Lorenzo Pavesi

When Microring Resonators Behave Like Neurons

Emergent photonic dynamics could help move AI hardware beyond the energy limits of conventional digital computing.

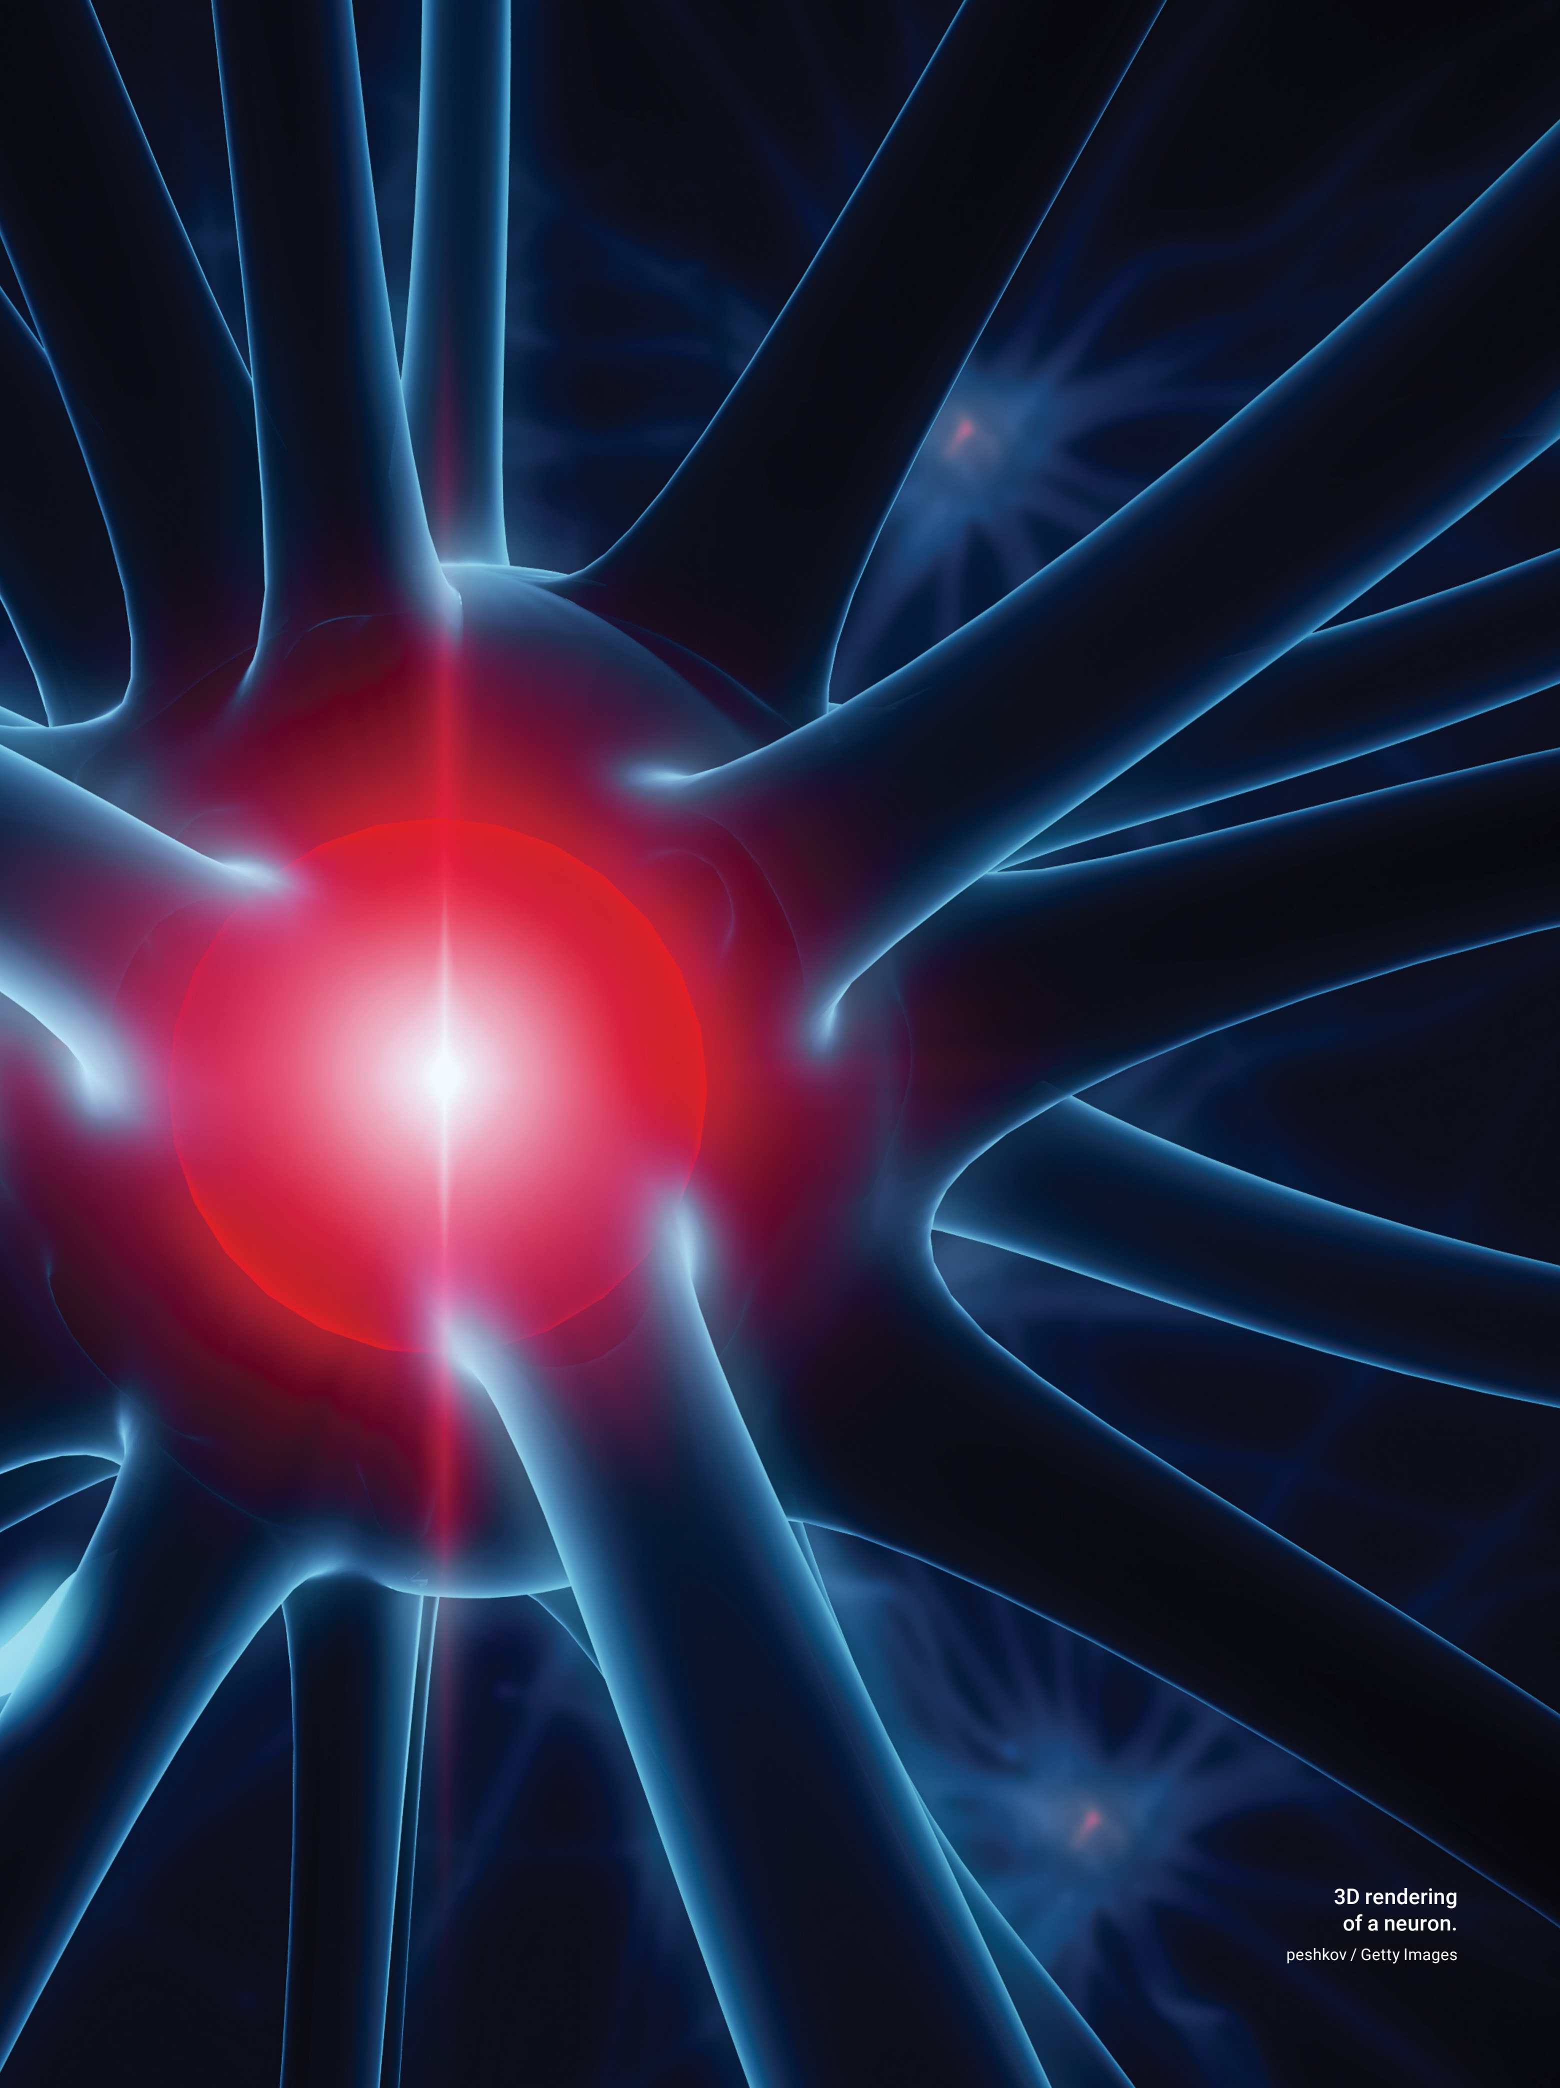

3D rendering of a neuron.
peshkov / Getty Images

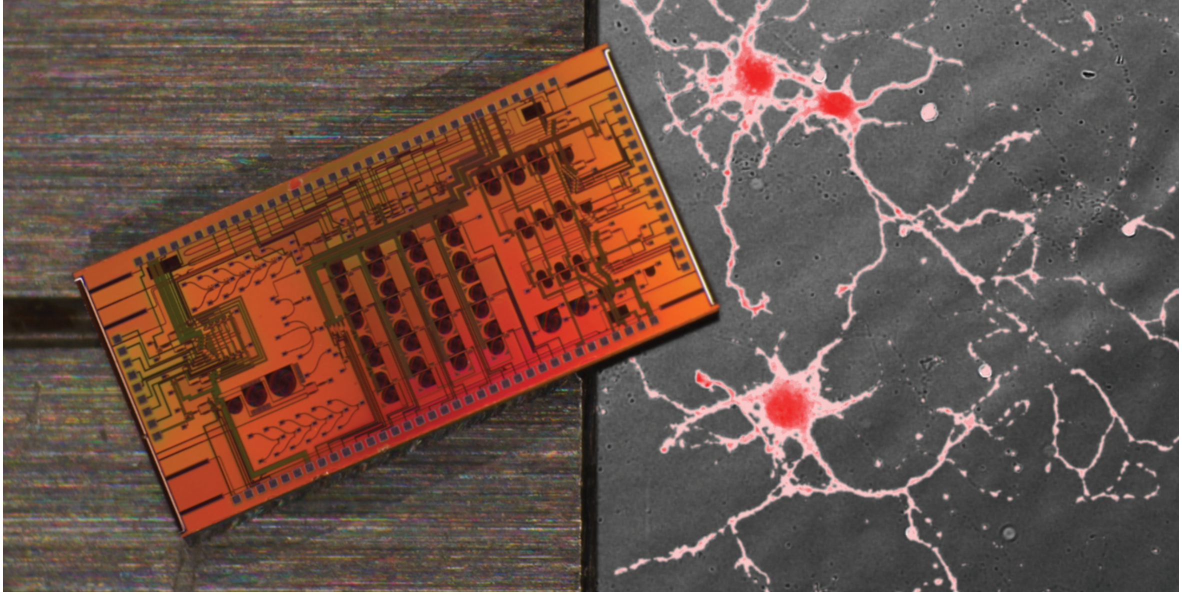

Left: A photonic integrated circuit fabricated with multiple neural-network configuration. Right: Neurons in a primary culture, showing their characteristic interconnections.
Courtesy of D. Bazzanella and C. Zaccaria

Artificial intelligence is advancing rapidly, but the hardware that supports it faces a growing energy and sustainability barrier. A central culprit is the "memory wall": the cost of moving data between memory and processors in traditional von Neumann architectures. Biological brains, by contrast, process information with unmatched power efficiency relying on event-driven, parallel and multiscale physical mechanisms.

Neuromorphic photonics, particularly platforms based on integrated silicon microring resonators, offers one route toward hardware that does not merely simulate neurons in software but uses the physics of light, carriers and heat to reproduce key neuronal functions. In such devices, the competition between ultra-fast electronic nonlinearities and slower thermal relaxation paths can generate spiking, temporal integration and refractory behavior in a sub-millimeter optical cavity.

Why now? AI, energy and the memory wall

The proliferation of deep learning frameworks—most notably massive generative models—has reshaped computing. Yet this revolution is underpinned by an unsustainable power footprint. Modern deep learning architectures require hundreds of megawatt-hours of electrical energy for a single training cycle, reflecting an exponential growth trajectory in computing demand that far outpaces traditional silicon scaling laws. This energy crisis is a fundamental architectural limitation tied to the structural foundations of contemporary digital hardware.

