outputs, and recalled exactly.	Fast short-term memory without changing any weights.
40	A causal self-map	The network lesions its own parts in a sandbox copy and protects what matters; it matches an oracle at the right budget.	Continual learners that know which parts carry their existing skills.

V. Limits

All forty capabilities are minimal versions: small layers, short runs and one measurement each. They establish that each capability runs and behaves as intended on the hardware. They do not establish performance at application scale. Several results hold within bounds stated in the tables. Event-driven dendrites (Capability 18) pay off only below about 12% input activity. Dreaming (Capability 37) helps only on blends of familiar material. The causal self-map (Capability 40) matches the oracle when its budget covers the parts at risk.

The control loops that join chips, FPGAs and hosts step in milliseconds, not microseconds, so capabilities that need fast dynamics run slower than biological time in some cases and faster in others. Power was not measured. No comparison with other neuromorphic platforms was attempted. The interest of each capability lies in what it makes possible on this hardware, not in a ranking.

Some capabilities compute part of their logic on the host or the FPGA, such as a learning rule, a statistic or a controller. The division is part of each design. The work treats the chips, FPGAs and hosts as one orchestrated system, and the results are results of that system.

Several aspects of these tests, however, have already been carried to fleet scale in the author's demonstration record [2]. Ten AKD1500s ran 20,000 complete fly connectomes at once, with every chip readout matching the integer reference across half a million comparisons (Demonstration 81), which is the many-networks-per-chip principle of the virtual cortex (Capability 7) and the exactness of Section III at scale. Four Tang Primer 20K FPGAs and ten AKD1500s swapped a hundred models across ten chips in 0.155 s and learned an emitter class absent from all of them after deployment, beating a frozen control on 80 of 80 runs (Demonstration 77), joining bank selection, neurogenesis and the chip-FPGA combination (Capabilities 7, 15 and 17). Ten nodes each learned their own normal on chip and flagged a driven node at 11.4 sigma while nine neighbours held still (Demonstration 72), the native learning of Capability 9 across a fleet. Ten AKD1000s carried a full functional-brain design, with structural plasticity, neuromodulation, sleep-time pruning, consolidation and a predict-sense-update loop (Demonstrations 49, 50 and 51), which spans Capabilities 4, 5, 12, 14 and 19. Twelve chips learned on chip to model and drive a race car, one of them learning from reward how hard to push (Demonstrations 60 and 63), as in three-factor learning (Capability 11). A council of chips across three sites kept reaching consensus while a link was jammed, a node was lost and the mesh was partitioned (Demonstration 54), which is self-repair and coordination over limited links at the scale of a network (Capabilities 6 and 35). An FPGA neural fabric beside an AKD1000 combined neuromodulated learning, homeostatic thresholds, local dendritic integration and stochastic neurons in one loop (Demonstration 68), as in Capabilities 3, 4 and 10.

VI. Portability to the AKD1000

All forty capabilities were run on the AKD1500. The AKD1000 carries the same Akida 1.0 neural processing architecture, and in earlier work the two parts ran an identical set of models with identical coverage. The forty capabilities could run on the AKD1000 as well.

VII. Conclusion

Forty brain-inspired capabilities, drawn from decades of computational neuroscience and machine learning, run on Akida neuromorphic silicon today, with every chip output exact. Several of them, among them exact unlearning, introspection of one's own synapses, exact perturbation, a rewindable brain and a causal self-map, depend on a system that is at once physical, exact, inspectable and rewritable. Orchestrated with FPGAs and hosts within a heterogeneous compute ontology, the chips provide that combination.

The results bear on a long-standing view of neuromorphic computing. Van der Made argued that a functional artificial brain would come from hardware neurons organised and adapted the way biological ones are [50]. A large part of that program is often placed several chip generations away. The forty capabilities reported here suggest that much of it is available now, through orchestration of current chips with programmable logic, and that the near-term work is to take these minimal versions to application scale.

