

## TOPICAL REVIEW

# Beyond Silicon: Materials, Mechanisms, and Methods for Physical Neural Computing

STEFAN FISCHER<sup>1</sup>, (Member, IEEE), NIHAT AY<sup>2</sup>, OLAF LANDSIEDEL<sup>2,3,4,5</sup>,  
ESFANDIAR MOHAMMADI<sup>1</sup>, SEBASTIAN OTTE<sup>1</sup>, BERND-CHRISTIAN RENNER<sup>2</sup>,  
AND NELE RUSSWINKEL<sup>1</sup>

<sup>1</sup>University of Luebeck, 23562 Lübeck, Germany

<sup>2</sup>Hamburg University of Technology, 21073 Hamburg, Germany

<sup>3</sup>United Nations University Hub on Engineering to Face Climate Change, Hamburg University of Technology, 21073 Hamburg, Germany

<sup>4</sup>United Nations University Institute for Water, Environment and Health (UNU-INWEH), 21073 Hamburg, Germany

<sup>5</sup>Kiel University, 24118 Kiel, Germany

Corresponding author: Stefan Fischer (stefan.fischer@uni-luebeck.de)

**ABSTRACT** Physical implementations of neural computation now extend far beyond silicon hardware, encompassing substrates such as memristive devices, photonic circuits, mechanical metamaterials, microfluidic networks, chemical reaction systems, and living neural tissue. By exploiting intrinsic physical processes, such as charge transport, wave interference, elastic deformation, mass transport, and biochemical regulation, these substrates can realize neural inference and adaptation directly in matter. As silicon GPU-centered AI faces growing energy and data-movement constraints, physical neural computation becomes increasingly relevant as a complementary path beyond conventional digital accelerators. This trend is driven in particular by pervasive intelligence, i.e., the deployment of on-device and edge AI across large numbers of resource-constrained systems. In such settings, co-locating computation with sensing and memory can reduce data shuttling and improve efficiency. Meanwhile, physical neural approaches have emerged across disparate disciplines, yet progress remains fragmented, with limited shared terminology and few principled ways to compare platforms. This survey unifies the field by mapping neural primitives to substrate-specific mechanisms, analyzing architectural and training paradigms, and identifying key engineering constraints including scalability, precision, programmability, and I/O interfacing overhead. To enable cross-domain comparison, we introduce a first-order benchmarking scheme based on standardized static and dynamic tasks and physically interpretable performance dimensions. We show that no single substrate dominates across the considered dimensions; instead, physical neural systems occupy complementary operating regimes, enabling applications ranging from ultrafast signal processing and in-memory inference to embodied control and in-sample biochemical decision making.

**INDEX TERMS** Physical neural computing, neuromorphic hardware, memristive systems, photonic neural networks, mechanical metamaterials, microfluidic and chemical computing, reservoir computing, in-memory computing, benchmarking frameworks.

## I. INTRODUCTION

In the domain of artificial intelligence, in-silico implementations of Artificial Neural Networks (ANN) are currently dominating. From the earliest perceptrons to modern Large Language Models (LLMs) and Foundation Models

The associate editor coordinating the review of this manuscript and approving it for publication was Domenico Rosaci<sup>1</sup>.

(FMs), the substrate of choice has been the transistor-based integrated circuit. Digital CMOS has dominated AI hardware mainly due to manufacturability, programmability, and ecosystem maturity—but not because alternative physical substrates are fundamentally incapable. As we approach the asymptotic limits of Moore's Law and Dennard Scaling [1], the research community is increasingly forced to look “beyond silicon.” However, the motivation to explore

alternative substrates extends far beyond the mere stagnation of lithographic improvements. Beyond data-center training, the rise of pervasive intelligence in the form of on-device and edge AI deployed across large numbers of resource-constrained systems makes energy and internal data-transport costs first-order design constraints. In this setting, physical neural computing becomes newly relevant as a complementary path to conventional digital accelerators by co-locating computation with sensing and memory and exploiting substrate dynamics to reduce data shuttling. This survey explores the emerging field of Physical Neural Networks (PNNs), systems that do not merely simulate neural dynamics digitally, but physically embody them in “wet” (biological/chemical) and “solid/fluid” (mechanical/microfluidic) media. These developments motivate treating a coherent research direction centered on physical neural networks, where substrate physics is not an implementation detail but a determinant of how learning itself can be realized [2].

### A. BEYOND THE TURING AND VON NEUMANN PARADIGMS

Contemporary Artificial Neural Networks are abstractions. They are mathematical constructs composed of weights, biases, and activation functions, executed on hardware that is not originally tailored for their mode of operation. The mainstream computing architecture, the Von Neumann machine, relies on a strict separation of processing units (CPU/GPU) and memory.

This separation gives rise to the “Von Neumann bottleneck,” where a significant fraction of both energy and execution time is primarily consumed by data movement rather than by the computation itself [3]. As ANN models continue to scale in size and complexity, this mismatch between algorithmic structure and hardware organization becomes increasingly salient. PNNs instead couple algorithmic operations to substrate dynamics, embedding computation directly in physical state evolution [2].

While neuromorphic engineering has attempted to bridge this gap using silicon (e.g. through digital layouts optimized for sparse, event-driven computation, analog sub-threshold circuits, or memristors) [4], [5], these approaches still largely rely on electron flow in solid semiconductors. They remain constrained by the physics of the electron in a rigid lattice. By contrast, biological brains (the inspiration for ANNs) utilize ions, chemical gradients, and fluid dynamics in a wet, three-dimensional medium.

Replicating and potentially even surpassing biological efficiency demands a fundamental rethinking of the physical substrate of neural computation and learning, way beyond evidently suboptimal electron transport in silicon-based hardware.

### B. THE NEED FOR PHYSICAL NEURAL NETWORKS (PNNs)

The motivation for non-silicon neural networks is often reduced to energy efficiency. While this is valid, it is

incomplete. We outline core arguments for PNNs that enable capabilities not achievable with conventional electronic computing alone.

#### 1) ELIMINATING THE “TRANSDUCTION TAX” (IN-SITU COMPUTATION)

Current cyber-physical pipelines rely on costly transduction chains. For example, biomarker detection with deep learning requires converting a chemical signal into an electrical current, amplifying and digitizing it (ADC), processing it on digital hardware, and often converting results back into analog actions. Each step adds latency, noise, and energy loss. Chemical PNNs can instead process molecules directly via reaction-diffusion dynamics [6], computing in the data’s native domain. This transduction overhead is quantifiable: ADC and digital signal processing can dominate the energy budget of sensing pipelines, in some systems accounting for most of total consumption [7]. Nucleic-acid-based molecular logic toolkits demonstrate diagnostic decisions directly in biological fluids without electronic readout [8], supporting in-situ computation as a practical way to bypass the transduction chain.

#### 2) THERMODYNAMIC COMPUTING: PHYSICS AS LOGIC

Digital logic expends energy to maintain stable “0” and “1” states against thermal noise. PNNs can instead exploit thermodynamic relaxation. Mechanical or chemical systems naturally converge to energy minima, which can encode solutions to optimization problems [9]. Spring glasses and elastic networks relax into states that map to Ising-type energy landscapes [10], [11]. In such systems, relaxation, diffusion, or wave propagation directly perform the computation.

#### 3) VOLUMETRIC DENSITY AND CONNECTIVITY

Conventional circuits are essentially planar, with 3D stacking constrained by heat and interconnect limits. Wet and molecular computing substrates are inherently volumetric. Even small droplets can host extremely large numbers of interacting molecules, with 3D diffusion providing connectivity densities that can exceed solid-state wiring schemes [12].

#### 4) THE “CLOCKLESS” CONTINUUM

Digital simulation of physics requires time discretization and numerical integration, introducing approximation error and clock overhead. Physical substrates operate in continuous time. A fluidic system implements Navier-Stokes [13] dynamics natively; the computational question is the I/O mapping and controllability of that dynamics. While temporal resolution is fundamentally bounded only by physical limits (Planck time), practical analog systems are constrained by thermal and material noise, yielding an effective signal-to-noise-limited bit depth often below 32-bit floating point precision [14].

### 5) RESILIENCE AND SELF-HEALING

Semiconductor hardware is structurally brittle, whereas biological and fluidic substrates can remain functional under deformation or partial damage. Reaction-diffusion media such as Belousov–Zhabotinsky systems can continue operating despite container deformation or division [15]. Synthetic biological implementations additionally offer self-repair and regeneration capabilities absent in silicon systems.

### 6) ENVIRONMENTAL INTEGRATION AND BIOCOMPATIBILITY

Silicon devices are poorly suited for saline or biological environments. PNNs based on DNA, proteins, or biocompatible fluids can operate inside living tissue, in the bloodstream, or in sensitive ecosystems [8]. This enables embedded intelligence for medical and ecological applications.

## C. SCOPE

This review focuses on *physical neural networks* (PNNs), defined as systems in which training and inference exploit the intrinsic dynamics of a material substrate rather than being realized solely through numerical operations on conventional digital processors [2], [16]. The emphasis is on computation implemented through physical state evolution, including charge transport, wave propagation, mechanical deformation, mass transport, chemical kinetics, and biological plasticity.

To maintain a clear technical focus, the following topics are explicitly excluded:

- **In-silico neural networks:** conventional artificial neural networks executed on CPUs, GPUs, or TPUs without essential reliance on substrate physics.
- **Standard CMOS neuromorphic hardware:** digital or mixed-signal silicon platforms such as Loihi [17] or TrueNorth [18], unless their operation fundamentally depends on non-silicon physical effects or materials.
- **Quantum neural computing:** neural models based on quantum superposition or entanglement, which constitute a distinct computational paradigm with separate theoretical and engineering challenges.

Physical neural computing, as used in this review, overlaps with but is not identical to several neighboring research areas. First, it is narrower than analog computing: while analog systems in general process continuous variables through physical dynamics, they are only considered physical neural systems here if these dynamics realize neural primitives such as weighted summation, nonlinearity, memory, or adaptation in a functionally neural architecture. Second, it overlaps only partially with neuromorphic hardware. Neuromorphic systems are inspired by neural organization, often through spiking, event-driven, or in-memory architectures, but many such systems remain conventional CMOS implementations and therefore fall outside our scope unless their computational advantage fundamentally relies on non-silicon substrate physics. Third, reservoir computing is treated here not as a separate substrate class, but as a computational regime in which a physical system with rich internal

dynamics serves as a fixed reservoir and only the readout is trained. Reservoir computing can therefore be realized within several of the substrate classes reviewed below.

Within these boundaries, the paper is organized into four conceptual parts that span the design space of physical neural computing.