Conventional digital computing has relied almost exclusively on the von Neumann model, in which processing units are physically isolated from memory storage banks. Consequently, mathematical operations require the continuous, high-frequency transport of data across narrow metal interconnects, causing significant Joule heating and parasitic power dissipation. As models scale to hundreds of billions of parameters, the energy expended simply shifting weight metrics and activations surpasses the energy used on the calculation itself. This systemic bottleneck, referred to as the "memory wall," introduces severe thermal dissipation constraints and caps the maximum effective throughput of electronic microprocessors. Major semiconductor companies are therefore investing heavily in optical interconnects to speed up data movement and reduce its power cost.

Furthermore, contemporary digital networks lack true generalization capabilities. They remain rigidly deterministic or rely on brute-force statistical correlations, requiring millions of labeled training examples to achieve proficiency in narrow domains. In contrast, biological neural networks perform complex tasks such as real-time sensory-motor integration, abstract reasoning and fluid learning while consuming approximately 20 watts of power. The human brain completely bypasses the memory wall by executing computation and memory storage within the same physical substrate: the synaptic connection. To overcome the structural limits of digital systems, we must look beyond standard binary abstraction and investigate alternative computing paradigms that harness the rich, intrinsic behavior of neuronal systems.

Biological inspiration

Biological neurons differ from digital hardware in significant ways. Digital systems rely on synchronized, high-frequency clocks that force every transistor across a wafer to switch uniformly, irrespective of data

1047-6938/26/09/26/9-$15.00 ©Optica

Integrated silicon microring resonators offer one route toward hardware that does not merely simulate neurons in software but uses the physics of light, carriers and heat to reproduce key neuronal functions.

relevance. Biological neural architectures, by contrast, are asynchronous and event-driven. They communicate through transient, discrete electrical signals, or action potentials. When a pathway receives no stimulus, it remains dormant, keeping power consumption low.

We can isolate seven core operational characteristics where biological systems fundamentally diverge from traditional digital hardware:

Event-based processing: Information is represented via sparse, discrete spikes, allowing biological systems to allocate energy only when meaningful environmental or systemic events occur.

Power-efficient, low-frequency operation: Biological neurons typically operate between 1 Hz and 100 Hz, far below the gigahertz switching rates of electronic logic. Their computational throughput comes from massive, parallel spatial interconnectedness rather than raw temporal speed.

Multisensorial input convergence: Biological nodes integrate audio, visual, somatosensory and other signals within shared cortical processing regions.

3D and folded topologies: Biological brains use an intricate, highly cross-linked 3D configuration. This minimizes interconnection distances and maximizes packaging density, avoiding the routing congestion common in flat 2D silicon layouts.

Multilayer heterogeneous architecture: Neural computing relies on complex interactions among heterogeneous neuron types and supportive glia cell networks, which regulate synaptic efficiency and maintain system balance over long periods.

Multiplicity of signaling carriers: Information transfer combines rapid electrical action potentials along axonal membranes with chemical neurotransmitter gradients across the synaptic cleft, establishing multiple operational timescales within the same physical network.

A physical substrate for memory: Memory is distributed across interacting cell assemblies, or engrams, where synaptic efficacies encode long-term information and internal state variables retain short-lived traces of recent activity. Unlike digital systems, storage is not confined to explicitly addressable memory locations separated from the processing unit.

To replicate biological behavior in a physical computing system, these neuro-computational requirements must be mapped to the device's operational characteristics. The structural mapping ensures that our system does not simply emulate mathematical abstractions but natively exploits its internal physics to execute computational primitives.

A quick map of neuromorphic photonics

Neuromorphic photonics is not a single architecture but a family of design philosophies tailored to different computational demands:

Photonic accelerators: Analog matrix–vector multiplication by exploiting light propagation and interference.

Photonic convolutional neural networks (CNNs): Spatial hierarchy of learned optical filters that automatically extract features before classification.

Reservoir computing (RC): A fixed, nonlinear dynamical system (the reservoir) that projects inputs into a rich feature space, meaning only a simple linear readout layer requires training.

Photonic-spiking neural networks (SNNs): Asynchronous systems that encode information via discrete time-dependent events (spikes), aiming for sparse, robust and energy-efficient processing.

Why light and why microresonators?