References

[1] BrainChip, "Akida neuromorphic processors: AKD1000 and AKD1500," product information. Available: https://brainchip.com

[2] K. D. Johnson, "Demonstration Record," 82 demonstrations, December 2025 to September 2026. Demonstration N is at https://kevindjohnson.org/#demoN. Full demonstration record available at: https://kevindjohnson.org/#demos

[3] K. D. Johnson and S. D. Harbour, "Neuromorphic Orchestration for Efficient LLM Infrastructure via Utility-Optimized Decision-Making," in *Proc. 2026 IEEE National Aerospace and Electronics Conference (NAECON)*, 2026. doi:10.1109/NAECON70028.2026.11675008.

[4] K. D. Johnson and S. D. Harbour, "Utility-Optimized Neuromorphic Orchestration for Real-Time ISR and Multi-Sensor Fusion," in *Proc. 2026 IEEE National Aerospace and Electronics Conference (NAECON)*, 2026. doi:10.1109/NAECON70028.2026.11675010.

[5] S. D. Harbour, K. D. Johnson, G. Silaparasetti, E. M. Ehmann, A. Bermudez, J. Scheitlin, P. Toutam, A. H. Bagheri and K. Wilson, "From Compute Scaling to Decision Optimization: Neuromorphic Architectures for Sustainable AI Infrastructure," in *Proc. 2026 IEEE National Aerospace and Electronics Conference (NAECON)*, 2026. doi:10.1109/NAECON70028.2026.11674860.

[6] K. D. Johnson, "Old Black Water Keep On Rollin': The Neuromorphic in a Heterogeneous Compute Ontology," 22 June 2026. Available: https://kevindjohnson.org/old-black-water-keep-on-rollin-the-neuromorphic-in-a-heterogeneous-compute-ontology/

[7] K. D. Johnson, "High Performance Quantum-Centric Supercomputing: A Working Implementation of Heterogeneous Orchestration across QPU, NPU, GPU, CPU, and Other Tiers," Technical Paper, Mar. 2026. doi:10.5281/zenodo.19888562.

[8] M. Larkum, "A Cellular Mechanism for Cortical Associations: An Organizing Principle for the Cerebral Cortex," *Trends in Neurosciences*, vol. 36, no. 3, 2013. doi:10.1016/j.tins.2012.11.006.

[9] A. Payeur, J. Guerguiev, F. Zenke, B. A. Richards and R. Naud, "Burst-Dependent Synaptic Plasticity Can Coordinate Learning in Hierarchical Circuits," *Nature Neuroscience*, vol. 24, 2021. doi:10.1038/s41593-021-00857-x.

[10] J. Schmidhuber, "Learning to Control Fast-Weight Memories: An Alternative to Dynamic Recurrent Networks," *Neural Computation*, vol. 4, no. 1, 1992. doi:10.1162/neco.1992.4.1.131.

[11] J. Ba, G. E. Hinton, V. Mnih, J. Z. Leibo and C. Ionescu, "Using Fast Weights to Attend to the Recent Past," in *Advances in Neural Information Processing Systems 29*, 2016. arXiv:1610.06258.

[12] J. J. Hopfield, "Neural Networks and Physical Systems with Emergent Collective Computational Abilities," *Proc. National Academy of Sciences*, vol. 79, no. 8, 1982. doi:10.1073/pnas.79.8.2554.

[13] M. S. Turan, E. Barker, J. Kelsey, K. A. McKay, M. L. Baish and M. Boyle, "Recommendation for the Entropy Sources Used for Random Bit Generation," NIST Special Publication 800-90B, 2018. doi:10.6028/NIST.SP.800-90B.

[14] E. Marder, "Neuromodulation of Neuronal Circuits: Back to the Future," *Neuron*, vol. 76, no. 1, 2012. doi:10.1016/j.neuron.2012.09.010.

[15] A. Holtmaat and K. Svoboda, "Experience-Dependent Structural Synaptic Plasticity in the Mammalian Brain," *Nature Reviews Neuroscience*, vol. 10, 2009. doi:10.1038/nrn2699.