### 1) PART 1: FOUNDATIONS AND CONCEPTUAL FRAMEWORK (SECTIONS I–II)

The first part motivates physical neural computing, defines core terminology, and maps neural primitives (weighted summation, nonlinearity, memory, adaptation) to physical mechanisms. It positions PNNs relative to neuromorphic engineering and digital ML accelerators, and highlights system-level challenges when computation is embedded in physical substrates.

### 2) PART 2: WETWARE AND MOLECULAR INTELLIGENCE (SECTIONS III–V)

The second part surveys chemical and biological substrates, from DNA strand displacement and molecular reaction circuits to reaction-diffusion systems and synthetic biological platforms. It emphasizes how biochemical kinetics, diffusion, and cellular regulation realize neural primitives and enable partially autonomous computation in chemical or living media.

### 3) PART 3: PHYSICAL NEURAL HARDWARE SUBSTRATES (SECTIONS VI–XI)

The third part reviews engineered non-biological platforms, including memristive, phase-change, and ferroelectric non-volatile memory devices, spintronic and superconducting systems, photonic, mechanical, metamaterial, microfluidic, and iontronic networks. The focus is on shared architectures, training and calibration strategies, and trade-offs in bandwidth, precision, scalability, robustness, and I/O overhead.

### 4) PART 4: CROSS-SUBSTRATE BENCHMARKING AND OUTLOOK (SECTIONS XII–XIII)

The final part introduces a comparative benchmarking framework based on standardized static and dynamic tasks and physically meaningful metrics. It then discusses design principles, application domains, and the key challenges in translating laboratory PNN concepts into scalable engineering technologies.

In our view, the recent proliferation of physical neural substrates represents not merely a diversification of hardware platforms, but a shift toward treating physical dynamics as a first-class computational resource. This perspective motivates the unified treatment adopted throughout the review.

## II. THE PHYSICS OF NEURAL COMPUTATION

To engineer intelligence in physical substrates, we must relate neural computation to the governing laws of physics. Conventional hardware implements neural models indirectly

through Boolean logic. In contrast, Physical Neural Networks (PNNs) implement them directly in the substrate: the physical interaction itself performs the computation. This section therefore does not only generalize the neural primitives of connectivity (weights), nonlinearity, and optimization (learning) across both “wet” (diffusion/reaction) and “solid/fluid” (transport/interaction) domains; it also makes explicit how material properties determine which neural primitives are natively available, and how device and network architecture determine how these primitives are organized into a computational system. In other words, the material sets the native computational vocabulary, whereas the architecture governs signal propagation, coupling structure, and the overall organization of computation.

### A. PHYSICS AS COMPUTATION: TRANSPORT, INTERACTION, AND RELAXATION

A generic neural network relies on three core operations: *weighted signal transmission*, *nonlinear activation*, and *state evolution toward a solution*. In Physical Neural Networks (PNNs), these primitives are not executed symbolically by Boolean logic; instead, they arise directly from substrate physics through (i) linear transport and superposition, (ii) intrinsic nonlinear interaction and thresholds, and (iii) relaxation dynamics that implement optimization.

#### 1) SIGNAL TRANSPORT AND WEIGHTING (THE LINEAR REGIME)

In ANNs, the weighted sum  $\sum_j w_{ij}x_j$  is the fundamental linear operation. In wetware, such linear accumulation is naturally mediated by diffusion: molecular connectivity is effectively “wireless” and governed by Fick’s laws, so concentration fields spread and mix without explicit wiring. At cellular scales, flows typically occur at very low Reynolds numbers where viscosity dominates inertia; motion is laminar and effectively time-reversible, making diffusion—rather than turbulent mixing—the dominant mechanism for distributing signals [19]. In engineered hardware, the same linear primitive is realized by conservation laws: Kirchhoff-type constraints make nodal summation of currents or flows thermodynamically “free” [20]. Here,  $w_{ij}$  becomes a physical parameter such as electrical conductance, hydraulic resistance, or (in mechanical networks) an entry of the stiffness matrix that maps forces and displacements. Across both wetware and hardware, the shared computational essence is linear superposition: physics performs accumulation by default, with weights encoded in geometry or material parameters.

#### 2) PHYSICAL INTERACTION (THE NON-LINEAR REGIME)

A purely linear network collapses to a single matrix multiplication and cannot represent functions requiring decision boundaries (e.g., XOR). Nonlinearity in PNNs therefore comes from exploiting material interactions rather than explicitly programming an activation function. In chemical

systems, reaction kinetics provide nonlinear transforms directly: the law of mass action yields native multiplicative terms (e.g., rates proportional to  $[A][B]$ ), enabling polynomial computation [21], while enzymatic or catalytic saturation produces transfer curves that can closely match sigmoidal neuron models [22]. In solid/fluidic hardware, nonlinearity is often engineered via instabilities and thresholds. Mechanical elements can exhibit bifurcations (e.g., buckling) that act as sharp threshold functions and thus enable discrete decisions from continuous inputs [23]; similarly, microfluidic valves can switch state above a pressure threshold, implementing physical gating. In all cases, the computational role of nonlinearity is the same: critical thresholds break superposition ( $f(a + b) \neq f(a) + f(b)$ ) and partition state space, enabling expressive function approximation.

#### 3) RELAXATION AS OPTIMIZATION (THE SOLVER)

In conventional computing, optimization is implemented as an explicit iterative algorithm that updates parameters step-by-step. In many physical substrates, by contrast, “optimization” is an intrinsic tendency to relax toward lower free energy or lower internal stress. In wetware, this is exemplified by molecular constraint satisfaction: DNA strands spontaneously bind to Watson–Crick complements to minimize free energy, and massive parallelism arises because many candidate assemblies explore the solution space simultaneously. Adleman’s seminal demonstration showed how such thermodynamic relaxation can be harnessed to solve combinatorial problems by encoding constraints into molecular binding and then letting the system settle [24]. In engineered hardware, analogous solver behavior appears in systems whose dynamics minimize an implicit cost function, such as spin glasses or elastic networks that relax to reduce frustration or stress. The connection between annealing in statistical physics and combinatorial optimization was formalized by Kirkpatrick et al. [25]. Conceptually, these systems admit a scalar “energy” (or Hamiltonian) that acts as a Lyapunov-like function with a non-increasing trajectory,  $\frac{dH}{dt} \leq 0$ , so that mapping problem constraints onto the substrate’s energy landscape turns computation into controlled relaxation rather than explicit iteration.

### B. TRAINING STRATEGIES FOR PHYSICAL SYSTEMS

In deep learning, training is typically cast as minimizing a loss function via gradient descent, with backpropagation providing exact gradients  $\nabla_{\theta} \mathcal{L}$  for each parameter  $\theta$ . Physical substrates do not natively expose a “backward pass”: atoms neither store gradients nor provide explicit credit assignment. The resulting *credit assignment problem* asks how a global output error can be attributed to local, embodied parameters (e.g., a specific channel conductance or beam stiffness). In practice, training strategies for Physical Neural Networks (PNNs) fall into three broad paradigms. At a high level, these strategies can be grouped into ex-situ digital-twin optimization, hybrid digital–physical training, and in-materio

learning based on physical gradient approximation or related local update rules. In addition, when gradients are unavailable or unreliable, some substrates rely on closed-box search procedures such as evolutionary optimization. Which of these approaches is most suitable depends strongly on substrate characteristics, including tunability, observability, relaxation time, and fabrication variability. The corresponding substrate-specific realizations are therefore discussed in the “Training and Design Paradigms” subsections of Sections III to XI.

### 1) IN-SILICO TRAINING (DIGITAL TWIN/EX-SITU OPTIMIZATION)

In-silico training decouples learning from execution by treating the device as an inference engine. A differentiable *digital twin* (e.g., implemented in PyTorch or JAX) is trained with standard backpropagation, and the converged parameters are then transferred to the physical instance in a one-shot programming step. The main advantage is immediate access to mature optimizers and architectures (e.g., Adam/RMSprop, CNNs/RNNs) without changing the hardware. The central limitation is the *reality gap*: unmodeled noise, tolerances, and nonlinearities can cause a model to overfit simulation artifacts and fail to transfer reliably [16]. Moreover, accurate deployment typically requires calibration of the twin to the specific fabricated instance, reducing plug-and-play operation [26]. This strategy is therefore particularly attractive for substrates whose forward dynamics can be modeled with sufficient fidelity and whose parameters can be programmed reproducibly, but that provide little support for direct in-device adaptation.

### 2) HYBRID DIGITAL–PHYSICAL TRAINING (RESERVOIR COMPUTING)

A second class comprises hybrid digital–physical schemes, in which optimization remains partly external while the physical substrate directly performs the forward computation. The most established example is Reservoir Computing (RC), which avoids training the substrate itself by exploiting its intrinsic nonlinear dynamics. The physical system acts as a fixed dynamical reservoir that projects inputs into a high-dimensional spatiotemporal state; only a simple linear readout (often digital linear regression) is trained to map these states to targets. Conceptually, the reservoir functions like a physical kernel, turning nonlinear separability into a linear readout problem. Because the reservoir is fixed, training is computationally light and typically admits a closed-form solution (no epoch-wise backpropagation). The approach is also substrate-agnostic: reservoirs have been demonstrated in diverse media such as water-bucket dynamics [27], soft robotic bodies [28], and optical delay systems [29]. The trade-off is task sub-optimality: since the internal features are not optimized end-to-end, comparable accuracy may require substantially larger reservoirs [30]. Hybrid schemes are thus especially appealing for substrates with rich native dynamics

but limited fine-grained weight programmability, whereas they are less suitable when peak task-specific performance requires end-to-end optimization of the internal state. As a side remark: in this review, reservoir computing is treated as a cross-cutting training and architectural paradigm within physical neural computing rather than as a separate material class.