This mapping requires a physical medium that offers speed, parallelism, nonlinearity, memory and scalability. Light provides an ideal candidate because photons are non-interacting particles in linear regimes and experience minimal losses traveling down a low-loss dielectric waveguide, which eliminates the Joule heating associated with charge transport in traditional copper interconnects. Integrated photonics

Mapping biological functions onto microring resonator behavior

Biological property	Biological role	Role in a neural network	Microresonator property
Neuronal polarity	Directs information flow from dendrites to axons.	Establishes directed graph connectivity and clear causality, defining distinct input nodes, processing layers and output ports.	Geometry-engineered layouts, such as nonlinear Taiji microresonators, create direction-dependent excitability, giving the optical node a functional polarity.
Spiking	Produces an action potential once input crosses a threshold.	Turns continuous state accumulations into discrete, sub-nanosecond communication events.	Two-photon absorption and free-carrier dispersion interact with thermal effects to generate fast optical pulses or self-pulsing oscillations.
Temporal integration	Sums incoming signals over time.	Enables short-term temporal memory of recent signals for time-series tasks.	Self-pulsing depends on intracavity energy, which increases when low-threshold input signals temporally overlap or constructively interfere.
Refractory period	Creates a recovery pause that prevents backward signaling and stabilizes the network.	Enforces a maximum firing rate and prevents feedback saturation loops.	Combined thermal and free-carrier relaxation set a finite recovery time after each pulse or oscillation cycle.
Local plasticity	Changes connection strength based on local experience.	Updates specific connection strengths (synaptic weights) based on local correlation of node actions (e.g., Hebbian or spike-timing-dependent plasticity (STDP)).	Nonvolatile phase or structural alterations via integrated material overlays, such as phase-change materials or ferroelectric layers, can semi-permanently shift resonances and encode transient weights.
Excitation and inhibition	Balances network activity to prevent signal saturation.	Excitatory links drive the node closer to its activation threshold, while inhibitory links actively suppress activation, establishing homeostatic balance.	Competing blueshift- and redshift-tuning mechanisms, or constructive versus destructive interference, can increase or suppress activation.
Threshold adaptability	Maintains network stability across varying input levels.	Dynamically repositions a neuron's trigger threshold relative to baseline drift, preserving high sensitivity and preventing complete blinding by ambient noise.	Active wavelength tuning via integrated micro-heaters or electrical bias ports (p-i-n junctions), altering the internal energy threshold as demonstrated in Taiji or DRUM microresonators.
Backaction (retrograde)	Backpropagating action potential allows post-synaptic neurons to modify synaptic weight updates.	Acts as a localized timing feedback signal that triggers STDP, adjusting weights based on the arrival order of pre- and post-synaptic spikes without a global supervisor.	Bidirectional optical pathways or backscattering effects that allow downstream nodes to modulate the field properties of upstream elements via counterpropagating modes or localized thermal/electrical cross-talk.
Synchronization	Coordinates network timing for feature binding.	Locks the phase or timing of multiple spiking nodes together, establishing coherent collective dynamics for sensory binding and feature extraction.	Injection locking and mutually coupled resonance modes where individual microresonators synchronize via shared waveguides or feedback paths.

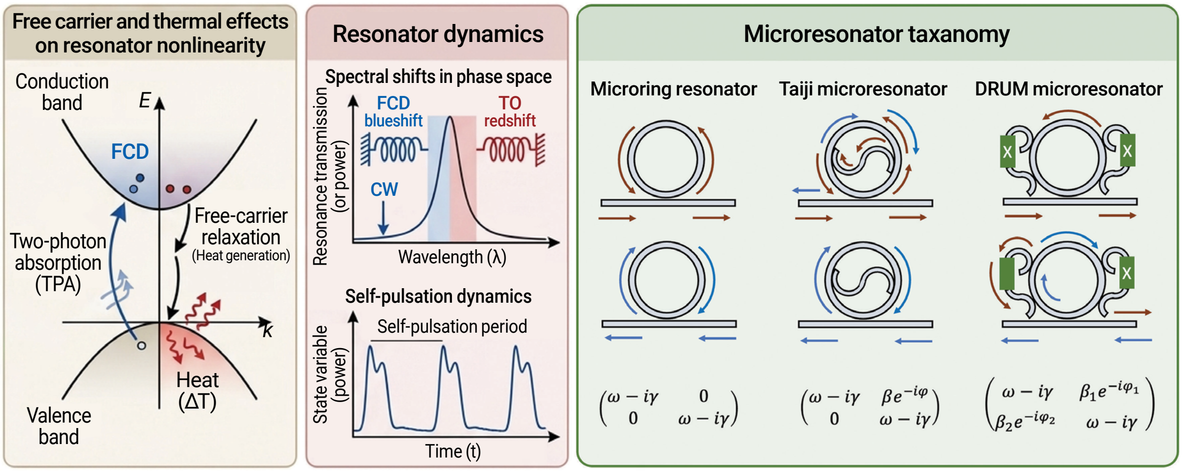

Left: Band diagram and phase-space signatures illustrating the origin of resonator nonlinearity. Center: Due to the opposing signs of free-carrier dispersion (FCD) and the thermo-optic (TO) effect, the microresonator resonance is push-pulled into continuous limit-cycle oscillations under continuous-wave (CW) excitation. Right: Microring configurations and their coupled-mode-theory matrix representations, including a standard microring, a Taiji microresonator with directional mode coupling (nonreciprocal) and a dynamically reconfigurable unified microresonator (DRUM) microresonator with tunable asymmetric intermodal coupling.