[16] S. Amari, "Dynamics of Pattern Formation in Lateral-Inhibition Type Neural Fields," *Biological Cybernetics*, vol. 27, 1977. doi:10.1007/BF00337259.

[17] G. G. Turrigiano, "The Self-Tuning Neuron: Synaptic Scaling of Excitatory Synapses," *Cell*, vol. 135, no. 3, 2008. doi:10.1016/j.cell.2008.10.008.

[18] N. Fremaux and W. Gerstner, "Neuromodulated Spike-Timing-Dependent Plasticity, and Theory of Three-Factor Learning Rules," *Frontiers in Neural Circuits*, vol. 9, 2016. doi:10.3389/fncir.2015.00085.

[19] R. P. N. Rao and D. H. Ballard, "Predictive Coding in the Visual Cortex: A Functional Interpretation of Some Extra-Classical Receptive-Field Effects," *Nature Neuroscience*, vol. 2, no. 1, 1999. doi:10.1038/4580.

[20] J. Hawkins and S. Ahmad, "Why Neurons Have Thousands of Synapses, a Theory of Sequence Memory in Neocortex," *Frontiers in Neural Circuits*, vol. 10, 2016. doi:10.3389/fncir.2016.00023.

[21] M. A. Wilson and B. L. McNaughton, "Reactivation of Hippocampal Ensemble Memories During Sleep," *Science*, vol. 265, 1994. doi:10.1126/science.8036517.

[22] J. B. Aimone, W. Deng and F. H. Gage, "Resolving New Memories: A Critical Look at the Dentate Gyrus, Adult Neurogenesis, and Pattern Separation," *Neuron*, vol. 70, no. 4, 2011. doi:10.1016/j.neuron.2011.05.010.

[23] P. Fries, "A Mechanism for Cognitive Dynamics: Neuronal Communication Through Neuronal Coherence," *Trends in Cognitive Sciences*, vol. 9, no. 10, 2005. doi:10.1016/j.tics.2005.08.011.

[24] K. O. Stanley and R. Miikkulainen, "Evolving Neural Networks Through Augmenting Topologies," *Evolutionary Computation*, vol. 10, no. 2, 2002. doi:10.1162/106365602320169811.

[25] W. C. Abraham and M. F. Bear, "Metaplasticity: The Plasticity of Synaptic Plasticity," *Trends in Neurosciences*, vol. 19, no. 4, 1996. doi:10.1016/S0166-2236(96)80018-X.

[26] T. J. Teyler and P. DiScenna, "The Hippocampal Memory Indexing Theory," *Behavioral Neuroscience*, vol. 100, no. 2, 1986. doi:10.1037/0735-7044.100.2.147.

[27] P. Kanerva, *Sparse Distributed Memory*. Cambridge, MA: MIT Press, 1988.

[28] M. K. Benna and S. Fusi, "Computational Principles of Synaptic Memory Consolidation," *Nature Neuroscience*, vol. 19, 2016. doi:10.1038/nn.4401.

[29] T. A. Plate, "Holographic Reduced Representations," *IEEE Transactions on Neural Networks*, vol. 6, no. 3, 1995. doi:10.1109/72.377968.

[30] P. Kanerva, "Hyperdimensional Computing: An Introduction to Computing in Distributed Representation with High-Dimensional Random Vectors," *Cognitive Computation*, vol. 1, no. 2, 2009. doi:10.1007/s12559-009-9009-8.

[31] D. H. Ackley, G. E. Hinton and T. J. Sejnowski, "A Learning Algorithm for Boltzmann Machines," *Cognitive Science*, vol. 9, no. 1, 1985. doi:10.1207/s15516709cog0901_7.

[32] B. L. McNaughton, F. P. Battaglia, O. Jensen, E. I. Moser and M.-B. Moser, "Path Integration and the Neural Basis of the 'Cognitive Map'," *Nature Reviews Neuroscience*, vol. 7, 2006. doi:10.1038/nrn1932.