### 3) IN-MATERIO LEARNING (PHYSICAL GRADIENTS/IN-SITU UPDATES)

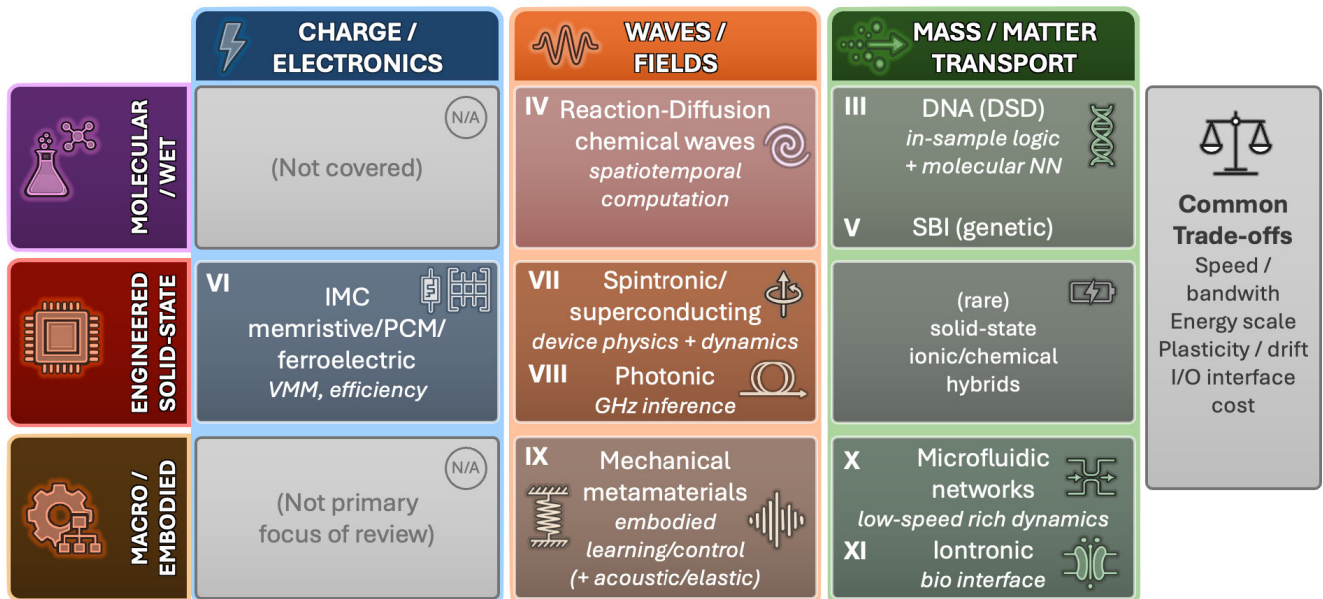
The most ambitious paradigm is to perform learning *within* the physics, thereby closing the reality gap by letting the device compute (or approximate) its own gradients. A prominent framework is *Equilibrium Propagation (EqProp)*: the system is observed in a “free” phase (inputs only) and in a “clamped” phase (inputs plus a small output nudging), and differences between the resulting physical states yield an estimate proportional to the gradient [31]. In principle, this enables local, contrastive updates without storing a global gradient matrix [32]. In practice, it raises substantial engineering challenges: implementing fast, stable, continuously tunable weights is difficult, and the need to physically relax (often twice per update) can make learning slow compared to purely digital training for small- to medium-scale systems [33]. More broadly, this category includes in-situ update rules that use physically measured state differences or other local surrogate signals to approximate credit assignment directly in the substrate. Overall, PNN training remains less mature than large-scale digital backpropagation, but the combination of hybrid and in-situ approaches demonstrates that physically embodied learning is feasible under realistic constraints [2]. Such approaches are most promising for substrates that offer stable, repeatable, and sufficiently local update mechanisms, and remain difficult to realize when observability is poor, relaxation is slow, or device drift dominates the dynamics.

## C. TAXONOMY OF PHYSICAL SUBSTRATES

Physical neural substrates span everything from molecular test tubes to macroscopic mechanical lattices. To navigate this diversity without over-indexing on “wet” versus “dry” materials, we classify substrates by computational characteristics that mirror core abstractions in computer science: the *data type* used for signal encoding, the *memory architecture* used for plasticity, and the *execution model* implied by their temporal dynamics.

### 1) SIGNAL ENCODING (THE DATA TYPE)

Most PNNs are intrinsically analog: information is carried by continuous thermodynamic variables (e.g., concentration, pressure, strain) whose effective precision is limited by noise rather than by a fixed word length. This maps naturally to continuous-time recurrent models, where graded states evolve under coupled dynamics; Hopfield’s analysis of networks with graded responses provides an early and influential



**FIGURE 1.** Landscape of physical substrates for computation and learning. The horizontal axis organizes approaches by the dominant coupling carrier (charge and electronics, waves and fields, and mass and matter transport), while the vertical axis groups them by substrate class and physical scale (molecular and wet, engineered solid-state, and macro and embodied systems). Each box summarizes a representative paradigm and indicates the corresponding chapter(s). Across this landscape, material physics determines which neural primitives are natively available in a substrate, whereas architecture determines how these primitives are arranged into scalable computational organizations.

formalization of this continuous regime [34]. A smaller but important class of substrates is engineered for bistability, enabling digital-style encodings where the presence/absence of a molecular species or a snapped mechanical mode represents a Boolean state; DNA strand displacement is a canonical example, and Qian and Winfree demonstrated that it can implement arbitrary Boolean logic circuits [12]. Finally, event-driven encodings represent information primarily in *timing* rather than amplitude, as in spiking or excitable media where propagating wavefronts act as discrete events; collision-based computing in reaction–diffusion systems illustrates how such impulse-based dynamics can support logic operations [6].

## 2) PLASTICITY IMPLEMENTATION (THE MEMORY ARCHITECTURE)

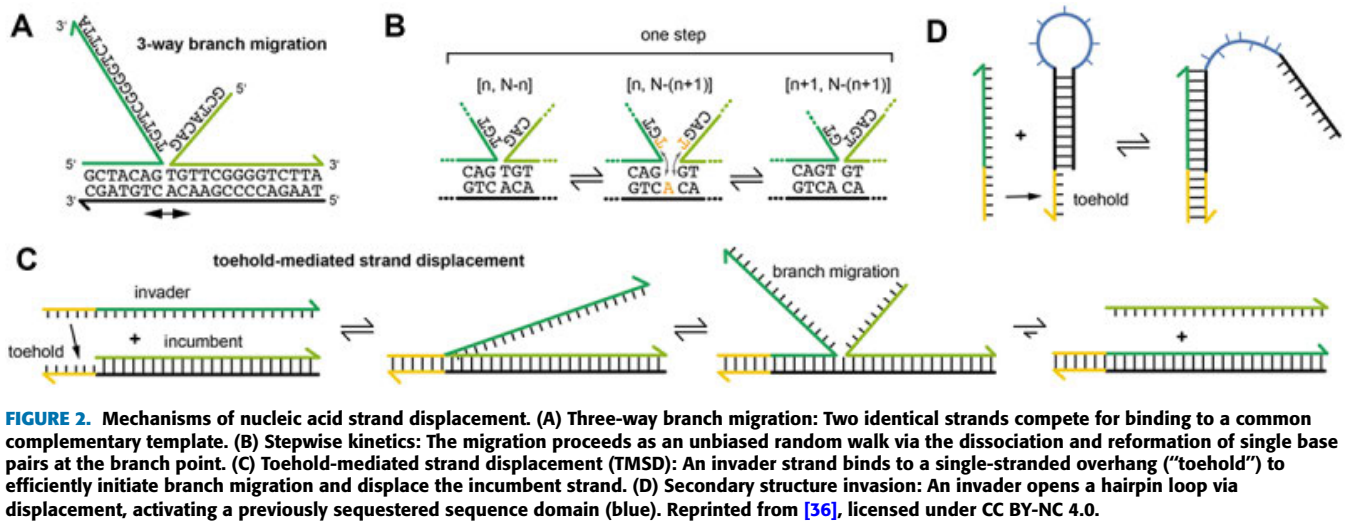
In digital learning systems, weights are explicit variables updated by an external algorithm. In physical substrates, by contrast, “memory” is realized only insofar as material properties can be configured and retained. Some platforms are effectively read-only: geometry fixes the connectivity and coupling strengths during fabrication, yielding ASIC-like inference engines commonly used in reservoir-style setups; passive photonic reservoirs are a representative example [29]. Other systems are programmable but only through external intervention (e.g., mechanical adjustment, chemical titration), enabling reconfiguration without true autonomy. The strongest notion of plasticity arises when the substrate locally adapts in response to its own signal flux, providing a material analogue of Hebbian update rules;

memristive devices exemplify this state-dependent memory, as highlighted in the foundational memristor formulation by Strukov et al. [5].

## 3) TEMPORAL DYNAMICS (THE EXECUTION MODEL)

A final differentiator is how computation uses time. In equilibrium-style systems, the input initializes the substrate, and the computation is the relaxation to a steady state; intermediate trajectories are not part of the “output” and are treated as transient. Molecular computing via thermodynamic equilibration is a classic illustration, including Adleman’s use of DNA to solve a Hamiltonian path instance by letting the system settle under designed constraints [24]. In dynamic systems, by contrast, information is encoded in the evolving trajectory itself, so outputs depend on input history through inertia, hysteresis, or internal state; this corresponds to sequential processing with fading memory. Jaeger’s “echo state” formulation provides a compact criterion for when such driven dynamics can support time-series processing reliably [35].

The following chapters survey the field along the major physical regimes introduced here, moving from molecular and biochemical substrates to engineered solid-state platforms and embodied fluidic or mechanical systems. Figure 1 provides a compact overview of this landscape and situates the reviewed approaches within a common comparative mapping. While Figure 1 organizes the field primarily by substrate class and dominant physical carrier (as described in Subsections II-A and II-C), two further cross-cutting distinctions are important. First, material



properties such as transport mechanism, switching behavior, intrinsic nonlinearity, stability, and temporal response determine which neural primitives are natively available in a given substrate, including weighted summation, thresholding, memory, or plasticity. Second, device and network architecture determine how these primitives are organized into computation, for example through crossbar arrays for parallel vector–matrix multiplication, coupled oscillators or reservoirs for dynamical processing, or spatially distributed cells, droplets, and reaction domains for propagating local interactions. As already explained in Section II-B, learning strategies form an additional dimension that cuts across these material categories: in the reviewed literature, ex-situ digital-twin optimization, hybrid reservoir/readout training, in-materio adaptation, and, where needed, closed-box evolutionary search recur as the main learning regimes; substrate-specific realizations are discussed in the “Training and Design Paradigms” subsections of Sections III to XI.

### III. MOLECULAR COMPUTING WITH DNA

DNA is often viewed merely as a passive storage medium for genetic information (comparable to a hard drive). However, it possesses thermodynamic properties that make it also a powerful substrate for active computation. By stripping DNA of its biological context and treating it purely as a programmable polymer, researchers have engineered systems capable of executing complex neural logic in a test tube.

#### A. THE SUBSTRATE PRINCIPLE

The dominant mechanism for implementing molecular circuits is DNA strand displacement (DSD), a process driven by the thermodynamic tendency of DNA strands to maximize base-pairing. By designing “toehold” sequences that initiate branch migration, an input strand can displace a pre-hybridized output strand from a gate complex, effectively releasing a signal into the solution [36], [37] (see Fig. 2). This mechanism functions analogously to signal transmission

at a synapse, but operates purely through enthalpy-driven kinetics.