Courtesy of the author

possesses three distinct qualities that make it uniquely well-suited for brain-inspired computing structures:

Multi-wavelength parallelism: By using wavelength division multiplexing (WDM), independent data signals can travel through a single optical channel without crosstalk. A single waveguide can therefore act as an expansive, parallel data bus, directly mimicking the highly interconnected wiring of biological dendritic bundles.

Multi-scale temporal dynamics: Integrated optical structures leverage a broad spectrum of physical mechanisms operating across highly diverse timescales. Ultra-fast optical switching can be driven at sub-picosecond speeds (~100 GHz) via second- or third-order electronic Kerr nonlinearities, providing an excellent substrate for high-speed mathematical operations. Conversely, relatively slow microsecond or millisecond responses can be controlled using localized thermo-optic effects. This dual-speed operation bridges high-speed computation and the slower timescales typical of real-world sensory environments.

Continuous-time waveform processing: Telecommunication waveforms, fiber optic sensor outputs, lidar returns and time-stretch imaging or spectroscopic data arrive natively as high-bandwidth, time-dependent series. Processing these signals directly in the optical domain can avoid the severe latency and energy penalties of electronic digitization.

Dimensionality expansion in three optical dimensions

In machine learning software, mapping data into a higher-dimensional feature space often makes complex, overlapping classes linearly separable using simple decision boundaries. Photonic neural networks can perform a related operation physically using three optical dimensions:

Time: Cavity dynamics allow a device to retain a fading trace of previous inputs.

Wavelength: Different spectral colors can probe different optical transfer functions across the device.

Space: Different output ports and paths can sample distinct interference patterns and output states.

The result is not just more data, but a set of concurrent nonlinear representations of the same input.

Silicon photonics has emerged as the dominant deployment platform due to its deep integration with mature silicon-on-insulator manufacturing lines. Complex photonic integrated circuits (PICs) can incorporate hundreds of high-Q resonators, phase modulators, heterogeneously integrated lasers and photodetectors on massive silicon wafers.

Beyond computing, this monolithic compatibility enables advanced optogenetic bio-interfaces. Because light can stimulate or probe neural tissue via

By tuning input power, laser-cavity frequency detuning and coupler configuration, a silicon microresonator can transition across a broad range of operational regimes.

Directionality and the "axon problem"

Biological action potentials generally propagate unidirectionally along the axon toward downstream synapses. This one-way street prevents destabilizing backward echo signals and allows clean, feedforward processing. Integrated photonic circuits, however, suffer from the "axon problem": They are usually reciprocal. A conventional microring responds identically to forward and backward excitation paths. In a scaled spiking network, this structural symmetry can become a major liability, as parasitic reflections and back-action blur signal directionality and make large networks difficult to stabilize.

Geometry-engineered resonators offer a possible solution. In a Taiji microresonator, an internal S-shaped waveguide loop creates asymmetric coupling between clockwise and counterclockwise modes. In a nonlinear regime, this spatial asymmetry can make the threshold for spiking-like behavior lower in one direction than the other, enabling a passive form of directional excitability.

light-sensitive proteins with exceptional spatial and temporal resolution, an integrated silicon photonic neural processor could potentially be coupled directly to live biological structures, paving the way for advanced bio-hybrid computing platforms.

Microring dynamics as a neuron analog

Silicon microring resonators (MRRs) and coupled MRR arrays have become important building blocks for photonic neural networks. Recent work demonstrates that an engineered nonlinear MRR can act as a fully integrated, multi-scale analog of a biological neuron. This capability stems from the rich, coupled nonlinear dynamics that occur when light is confined within a tiny circular, sub-millimeter cavity. Rather than serving as passive spectral filters, microresonators function as active, state-dependent dynamical nodes governed by the interplay of light, free carriers and heat.

When a continuous-wave probe laser is coupled into an MRR near one of its resonant modes, the intracavity optical power density increases drastically. At these elevated field intensities, two-photon absorption can generate a transient population of free electrons in the conduction band and free holes in the valence band. These carriers induce two primary physical responses: free-carrier absorption (FCA), which introduces temporary optical attenuation, and free-carrier dispersion (FCD), which alters the refractive index, shifting the resonator's resonance wavelength toward shorter wavelengths—a blueshift. The carriers recombine and thermalize over a characteristic timescale of about 2 nanoseconds. Crucially, this lifetime can be tuned down to just a few tens of picoseconds by integrating a reverse-biased p-i-n junction across the waveguide architecture to actively sweep the carriers out of the core.

Concurrently, a slower, opposing physical process takes place. Carrier recombination and parasitic absorption generate a localized temperature increase. Because silicon exhibits a strong positive thermo-optic coefficient, heating increases the refractive index, pulling the cavity's resonance peak toward longer wavelengths (a redshift). This thermal process occurs over a slower timescale of approximately 50 to 100 nanoseconds.