[33] A. Gidon, T. A. Zolnik, P. Fidzinski, F. Bolduan, A. Papoutsi, P. Poirazi, M. Holtkamp, I. Vida and M. E. Larkum, "Dendritic Action Potentials and Computation in Human Layer 2/3 Cortical Neurons," *Science*, vol. 367, 2020. doi:10.1126/science.aax6239.

[34] K. C. Bittner, A. D. Milstein, C. Grienberger, S. Romani and J. C. Magee, "Behavioral Time Scale Synaptic Plasticity Underlies CA1 Place Fields," *Science*, vol. 357, 2017. doi:10.1126/science.aan3846.

[35] S. Thrun and L. Pratt, Eds., *Learning to Learn*. Boston, MA: Springer, 1998. doi:10.1007/978-1-4615-5529-2.

[36] J. Aru, M. Suzuki and M. E. Larkum, "Cellular Mechanisms of Conscious Processing," *Trends in Cognitive Sciences*, vol. 24, no. 10, 2020. doi:10.1016/j.tics.2020.07.006.

[37] E. Jonas and K. P. Kording, "Could a Neuroscientist Understand a Microprocessor?," *PLoS Computational Biology*, vol. 13, no. 1, 2017. doi:10.1371/journal.pcbi.1005268.

[38] J. Schmidhuber, "A Self-Referential Weight Matrix," in *Proc. International Conference on Artificial Neural Networks (ICANN)*, 1993. doi:10.1007/978-1-4471-2063-6_107.

[39] S. M. Fleming and R. J. Dolan, "The Neural Basis of Metacognitive Ability," *Philosophical Transactions of the Royal Society B*, vol. 367, 2012. doi:10.1098/rstb.2011.0417.

[40] Y. Cao and J. Yang, "Towards Making Systems Forget with Machine Unlearning," in *Proc. IEEE Symposium on Security and Privacy*, 2015. doi:10.1109/SP.2015.35.

[41] L. Bourtoule, V. Chandrasekaran, C. A. Choquette-Choo, H. Jia, A. Travers, B. Zhang, D. Lie and N. Papernot, "Machine Unlearning," in *Proc. IEEE Symposium on Security and Privacy*, 2021. doi:10.1109/SP40001.2021.00019.

[42] N. C. Rattenborg, C. J. Amlaner and S. L. Lima, "Behavioral, Neurophysiological and Evolutionary Perspectives on Unihemispheric Sleep," *Neuroscience and Biobehavioral Reviews*, vol. 24, no. 8, 2000. doi:10.1016/S0149-7634(00)00039-7.

[43] M. S. Gazzaniga, "Forty-Five Years of Split-Brain Research and Still Going Strong," *Nature Reviews Neuroscience*, vol. 6, 2005. doi:10.1038/nrn1723.

[44] M. Jabri and B. Flower, "Weight Perturbation: An Optimal Architecture and Learning Technique for Analog VLSI Feedforward and Recurrent Multilayer Networks," *IEEE Transactions on Neural Networks*, vol. 3, no. 1, 1992. doi:10.1109/72.105429.

[45] R. J. Williams, "Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning," *Machine Learning*, vol. 8, 1992. doi:10.1007/BF00992696.

[46] E. Hoel, "The Overfitted Brain: Dreams Evolved to Assist Generalization," *Patterns*, vol. 2, no. 5, 2021. doi:10.1016/j.patter.2021.100244.

[47] G. M. Edelman, *Neural Darwinism: The Theory of Neuronal Group Selection*. New York: Basic Books, 1987.

[48] G. Mongillo, O. Barak and M. Tsodyks, "Synaptic Theory of Working Memory," *Science*, vol. 319, 2008. doi:10.1126/science.1150769.

[49] J. Kirkpatrick et al., "Overcoming Catastrophic Forgetting in Neural Networks," *Proc. National Academy of Sciences*, vol. 114, no. 13, 2017. doi:10.1073/pnas.1611835114.

[50] P. van der Made, *Higher Intelligence: How to Create a Functional Artificial Brain*. Vivid Publishing, 2013.