While early models abstracted DSD as digital logic, contemporary approaches exploit the continuous nature of reaction rates to implement analog computation, where signal strength is encoded in molar concentrations rather than binary presence. Recent advances in “molecular commutation” further abstract away the structural details, utilizing inverse design algorithms to compute optimal dissociation constant matrices for a desired network behavior [38]. This shifts the design paradigm from manually assembling logic gates to mathematically optimizing the free-energy landscape of the reaction network.

#### B. MAPPING NEURAL PRIMITIVES

To implement neural architectures in DNA, abstract mathematical operations must be mapped isomorphically to specific biochemical reaction kinetics.

##### 1) WEIGHTS AND SIGNALS

Information is encoded in the continuous molar concentration of specific DNA strands. Synaptic weighting is realized through the stoichiometry of strand displacement reactions: the concentration of intermediate “gate” complexes modulates the gain between an input strand and the release of a downstream output strand [12]. Alternatively, DNA origami offers a spatially addressable substrate where binary or analog weights can be encoded via the physical arrangement of molecules on a nanoscale scaffold [39], [40]. Summation occurs naturally through the parallel contribution of multiple reactant pathways to the production of a common output species.

##### 2) ACTIVATION AND NON-LINEARITY

To enable decision boundaries and deep logic, linear hybridization kinetics are augmented with molecular sequestration. By introducing high-affinity “sink” strands that

bind and neutralize signal strands, the system enforces a concentration threshold below which no output is generated. This competitive annihilation process implements a function mathematically equivalent to the Rectified Linear Unit (ReLU) [21]. More advanced “winner-take-all” mechanisms group signals to compete against each other, allowing for efficient non-linear decision making in classification tasks [41].

### C. ARCHITECTURES AND I/O INTERFACES

The topology of DNA neural networks is not defined by physical placement but by the specificity of molecular interactions in a well-mixed solution. Research has evolved from static logic gates to dynamic systems capable of temporal processing and cellular integration.

Early architectures focused on implementing static input-output mappings akin to Boolean logic or feed-forward classifiers. The “Winner-Take-All” (WTA) networks demonstrated by Cherry and Qian utilize competitive hybridization to implement non-linear classification boundaries, successfully classifying 100-bit patterns (handwritten digits) [42]. Similarly, ‘seesaw gate’ networks, which established a framework for scaling up digital circuit computation [43], exploit reversible strand displacement to realize analog weighted summation and thresholding, enabling the construction of sigmoid-like transfer functions purely through mass-action kinetics [12].

A long-time bottleneck for DNA computing has been signal degradation and “crosstalk” in deep networks. The CALCUL system (Classified Allosteric-toehold based Continuous and Ultra-accurate Learning) addresses this by introducing spatial isolation via magnetic beads. This separation allows for layer-by-layer processing and the implementation of convolutional operations (weight sharing) that were previously hindered by the “soup” nature of bulk solutions. This architecture has demonstrated high-accuracy classification of image patterns by physically preventing unwanted feedback loops [44].

To process time-varying signals or solve optimization problems, architectures have incorporated feedback loops and autocatalysis. Building on the foundational associative memory models proposed by Hopfield [45], DNA-based Hopfield networks utilize the energy landscape of the reaction system to store state information. Recent work demonstrated a discrete Hopfield network capable of solving combinatorial optimization tasks (such as Sudoku puzzles) by relaxing into a stable energy minimum representing the solution [46].

Data entry is typically achieved by introducing specific strand concentrations [12] or via light-sensitive molecular switches [47]. The computational result is read out through fluorescence kinetics using fluorophore-quencher pairs, where the intensity of the emitted light corresponds to the final output activation level [12].

### D. TRAINING AND DESIGN PARADIGMS

Training strategies for DNA neural networks are evolving from static design to dynamic adaptation. At present, the dominant paradigm is ex-situ optimization via in-silico modeling and inverse design, whereas fully in-materio adaptation remains an emerging capability.

The currently dominant approach relies on in silico optimization, where the desired network behavior is modeled mathematically. Algorithms compute the optimal matrix of dissociation constants or reaction rates, which are subsequently translated into DNA sequences using inverse design tools [38]. The physical molecules are then synthesized to embody these pre-calculated weights.

Emerging paradigms demonstrate autonomous learning directly within the substrate. Recent experimental work has realized supervised learning in test tubes, where “training strands” (representing data and labels) interact with the network to physically adjust the concentration of weight molecules. This allows the system to update its decision boundaries purely through thermodynamic equilibration, without external digital processing [48]. This makes DNA particularly attractive when reaction pathways can be modeled and compiled into molecular programs with high specificity, but less suitable for fast, repeated, and fine-grained in-situ weight adaptation because kinetics remain slow and parameter tuning is tightly coupled to sequence design.

### E. REPRESENTATIVE DEMONSTRATIONS

Representative demonstrations of molecular neural computation span cell-free pattern classification, multiplexed diagnostic decision-making, and intracellular neuromorphic control. Cherry and Qian’s molecular classifier for handwritten digits (MNIST) showed that complex non-linear decision boundaries can be realized in a well-mixed liquid medium, demonstrating that DNA strand-displacement networks can perform genuine pattern classification rather than merely simple logic operations [42]. In biosensing, molecular logic circuits have been used for multiplexed diagnostic decision-making by integrating several microRNA inputs into a single molecular output, thereby improving specificity over single-marker assays [49], [50], [51]. Moving from cell-free systems into living cells, synthetic gene circuits further show that neuromorphic computation can be embedded into cellular regulation: “perceptgenes” in bacteria implement transcriptionally regulated input integration [52], while CRISPR-Cas9/gRNA architectures combined with RNA sequestration realize perceptron-like threshold units for classification and regression tasks [53]. Taken together, these studies show that the same general molecular design principles can support pattern recognition, diagnostic inference, and intracellular computation across increasingly application-near settings.

## F. ENGINEERING CONSTRAINTS

The transition from theoretical molecular algorithms to physical implementation imposes specific thermodynamic and kinetic constraints that define the operating regime of DNA computing.

The most distinct feature of the DNA substrate is the extreme disparity between information density and propagation speed. Molecular systems offer storage densities theoretically approaching  $10^{19}$  bits/cm<sup>3</sup>, with a single microliter capable of containing trillions of independent processing threads [12]. This allows for massive parallelism at thermodynamic limits of energy efficiency. However, operation bandwidth is diffusion-limited, typically restricting complex strand-displacement cascades to timescales of minutes or hours to reach equilibrium [42].

Another long-time bottleneck for molecular logic has been the “single-use” nature of the hardware. However, recent developments in 2025 have begun to mitigate this constraint. Sun et al. demonstrated that non-complementary DNA neural networks can decouple the “weight” molecules from the input signals, allowing inputs to be chemically washed from the system while preserving the trained parameters [54]. Furthermore, “heat-rechargeable” circuits have been proposed that utilize thermal cycling not merely for resetting, but to actively drive the system into high-energy non-equilibrium states, effectively recharging the circuit’s computational potential [55].

## G. POSITION WITHIN PHYSICAL NEURAL COMPUTING

DNA-based neural systems occupy an extreme position in the landscape of physical neural computing. They operate at the smallest physical scales and lowest characteristic bandwidths, while offering exceptionally high intrinsic parallelism. Architecturally, they realize chemically embedded computation: weighted summation and nonlinearity are implemented directly by reaction stoichiometry, distinguishing them from electronic or photonic substrates, where physical processes accelerate numerical operations but do not replace symbolic computation. Their relevance lies in a complementary regime characterized by direct processing of molecular inputs, biochemical compatibility, and extreme energy efficiency.

## IV. CHEMICAL AND REACTION-DIFFUSION SYSTEMS

While DNA computing typically relies on thermodynamic equilibrium and the static storage of genetic information, chemical computing exploits the non-equilibrium dynamics of dissipative structures. By utilizing reaction-diffusion (RD) systems and active matter, computation is shifted from the passive readout of code to the active propagation of information in space and time [6]. In this paradigm, the material substrate behaves as a continuous excitable medium, capable of processing information through wave propagation and synchronization rather than symbolic manipulation.

## A. THE SUBSTRATE PRINCIPLE

Unlike the discrete logic of silicon or the hybridization kinetics of DNA, chemical computers operate by physically embodying partial differential equations. The fundamental “machine language” of this substrate is the Reaction-Diffusion equation:

$$\frac{\partial u}{\partial t} = D\nabla^2 u + R(u) \quad (1)$$

where  $u$  represents the vector of chemical concentrations. The diffusion term  $D\nabla^2 u$  mediates spatial coupling (signal transport), while the reaction term  $R(u)$  dictates the non-linear local kinetics (processing) [6].

Historically investigated in macroscopic vessels, recent advances have shifted towards “On-Chip Chemical Computing”. Agostini et al. demonstrated the integration of oscillating Belousov-Zhabotinsky (BZ) reaction chambers directly onto silicon microchips, effectively creating hybrid architectures where the chemical medium performs massive parallel processing while electronic overlays handle I/O [56]. A defining thermodynamic feature of these systems is their dissipative nature: computation is a continuous dynamic process maintained by a constant flux of reagents. Unlike non-volatile memory, if the energy supply ceases, the system relaxes to equilibrium and the computational state is lost [57].

## B. MAPPING NEURAL PRIMITIVES

To engineer intelligence into fluid substrates, the abstract mathematical operations of artificial neural networks must be mapped to the kinetic laws governing chemical reactions.

### 1) SIGNALS AND WEIGHTS

In chemical neural networks, information is typically encoded in the molar concentration of molecular species or the phase of an oscillatory wave. Synaptic weights, which modulate signal transmission, are physically instantiated by reaction rate constants ( $k$ ) and diffusion coefficients ( $D$ ). While traditional homogeneous systems rely on intrinsic kinetic parameters, recent microfluidic implementations allow for spatially programmable weights by altering channel geometries or utilizing responsive hydrogels that modulate local diffusion rates [58].

A key advantage of this chemical signaling is its inherent robustness. Comparative studies suggest that chemical synapses exhibit superior anti-interference characteristics compared to electrical coupling, naturally filtering high-frequency noise and stabilizing synchronization in a manner that is difficult to replicate in standard electronics [59].