The microresonator operates as a competitive, multi-scale dynamical system. The fast-electronic blueshift (~2 ns) and the slow thermal redshift (~100 ns) act as coupled, opposing internal feedback loops. This physical interaction provides a functional analogy to the coupled ion-channel dynamics found in biological cellular membranes, such as the fast sodium activation and slow potassium relaxation loops described by the classic Hodgkin–Huxley neurobiological model.

By exploiting these intrinsic physical mechanisms, the microresonator natively unifies temporal integration and spiking dynamics. An incoming optical pulse sequence entering the input waveguide couples into the ring cavity, incrementally generating free carriers and localized heat. Their dynamical interplay causes the MRR resonance to oscillate with a characteristic frequency governed by the input signals. Because the localized thermal footprint requires roughly 100 nanoseconds to dissipate, it establishes a native refractory period that sets the recovery timescale of each oscillation cycle, gracefully replicating the relaxation phases observed in biological neurons.

These dynamics are useful because they transform high-bandwidth temporal signals into task-relevant features, such as thresholds, delays, correlations and oscillatory patterns. These features can then support computational tasks in communications, sensing and edge-AI applications.

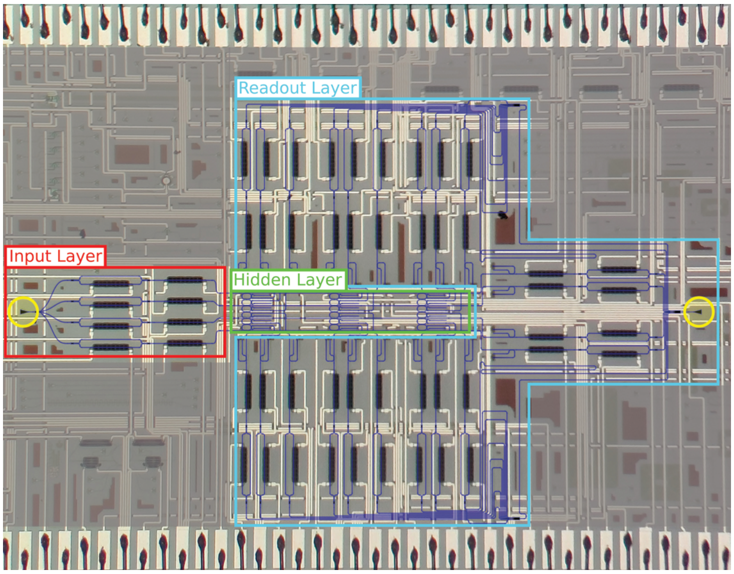

Optical micrograph of a fabricated silicon photonic integrated circuit configured as a high-speed photonic extreme learning machine (PELM). The input layer splits incoming optical data streams into parallel processing channels from a single on-chip injection port (yellow circle). The hidden layer houses an integrated array of nonlinear microresonators that map the input data into a high-dimensional computational state. The readout layer comprises a complex network of routing waveguides and integrated on-chip photodetectors. This layer simultaneously senses the transient nonlinear state of each individual neuron in the hidden layer, converting the optical power distribution into parallel electronic signals ready for linear classification and final output readout.

Courtesy of R. Franchi and S. Biasi

From linear filters to collective intelligence

By tuning input power, laser-cavity frequency detuning and coupler configuration, a silicon microresonator can transition across a broad range of operational regimes. The same physical node on a single monolithic chip can be repurposed for entirely distinct computational tasks.

Linear filtering: At low input powers, the intra-cavity energy remains below the nonlinear activation threshold. The microresonator operates as a standard, state-independent Lorentzian filter for spectral multiplexing, wavelength routing and linear signal weighting.

Multistability and hysteresis: As input power increases and the laser is red-detuned, the slow thermo-optic effect begins to dominate. The device can then exhibit multiple stable optical output states for a single input power level, depending on historical trajectory. This native hysteresis loop provides co-localized short-term optical memory.

Continuous self-pulsing: At higher continuous-wave input power and appropriate detuning, the steady-state operating point becomes unstable. Fast electronic and slow thermal feedback loops drive continuous, self-sustained optical pulse generation at megahertz to gigahertz frequencies. This state can be leveraged as a continuous pacing neuron or oscillatory memory cell.

Collective dynamics and synchronization: When multiple self-pulsing microresonators are interconnected via shared bus waveguides or feedback routing loops, they exchange optical signals. Depending on coupling strength, the system can exhibit periodic, aperiodic or chaotic dynamics. Through mutual injection locking, separate nodes might coordinate their individual spiking profiles, increasing the expressive power of the physical system. This emergent synchronization replicates the collective neural binding processes that biological brains use for complex pattern recognition and sensory feature integration.

Scale, learning and hybrid integration

Microring resonators represent a niche within the broader neuromorphic photonics landscape. Mach-Zehnder interferometer meshes, for instance, are used to execute programmable linear transformations and matrix-vector multiplication, but they generally require separate strategies for nonlinearity and memory storage. Laser-based neuromorphic systems offer native oscillatory or active spiking behavior, but complex integration demands, optical noise and platform-specific gain media can complicate large-scale manufacturing. Electro-optic approaches provide exceptionally fast and controllable modulation, while phase-change and ferroelectric materials are strong candidates for nonvolatile weight storage.

Laser-based neuromorphic systems offer native oscillatory or active spiking behavior, but complex integration demands, optical noise and gain media can complicate large-scale manufacturing.

Where microrings can have an early impact

Intelligent front-end modules: Compact nonlinear front ends performing physical feature expansion on raw data before electronic readout, reducing data movement downstream.

Intelligent sensing: Photonic front ends that extract and retain transient information from optical sensors, identifying critical patterns without requiring immediate, full-rate digitization.

High-speed communications: Real-time channel equalization, dispersion impairment compensation and anomaly detection for high-speed optical telecommunication links.

Ultrafast optical diagnostics: On-the-fly feature extraction for time-stretch imaging and flow cytometry, filtering blank frames before electronic digitization.

The appeal of microrings lies in their convergence of compactness, resonance enhancement and wavelength selectivity. A single microring can perform optical filtering, temporal memory accumulation and threshold nonlinearity within the same footprint. Taiji-style and related non-Hermitian ring designs, including DRUM resonators, add a structural control knob by making excitability direction-dependent. In short, microring networks are especially attractive when compact resonant memory, dense WDM wavelength parallelism and native nonlinear feature generation are required simultaneously within a single, scalable PIC element.

Scaling these individual resonators into dense, multi-node spiking neural networks is the next major leap for the field. By pairing expansive WDM routing grids with arrays of active MRRs, we can build scalable, high-throughput processors capable of analyzing complex, high-bandwidth time-series data in real time. The primary obstacle is control. Resonant devices are highly sensitive to environmental variations. Nanoscale fabrication tolerances can shift cold resonances away from target channels, local thermal drift continuously moves operating points, and coupled arrays can develop unexpected inter-ring crosstalk.

Learning is another challenge. Conventional training techniques, such as backpropagation through time, require tracking and storing every internal physical state variable digitally over time, which can reintroduce the memory wall during training. To bypass this bottleneck, future systems must transition toward hardware-native, physics-inspired learning models. By pairing microresonators with nonvolatile material overlays (e.g. phase-change materials or ferroelectric cladding layers) we can implement local rules such as spike-timing-dependent plasticity directly on chip. Passing optical spikes would modify the physical phase of the cladding material, updating the synaptic connection weights in real time without requiring external digital control pipeline or memory manipulation.

The ultimate evolution is toward hybrid integration. By blending high-speed photonic processing cores with alternative physical substrates (such as electronic memristors, elastomers or optomechanical systems), we can develop highly adaptive, multi-scale computing systems. In these hybrid configurations, optical circuits handle massive, high-bandwidth front-end data signals, while slower, nonvolatile electronic or ionic memory arrays manage long-term structural weight updates and keep the entire system stable and running smoothly. Microring resonators offer an extraordinary combination of compactness, resonance enhancement and controllable coupling for the expanding field of neuromorphic photonics. This biomimetic framing helps directly bridge device-level physics and computational function. By intentionally codesigning hardware to exploit the multi-scale interplay of light and matter, neuromorphic photonics bypasses traditional architectural limits to offer a sustainable path toward low-power artificial intelligence. OPN

Lorenzo Pavesi (lorenzo.pavesi@unitn.it) leads the Nanoscience Laboratory at the University of Trento, Italy.

For references and resources, go online: optica-opn.org/link/neuromorphic-photonics.

References and Resources

S. Biasi et al. "Photonic neural networks based on integrated silicon microresonators," *Intell. Comput.* **3**, 0067 (2024).

A. Lugnan et al. "Reservoir computing with all-optical non-fading memory in a self-pulsing microresonator network," *Adv. Opt. Mater.* **13**, 2403133 (2025).

G. Donati et al. "All-optical spiking processing and reservoir computing with a passive silicon microring and wavelength-time division multiplexing," Photon. Res. **13**, 2641 (2025).

B. Aslan et al. "Coherent control of (non-)Hermitian mode coupling: tunable chirality and exceptional point dynamics in photonic microresonators," Light Sci. Appl. **15**, 150 (2026).

S. Biasi et al. "Mimicking biological neurons with non-hermitian silicon microresonators," *preprint* (2025).

A. Foradori et al. "Memory in integrated photonic neural networks: From physical mechanisms to neuromorphic architectures," APL Photonics *arXiv:2604.22620* (2026).