### 2) ACTIVATION AND NON-LINEARITY

The linear accumulation of signals is handled natively by diffusion and mixing. The essential non-linearity required for deep computation arises from the Law of Mass Action and enzyme kinetics. Since reaction rates are proportional to the product of reactant concentrations, chemical systems provide

a native multiplication operation, enabling the computation of polynomial functions [21].

Furthermore, saturation effects in enzymatic reactions yield concentration curves mathematically identical to sigmoidal transfer functions. Sharp thresholding behaviors, equivalent to the Rectified Linear Unit (ReLU), are implemented via chemical titration mechanisms, where a signal molecule must neutralize a specific inhibitor concentration before triggering a downstream reaction [22].

### C. ARCHITECTURES AND I/O INTERFACES

The architectural landscape of chemical computing has shifted from rigid, geometrically constrained logic circuits to amorphous, self-organizing dynamical systems.

Early implementations focused on collision-based computing, where traveling chemical wavefronts in excitable media (such as the above-mentioned BZ reaction) act as information carriers. Logical operations are realized through the precise geometrical routing of these waves: annihilation of colliding wavefronts implements NOT or XOR functions, while channel junctions perform OR logic [6], [60]. While conceptually foundational, these systems face scalability challenges due to the need for precise spatial structuring and the slow propagation speed of diffusion waves.

To overcome the constraints of explicit circuit design, the field has largely adopted the reservoir computing paradigm. Here, the “messy” complexity of non-linear reaction networks, such as the autocatalytic Formose reaction or electrochemical oscillators, serves as a high-dimensional projection space. The chemical substrate naturally maps low-dimensional inputs (e.g., flow rates) into a vast state space of intermediate concentrations, allowing a simple trained linear readout to extract complex temporal features [57], [61]. This approach utilizes the intrinsic thermodynamic relaxation of the medium as a computational resource.

Moving beyond bulk solutions, compartmentalized architectures utilize networks of interacting droplets or “active matter” swarms. In droplet-based systems, individual reaction vessels act as discrete nodes coupled via diffusion across lipid interfaces, forming programmable chemical lattices [62]. Conversely, active matter reservoirs employ self-propelled particles (e.g., bacteria or Janus colloids) whose collective hydrodynamics create reconfigurable network topologies, introducing advection as a signal transport mechanism to potentially surpass diffusion limits [63].

### D. TRAINING AND DESIGN PARADIGMS

Training chemical systems presents a unique challenge: physical atoms do not store gradients, and the “credit assignment problem” is complicated by the opacity of the reaction mixture. Accordingly, chemical substrates are currently dominated by ex-situ digital-twin modeling and hybrid reservoir-style training, while closed-box evolutionary optimization is often used when internal kinetics are too complex for direct gradient-based design.

To bridge the “reality gap” between simulation and wetware, recent approaches utilize Physics-Informed Neural Networks (PINNs) and Neural Ordinary Differential Equations (ODEs) [64]. Frameworks such as SPIN-ODE embed the governing chemical rate equations directly into the training loop, inferring the underlying kinetics from observational trajectories. This allows the construction of “digital twins” that are robust to noise and experimental variability, enabling offline optimization of control parameters that transfer reliably to the physical reactor [65], [66]. This paradigm is particularly attractive when reaction kinetics can be modeled with sufficient fidelity, but it remains limited by partial observability, stochastic fluctuations, and calibration effort for specific experimental instances.

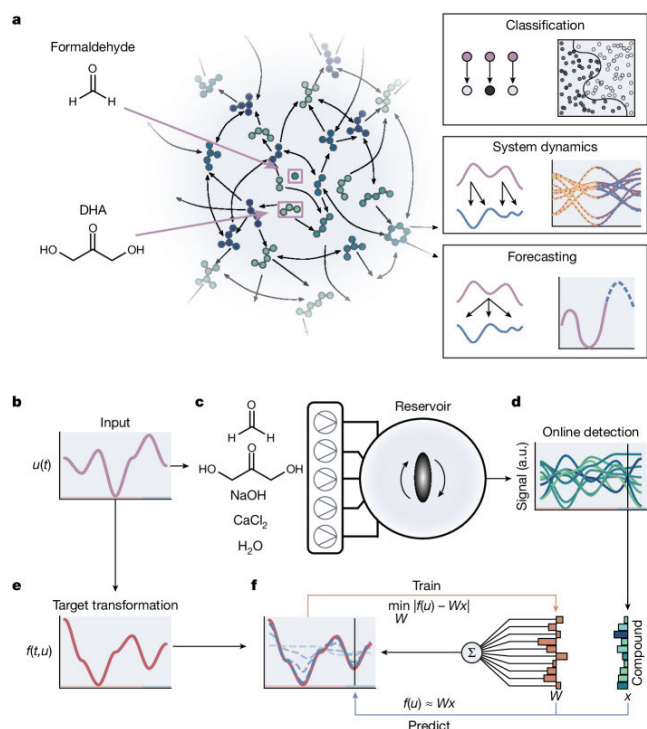
In the reservoir regime, training is restricted to the readout layer (typically a mass spectrometer or optical sensor), which is optimized via linear regression. For the substrate itself, evolutionary algorithms are often employed to select reagents or droplet topologies that maximize the “richness” or entropy of the chemical dynamics, effectively evolving the material’s computational capacity without requiring a differentiable model of the complex internal kinetics [67]. Such hybrid and closed-box strategies are well matched to chemically rich but only weakly observable substrates, although they generally sacrifice end-to-end optimality and fine-grained in-situ adaptation.

### E. REPRESENTATIVE DEMONSTRATIONS

Representative demonstrations of chemical neural computation span four regimes: temporal reservoir computing, embodied chemo-mechanical control, spatial analog problem solving, and structured oscillatory processing. The already mentioned “Formose Reservoir” exemplifies the use of molecular chaos for computation. Baltussen et al. showed that the autocatalytic polymerization of formaldehyde, previously regarded as an uncontrollable side reaction, exhibits rich transient dynamics that can serve as a physical reservoir. By modulating the inflow rates of reagents in a continuous stirred-tank reactor, the system generates characteristic molecular states that enable chaotic time-series prediction and non-linear classification, outperforming linear models by exploiting the intrinsic complexity of the chemical substrate [57]. Figure 3 illustrates this principle and highlights how low-dimensional inputs are mapped into a high-dimensional chemical state space from which only the readout layer is trained.

In the domain of soft robotics, chemical intelligence enables control without microcontrollers. Self-oscillating gels driven by the Belousov–Zhabotinsky reaction have been engineered to function simultaneously as sensor, controller, and actuator. These materials exhibit autonomous peristaltic motion and phototactic behavior, thereby embodying the control loop directly in the chemo-mechanical properties of the material [58].

The classical application of excitable chemical media remains spatial computing. Reaction-diffusion waves



**FIGURE 3.** Architecture of a chemical reservoir computer based on the Formose Reaction. (a) The chemical reaction network functions as a physical reservoir (a fixed, non-linear hidden layer). It projects low-dimensional inputs (reactants) into a high-dimensional chemical state space. This intrinsic non-linearity enables the linear separability of complex computational tasks, such as classification and time-series forecasting. (b–f) signal flow and readout training. A time-dependent input signal  $u(t)$  modulates the inflow rates of reagents into a continuous stirred-tank reactor (CSTR). The resulting non-linear chemical dynamics (the state vector  $x$ ) are monitored in real-time via mass spectrometry. Consistent with the reservoir computing paradigm, training is restricted to the readout layer: a linear weight matrix  $W$  is optimized via regression to map the high-dimensional chemical state to the target output function ( $y \approx Wx$ ). Reprinted from [57], licensed under CC BY 4.0.

naturally explore geometric spaces in parallel. Steinbock et al. and later Lagzi et al. demonstrated that chemical wavefronts can solve maze-navigation and Voronoi-tessellation problems: the wave propagates through all possible paths simultaneously, and the first wavefront reaching the exit identifies the optimal trajectory [68], [69].

In contrast to the unstructured dynamics of the Formose system, Parrilla-Gutierrez et al. demonstrated a programmable chemical processor based on a structured Belousov–Zhabotinsky cell array. Their system realized a chemical autoencoder for pattern recognition and encoded addressable memory states directly in the oscillatory phase and amplitude of the reaction vessels [67].

### F. ENGINEERING CONSTRAINTS

The practical deployment of chemical neural networks is governed by thermodynamic necessities and interface limitations that distinguish them sharply from solid-state electronics.

Unlike DNA computing (which often relies on equilibrium states) or CMOS (where static power can be minimized), reaction-diffusion systems are inherently dissipative structures. Computation is a dynamic process maintained only by a continuous flux of free energy (reagents). This creates a fueling constraint: the system requires a constant supply of chemical “fuel” to sustain the non-equilibrium states necessary for wave propagation and oscillation. Consequently, these systems are best suited for environments where chemical energy is abundant and naturally replenished, such as within biological fluids or waste streams, rather than in battery-constrained portable devices [57].

The internal parallelism of a chemical droplet is immense (processing on the order of  $10^{20}$  molecular interactions simultaneously). However, accessing this state remains a bottleneck. Reading out the complex spectrum of a Formose reservoir typically requires mass spectrometry, a slow and bulky process that negates the miniaturization benefits. However, the recent move toward “lab-on-a-chip” architectures that integrate electronic sensors directly with BZ reaction chambers offers a path to mitigate this transduction tax, enabling real-time electrical readout of chemical states [56].

### G. POSITION WITHIN PHYSICAL NEURAL COMPUTING

Chemical neural systems occupy a distinct position between molecular and macroscopic physical computation. They operate at spatial scales ranging from micrometers to centimeters and at time scales set by diffusion and reaction kinetics, resulting in low bandwidth but rich spatiotemporal dynamics.

In contrast to DNA-based networks, chemical systems implement neural operations through continuous concentration fields and propagating wavefronts rather than discrete molecular interactions. Weighted interactions and nonlinear activation arise from diffusion coupling and chemical feedback loops. These properties make chemical substrates especially suitable for spatial pattern processing. Their main relevance is therefore not high-speed inference, but applications in which computation over space and time is intrinsic to the problem, such as morphological pattern formation, distributed sensing, and control in chemically active environments.

### V. SYNTHETIC BIOLOGICAL INTELLIGENCE

Synthetic Biological Intelligence (SBI) extends chemical computation into living matter. Unlike molecular or chemical substrates, which are programmable but passive, SBI exploits intrinsic biological functions such as metabolism, homeostasis, self-repair, and plasticity as computational resources. As a result, learning is not only imposed externally through weight updates, but can also emerge from biological mechanisms of adaptation and free-energy minimization [70].

#### A. THE SUBSTRATE PRINCIPLE

Research in SBI currently focuses on three distinct substrate classes representing increasing levels of biological

complexity. At the genetic level, engineered bacterial consortia implement distributed logic. Early work by Elowitz and Leibler demonstrated synthetic oscillatory networks (“Repressilator”), while Tamsir et al. and Danino et al. realized synchronized quorum sensing and Boolean gates within bacterial populations [71], [72], [73]. At the morphological level, the slime mold *Physarum polycephalum* acts as a macroscopic, multinucleate single cell capable of solving spatial optimization tasks by dynamically reconfiguring its protoplasmic tube network [74], [75]. Most recently, “Organoid Intelligence” (OI) has emerged, utilizing three-dimensional neural cultures derived from human induced pluripotent stem cells (iPSCs). These organoids develop dense and interconnected neuronal networks with spontaneous spiking activity and long-term potentiation. Among the SBI substrates discussed here, they most closely mimic the cytoarchitecture of the human cortex [70], [76].

## B. MAPPING NEURAL PRIMITIVES

Living substrates realize neural primitives through physical growth, electrophysiology, and biochemical regulation.

### 1) WEIGHTS AND PLASTICITY

In *Physarum*, the “weight” of a connection is physically instantiated by the diameter of the protoplasmic tube, which determines the cytoplasmic flux. This adaptive resizing follows a positive feedback loop:  $\Delta w_{ij} \propto \text{flux}_{ij}$  [74]. In neuronal organoids, weights correspond to synaptic efficacy. Their adaptation follows biologically conserved rules such as spike-timing-dependent plasticity (STDP), in which the temporal correlation between pre- and post-synaptic activity determines potentiation or depression [77].

### 2) ACTIVATION AND DYNAMICS

Activation in slime molds is continuous and oscillatory, driven by rhythmic contraction waves. Organoids, conversely, operate in the spiking regime, encoding information in discrete action potentials and population firing rates. In genetic circuits, activation is transcriptional, defined by the concentration of repressor or activator proteins binding to DNA promoters [71]. Going beyond digital logic, Daniel et al. demonstrated that these transcriptional networks can also be engineered to perform synthetic analog computation, such as addition and division in the logarithmic domain, directly within living cells [78].

### 3) OPTIMIZATION VIA HOMEOSTASIS

Unlike artificial systems that minimize an explicit loss function, biological networks appear to optimize for homeostasis and predictability. The Free Energy Principle proposes that living systems adapt their internal states to minimize the difference between predicted and sensed sensory inputs (sensory surprise), effectively performing active inference to maintain their structural integrity [79].

## C. ARCHITECTURES AND I/O INTERFACES

Architectures in SBI are typically emergent rather than lithographically defined, requiring sophisticated interfaces to guide computation.

*Physarum* naturally forms spatial graphs. When food sources are arranged to represent nodes in a geometric problem, the organism relaxes into a Steiner tree configuration, effectively computing the shortest path that connects all points while balancing metabolic cost against transport efficiency [74].

Brain organoids are typically interfaced using High-Density Microelectrode Arrays (HD-MEAs). These devices provide thousands of bidirectional channels that allow for both the recording of population spikes and the delivery of patterned electrical stimulation [80]. The organoid functions as a recurrent reservoir, transforming input stimulation patterns into high-dimensional spatiotemporal readouts [76].

To enable systematic experimentation at scale, Jordan et al. introduced an open and remotely accessible neuroplatform that integrates long-term electrophysiological recording, automated microfluidic control, continuous environmental regulation, and closed-loop stimulation for large numbers of neural organoids [81].

## D. TRAINING AND DESIGN PARADIGMS

Since the internal state of a living cell is largely opaque (“closed box”), training relies on closed-loop interaction and environmental shaping. Accordingly, living substrates are currently dominated by hybrid closed-loop training and ex-situ circuit design, whereas direct in-materio credit assignment remains largely out of reach.

The most promising training paradigm involves placing the biological culture in a closed feedback loop with a simulated environment (“DishBrain”). Sensory inputs are provided through electrical stimulation, and motor outputs are read from neural activity. The culture then receives structured feedback: predictable stimuli for correct behavior and noisy stimuli for incorrect behavior. This closed loop encourages the system to reorganize its connectivity toward more predictable states [82]. This paradigm is attractive because it exploits the adaptive plasticity of living neural tissue without requiring an explicit internal model. However, it offers only limited interpretability, weak control over credit assignment, and modest reproducibility across biological instances.

For bacterial systems, “training” corresponds to the iterative design of plasmid sequences. Combinatorial logic is implemented by layering transcriptional repressors (e.g., NOR gates), allowing for the construction of asynchronous digital circuits inside living cells [72], [83]. To scale such designs, Nielsen et al. introduced “Cello”, a design automation framework that compiles abstract Verilog specifications directly into functional DNA sequences, thereby solving the problem of manual interference between gates [84]. In such systems, the dominant paradigm is therefore closer to ex-situ

digital design and compilation than to learning in the neural-network sense: biological function is programmed into the genetic architecture before deployment. This is powerful for constructing reliable logic behavior, but less suitable for fast online adaptation or repeated gradient-based weight updates.

### E. REPRESENTATIVE DEMONSTRATIONS

Representative demonstrations of SBI span two characteristic regimes: closed-loop learning in neuronal cultures and self-organized spatial optimization in living networks. Kagan et al. demonstrated that human cortical neurons cultured on a multielectrode array can be embedded into a real-time sensorimotor loop with the arcade game “Pong” [82]. Electrical stimulation encoded the game state, neural activity was decoded as motor output, and structured feedback drove rapid adaptation within minutes, showing that living neural tissue can perform goal-directed learning when coupled to an interactive environment. Figure 4 illustrates this architecture and highlights that sensing, computation, and adaptation emerge jointly from the biological culture and its interface.

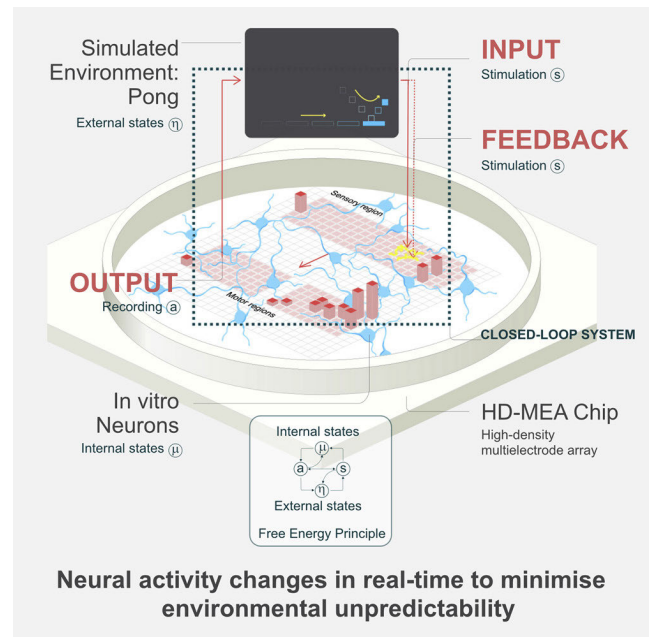
In a complementary form of SBI, *Physarum* was used to reproduce the topology of the Tokyo railway system [74]. When oat flakes were placed at locations corresponding to major cities, the organism grew a transport network that balanced efficiency, connectivity, and fault tolerance, closely matching the engineered rail layout. This demonstration is representative because it shows that living matter can solve a nontrivial spatial optimization problem through growth and adaptive reconfiguration rather than through explicitly programmed digital control.

### F. ENGINEERING CONSTRAINTS

The utilization of living matter introduces unique engineering challenges centered on maintenance and reproducibility. First, the energetic advantage of biological tissue (which operates at orders of magnitude lower power than silicon) is offset by the significant overhead of life support systems (sterile environments, nutrient perfusion). Second, biological variability remains a critical bottleneck. Unlike transistors, no two organoids are identical in cytoarchitecture or connectome, making standardized performance difficult to guarantee. Third, the “I/O bandwidth” is constrained by the bio-electronic interface; while the tissue processes information massively in parallel, the number of electrodes available for reading and writing state remains limited compared to the number of neurons [70], [80].

### G. POSITION WITHIN PHYSICAL NEURAL COMPUTING

SBI occupies the extreme end of the physical intelligence spectrum, characterized by maximal adaptivity and minimal programmability. It offers unparalleled learning efficiency and structural plasticity but lacks the speed and deterministic reliability of electronic hardware. As such, SBI is best viewed as a complementary paradigm for applications requiring high sensitivity, biocompatibility, or the study of intelligence itself in its native substrate.



**FIGURE 4.** Schematic of the closed-loop “Pong” experiment. Human cortical neurons cultured on a high-density multielectrode array (HD-MEA) are bidirectionally coupled to a simulated game environment: game states are encoded as electrical stimulation patterns, neural responses are recorded and decoded as motor output, and structured feedback is returned to the culture. The setup illustrates how synthetic biological intelligence can realize real-time learning through closed-loop interaction aimed at reducing environmental unpredictability. Reprinted from [82], licensed under CC BY 4.0.

## VI. RESISTIVE AND FERROELECTRIC IN-MEMORY COMPUTING

Memristive, phase-change, and ferroelectric systems implement computation through the physics of non-volatile conductance and polarization modulation. Unlike von Neumann architectures that shuttle data between distinct processing and storage units, these substrates realize “in-memory computing” (IMC), where the memory element itself performs the computation. By organizing resistive switching devices, which are typically based on filamentary metal-oxides (ReRAM), chalcogenide phase-change materials (PCM), or ferroelectric oxides (FeRAM), into dense crossbar arrays, vector-matrix multiplication (VMM) is executed directly through the relaxation of electrical currents. This offers a path to massive parallelism and energy efficiency that circumvents the memory wall [5], [85]. Recently, this paradigm has expanded to “in-sensor computing,” where ferroelectric materials simultaneously sense physical stimuli and perform synaptic weighting, eliminating the analog-to-digital conversion overhead at the edge [86].

### A. THE SUBSTRATE PRINCIPLE

The fundamental principle relies on the tunable electrical state of a two-terminal device, governed by history-dependent hysteresis loops.

In ReRAM, the conductance is modulated by the field-driven migration of oxygen vacancies, which form conductive filaments. In PCM, programming pulses induce a reversible phase transition between a highly conductive crystalline state and a highly resistive amorphous state [87]. Beyond resistive switching, nanoscale spintronic oscillators have also been shown to serve as non-linear neurons [88].

Ferroelectric devices (FeFETs, FTJs) store information in the remanent polarization of the material lattice. Unlike ReRAM, which relies on atomic filament formation, ferroelectrics rely on the displacement of ions within the unit cell. Applying an electric field switches the polarization domains. Crucially, partial switching of these domains allows for the storage of continuous analog values. In their review, Khan et al. identify this technology as a key enabler for future non-volatile logic, arguing that HfO<sub>2</sub>-based FeFETs provide the necessary scalability to merge logic and memory at the transistor level [89]. Recent advances in flexible electronics utilize ultra-thin Indium Tin Oxide (ITO) channels combined with ferroelectric gates to realize transparent, high-performance synaptic devices suitable for wearable applications [90]. Furthermore, the movement of “Domain Walls” (interfaces between regions of different polarization) exhibits rich, non-linear dynamics that can be exploited for reservoir computing, where the position of the wall naturally encodes the reservoir state [91].

When arranged in a crossbar topology, all these devices exploit Kirchhoff’s circuit laws to perform analog arithmetic. An input voltage vector applied to the rows generates currents proportional to the conductance (weight) at each crosspoint (Ohm’s Law). These currents sum naturally along the columns (Kirchhoff’s Current Law), physically computing the dot product  $\sum_j V_i G_{ij}$  in a single time step, regardless of the array size [85].

## B. MAPPING NEURAL PRIMITIVES

To function as a neural network, the analog physics of the crossbar must be mapped to the linear and non-linear operations of deep learning.

### 1) WEIGHTS AND SUMMATION

Synaptic weights are encoded as the analog conductance values ( $G$ ) or polarization states ( $P$ ) of the devices. Since conductance is strictly positive, signed weights are typically realized using a differential pair architecture, where the effective weight is the difference between two conductances ( $W = G^+ - G^-$ ). In ferroelectric tunnel junctions (FTJs), the weight is determined by the tunneling probability, modulated by the height of the ferroelectric barrier potential.

### 2) ACTIVATION AND NONLINEARITY

Unlike chemical or mechanical substrates, where non-linearity is often intrinsic to the material, memristive arrays typically perform only the linear transformation. The non-linear activation function (e.g., ReLU, Sigmoid) is usually

implemented at the array periphery using mixed-signal circuits. However, ferroelectric materials offer a distinct advantage here: the intrinsic  $P - E$  (Polarization-Electric Field) hysteresis loop is naturally sigmoidal. Recent reviews highlight how this property, combined with topological textures like domain walls, allows for fully passive non-linear layers suitable for reservoir computing [91].

## C. ARCHITECTURES AND I/O INTERFACES

The architectural paradigm is shifting from pure digital accelerators to heterogeneous systems that integrate sensing and computing.

To mitigate device variability and limited dynamic range, architectures employ redundancy. A single synaptic weight may be represented by multiple physical devices (multi-memristive synapses). Arbitration logic or averaging mechanisms are then used to combine these conductances, effectively increasing the bit-precision and reducing the impact of stochastic switching noise inherent to the nanoscale physics [92].

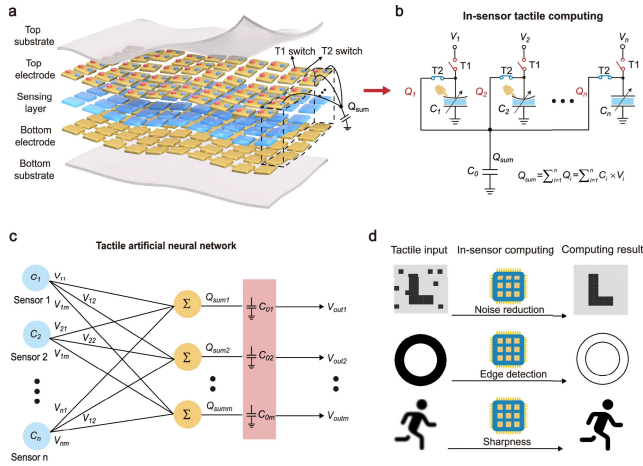
As already discussed above, a critical bottleneck in standard neuromorphic hardware is the transduction tax. Ferroelectric and capacitive architectures address this via “in-sensor computing.” Chen et al. demonstrated a system where the sensor elements themselves perform early-stage neural processing. By utilizing a flexible capacitive pressure sensor array, the system executes in situ analog multiplication and accumulation directly within the sensor network. This architecture bypasses the need for redundant data transfer to central processing units, effectively merging the “skin” and the “brain” of the device [7]. The idea of this approach is shown in Figure 5.

For standard crossbars, the interface remains a challenge. Input vectors are converted via DACs, and output currents are digitized by ADCs. Optimizing the energy and area of these peripheral circuits is critical, as they can dominate the total power budget [85].

## D. TRAINING AND DESIGN PARADIGMS

Training requires navigating the “non-idealities” of the substrate, such as asymmetric update/decay curves and cycle-to-cycle variability. Accordingly, resistive and ferroelectric substrates support a comparatively broad spectrum of training strategies, ranging from ex-situ device-aware optimization to hybrid in-situ updates and reservoir-style use of intrinsic device dynamics.

In the ex-situ training paradigm, the network is trained on a GPU using a high-precision digital model. The resulting weights are mapped onto the physical conductances. “Device-aware” training techniques explicitly model noise and drift during the training phase, making the network robust to the imperfections of the hardware deployment [93]. This approach is attractive because it leverages mature digital optimization while compensating for substrate imperfections,



**FIGURE 5. Capacitive in-sensor tactile computing architecture.** (a) Schematic of the flexible capacitive sensor array integrated with electrical switches. (b) Circuit diagram demonstrating in-situ multiplication and accumulation (MAC) operations within a sensor subregion. (c) Equivalent tactile neural network mapping input sensor capacitance vectors  $C$  to voltage weights  $V$ . (d) Examples of low-level sensory tasks, including noise reduction, edge detection, and sharpness extraction, computed directly within the sensor layer. The figure illustrates how early neural processing can be embedded directly into the sensing substrate, thereby reducing data transfer to external processing units. Reprinted from [7], licensed under CC BY 4.0.

but it remains limited by write variability, conductance drift, and the reality gap between simulated and fabricated devices.

In-situ approaches update the weights directly on the chip. A powerful hybrid approach is “Mixed-Precision Computing,” where the crossbar performs the bulk of the low-precision forward/backward pass arithmetic, while a digital unit accumulates the high-precision gradient updates. This strategy combines the efficiency of analog physics with the numerical stability of digital algorithms [94]. Such hybrid schemes are especially well matched to crossbar arrays because they exploit native analog parallelism while offloading numerically delicate update accumulation to digital control.

A particularly memristor-friendly learning regime is *physical reservoir computing*, where the memristive substrate is not trained as a weight array but exploited for its intrinsic nonlinear, history-dependent dynamics. In a representative demonstration, [95] experimentally implemented reservoir computing using volatile (“dynamic”)  $WO_x$  memristors whose internal ionic relaxation provides short-term memory, mapping temporal input streams into separable conductance-state trajectories. Only a simple readout is trained (e.g., linear regression), enabling temporal inference without iterative high-precision conductance programming. More recently, memristive reservoir-style models have also been explored at the algorithmic and modeling level for time-series processing [96]. This regime is particularly attractive for temporal tasks, although it generally sacrifices full end-to-end optimization of the internal device states.

Building on this memristive-reservoir perspective, [97] shows how memristive-inspired temporal dynamics can be

reorganized into a parallelized reservoir architecture for scalable time-series processing. Empirically, this substrate-aware reformulation achieves competitive, and in several cases superior, accuracy relative to modern gradient-based sequence models while reducing training from minutes or hours to the order of seconds, and in some cases to sub-second runtimes. Taken together, these results point to a broader co-design direction for physical neural computing: efficiency on memristive and in-memory substrates may depend not only on improved devices or calibration schemes, but also on model architectures whose computational structure is shaped from the outset by the dynamics and constraints of the target substrate.

### E. REPRESENTATIVE DEMONSTRATIONS

Representative demonstrations of resistive and ferroelectric neural hardware span four regimes: in-situ learning in small arrays, large-scale analog inference, flexible synaptic devices for wearable electronics, and integrated in-sensor computing. Prezioso et al. demonstrated one of the first integrated memristor perceptrons, showing that simple pattern classification can be achieved directly in hardware with in-situ training rather than by offline weight transfer alone [98]. Scaling up, Burr et al. used large PCM crossbar arrays to perform high-accuracy inference on MNIST and CIFAR while introducing drift-compensation strategies that addressed one of the key non-idealities of resistive memory hardware [87], [93]. In the domain of flexible electronics, Li et al. realized high-performance FeFET devices based on ultra-thin indium tin oxide channels, demonstrating that synaptic functionality can be combined with transparency, flexibility, and suitability for wearable systems [90]. At the system level, Chen et al. implemented an integrated in-sensor tactile computing architecture that performs low-level operations such as noise reduction and edge detection directly within the sensor layer, reducing data movement and achieving more than a 22-fold lower power consumption than conventional mixed-signal approaches [7].

### F. ENGINEERING CONSTRAINTS

Despite their promise, these systems face significant physical limitations. Particularly in PCM, stored weight values relax over time due to conductance drift, degrading accuracy. In contrast, ferroelectric devices (FeFETs) typically exhibit superior retention characteristics and lack the significant resistance drift of PCM, though they suffer from “imprint” effects. Furthermore, endurance poses a challenge for ReRAM and PCM, which suffer from limited cycling endurance (often  $10^8$ – $10^9$  cycles). Ferroelectric oxides, however, demonstrate significantly higher endurance ( $> 10^{10}$  cycles) and lower switching energies, making them preferable for applications requiring frequent weight updates [85]. Finally, voltage drops (IR drop) along the nanowires of large crossbars distort the effective voltage seen by devices, limiting the maximum feasible array size.

### G. POSITION WITHIN PHYSICAL NEURAL COMPUTING

Resistive and ferroelectric systems occupy the most technologically mature position in the landscape of physical neural computing. They are the primary candidates for solid-state “In-Memory Computing,” offering a direct path to break the von Neumann bottleneck. While PCM and ReRAM offer high density, the inclusion of ferroelectrics introduces superior linearity and energy efficiency, particularly for edge sensing applications. Their success, however, is inseparable from algorithm-hardware co-design: they are not ideal analog calculators, but noisy physical systems that require robust algorithmic wrappers to function reliably.

## VII. SPINTRONIC AND SUPERCONDUCTING NEURAL SYSTEMS

While resistive memory architectures (Section VI) focus on density and in-memory multiplication, a distinct class of physical substrates targets the extremes of temporal resolution and energy efficiency. Spintronic and superconducting systems diverge from charge-based electronics by utilizing magnetic state variables as the carriers of information, namely the electron spin and the magnetic flux quantum. These substrates are uniquely characterized by their ability to support ultrafast, oscillator-based, and spike-based computing paradigms that mimic the pulse-coded communication of biological neurons at speeds approaching the gigahertz and terahertz regimes [99]. We group them here as the “Magnetic and Quantum-Electronic Class,” unified by their reliance on collective quantum phenomena (ferromagnetism and superconductivity) to realize non-volatile dynamics with minimal dissipation.<sup>1</sup>

### A. THE SUBSTRATE PRINCIPLE

This class exploits two fundamental physical regimes: room-temperature magnetic dynamics and cryogenic macroscopic quantum coherence.

Spintronics utilizes the intrinsic angular momentum (spin) of electrons. The core computational element is often the Spin-Torque Nano-Oscillator (STNO). When a DC current passes through a magnetic junction, it transfers angular momentum, inducing a steady precession of the magnetization. This converts a constant electrical input into a microwave frequency output, acting as a tunable non-linear oscillator [100]. Moving beyond single devices, topologically stable, vortex-like spin textures such as magnetic Skyrmions exhibit strong non-linear dynamics that can be harnessed for neuromorphic and reservoir computing schemes [101], [102], [103].

Superconducting computing operates at cryogenic temperatures (typically 4 K), where resistance vanishes, and quantized magnetic flux pulses (Single Flux Quanta, SFQ) represent the information carriers. Josephson Junctions (JJs)

<sup>1</sup>Here, “quantum” refers to quantum-mechanical device physics (e.g., flux quantization in superconducting circuits), not to qubit-based quantum neural computing that relies on superposition/entanglement, which is excluded in Section I-C.

inherently generate such quantized SFQ voltage pulses and exhibit natural spiking dynamics analogous to biological neurons when biased appropriately. SFQ circuits have been proposed and demonstrated as neuromorphic hardware because they can operate at very high clock rates (tens to hundreds of GHz) with extremely low energy per pulse, orders of magnitude below typical CMOS spiking implementations [104], [105], [106].

### B. MAPPING NEURAL PRIMITIVES

The mapping of neural functions in these substrates shifts from static weights to dynamic synchronization and pulse interactions.

#### 1) WEIGHTS AND SYNCHRONIZATION

In spintronics, synaptic weights are often encoded in the coupling strength between oscillating STNOs. Computation emerges from the synchronization (phase-locking) of these oscillators, mimicking the temporal binding hypothesis of neuroscience [99]. In skyrmion reservoirs, the “weight” is implicit in the complex dynamical response of the magnetic textures; input signals perturb the skyrmion fabric, and the resulting high-dimensional non-linear modes encode the information [102], [107].

#### 2) ACTIVATION AND SPIKING

Superconducting Josephson Junctions provide the most physically faithful realization of a spiking neuron. A JJ integrates incoming flux pulses (current) until a critical current  $I_c$  is exceeded. At this threshold, the junction switches, releasing exactly one flux quantum (a voltage spike) and resetting. This “Integrate-and-Fire” dynamics is intrinsic to the device physics and does not require complex circuitry to emulate [105]. Similarly, spintronic devices utilize the non-linear threshold of magnetic precession to implement activation functions.

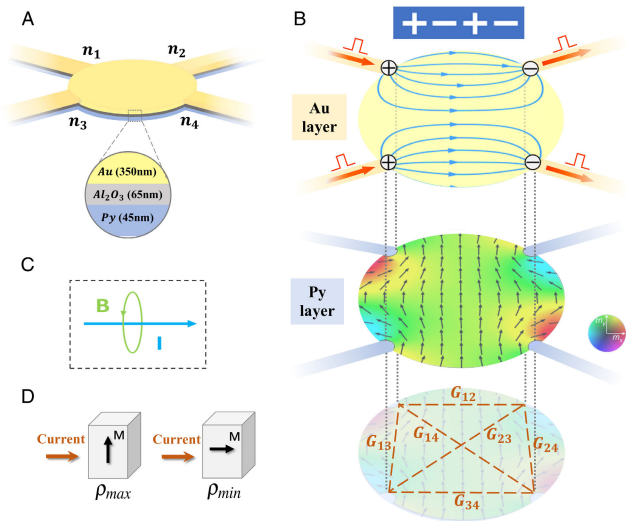
#### 3) PLASTICITY AND SELF-ORGANIZATION

A major breakthrough in 2025 was the demonstration of intrinsic learning in magnetic materials. Niu et al. showed that magnetic textures can minimize their internal free energy in response to clamped boundary conditions. This physical relaxation process is mathematically isomorphic to the update rules of a Hopfield network, allowing the material to “learn” patterns via intrinsic gradient descent without external digital calculation [108]. The basic principle is shown with a sample 4-node Hopfield network in Figure 6.

### C. ARCHITECTURES AND I/O INTERFACES

Architectures in this domain favor reservoir computing and event-driven spiking networks over static feedforward layers.

Pinna et al. proposed that disordered (random) skyrmion textures can serve as a high-dimensional physical reservoir. Inputs (e.g., current- or field-induced perturbations) locally distort the magnetic texture and excite non-linear collective



**FIGURE 6.** Working principle of a 4-node self-learning magnetic Hopfield neural network (HNN). (A) Schematic of the experimental device structure. (B) Voltage pulses applied to the top Au layer encode the input or boundary conditions and generate current distributions that configure the spin orientation in the Py layer via the Oersted field (C). The resulting magnetic texture relaxation defines the effective conductance matrix, mimicking self-adjusting synaptic weights. (D) The anisotropic magnetoresistance (AMR) effect is used to read out the evolving network state through the angle-dependent relation between current and magnetization. The figure illustrates how associative memory and gradient-descent-like learning emerge directly from ferromagnetic energy minimization. Reprinted from [108], licensed under CC-BY-NC-ND 4.0.

dynamics. The resulting interactions between skyrmions and the surrounding spin background (e.g., repulsion, deformation, mode coupling) embed the input history into a rich spatiotemporal state that can be linearly read out, for instance via electrical transduction such as magnetoresistive signals [102]. This approach does not rely on explicitly training the substrate itself; instead, it exploits the intrinsic complexity of the skyrmion texture and trains only the readout layer.

Superconducting architectures arrange Josephson junctions into large-scale spiking neural networks (SNNs) operating with Single Flux Quantum (SFQ) pulses. Signals propagate as ballistic quantized flux pulses along superconducting transmission lines with negligible resistive loss and minimal dispersion. Such networks are particularly suited for processing ultrafast temporal data streams (e.g., from RF or cryogenic sensors) in real time, enabling operations such as coincidence detection and temporal pattern matching at clock rates far beyond conventional CMOS implementations [106].

#### D. TRAINING AND DESIGN PARADIGMS

Training methods differ significantly between the two pillars of this chapter. At a high level, magnetic systems are increasingly explored in the in-materio regime, where optimization is partly delegated to the intrinsic energy landscape of the substrate, whereas superconducting neural systems more often rely on local in-situ update rules or hybrid training schemes.

For magnetic systems, the “In-Materio” learning paradigm is gaining traction. By mapping the loss function of a neural network to the Hamiltonian (energy function) of the ferromagnet, the training process becomes a physical relaxation to the ground state. This physics-driven strategy is attractive because it exploits the natural relaxation dynamics of the substrate and can partially replace a fully external optimizer loop. However, its practical effectiveness depends on how faithfully the physical energy landscape represents the computational objective and on whether the relaxation reliably reaches useful minima [108].

In superconducting SNNs, training often utilizes Spike-Timing-Dependent Plasticity (STDP). This can be implemented using variable inductive loops (SQIFs) or by hybridizing with magnetic Josephson junctions, where the magnetic state (and thus the critical current) is modified by the timing of arriving flux quanta [99]. Such local update rules are well matched to event-driven superconducting hardware because they directly exploit spike timing and device-level dynamics, but they currently offer less flexibility than large-scale gradient-based digital optimization and remain constrained by hardware complexity, tunability, and reproducibility.

#### E. REPRESENTATIVE DEMONSTRATIONS

Representative demonstrations in this class span three characteristic regimes: magnetic reservoir computing, ultrafast superconducting spiking networks, and intrinsic self-learning in magnetic energy landscapes. Pinna et al. investigated disordered (random) skyrmion textures as physical reservoirs and showed that their intrinsic non-linear dynamics generate high-dimensional spatiotemporal responses suitable for temporal signal processing tasks, supporting the use of magnetic textures as substrates for reservoir computing [102]. In the superconducting domain, Schneider et al. implemented a spiking neural network based on Josephson junctions and Single Flux Quantum (SFQ) signaling, demonstrating operation at multi-tens of gigahertz clock rates with attojoule-scale switching energies [106]. This work provides a concrete example of how superconducting device physics can be exploited for high-speed, energy-efficient neuromorphic architectures. Finally, Niu et al. realized a magnetic Hopfield network in which synaptic adaptation is governed by the intrinsic energy landscape of a ferromagnetic system, thereby demonstrating associative memory and gradient-descent-like learning directly at the device level [108]. Figure 6 illustrates this principle and shows how boundary-conditioned magnetic texture relaxation defines the effective conductance matrix and can be read out electrically via anisotropic magnetoresistance.

#### F. ENGINEERING CONSTRAINTS

The impressive speed and efficiency of these substrates come with distinct physical barriers. In spintronics, thermal fluctuations can destabilize small magnetic textures (superparamagnetism), imposing a trade-off between device miniaturization

and information retention, although stochastic resonance can also be exploited for probabilistic computing [100].

In the superconducting domain, the “cryogenic penalty” remains a major hurdle, as circuits require cooling to temperatures around 4 K. While the switching energy of SFQ pulses is in the attojoule range and spiking networks can operate at clock rates of tens to hundreds of gigahertz [106], the required cryogenic infrastructure (cryocoolers, thermal shielding, bias distribution, and room-temperature I/O electronics) dominates the overall energy budget. System-level analyses show that the cooling and support electronics typically consume from tens of watts up to kilowatts of electrical power, depending on system scale and cooling technology, overwhelming the energy dissipated by the Josephson junctions themselves [109].

As a result, superconducting neuromorphic hardware is currently viable primarily for large-scale high-performance computing installations or specialized environments such as quantum-computer control systems, rather than for mobile or embedded devices. In addition, impedance and voltage-level mismatches between superconducting SFQ circuits (millivolt-level signals) and room-temperature CMOS logic (volt-level signals) make interfacing complex and energetically costly, favoring architectures in which superconducting processors operate as cryogenic “islands” connected to conventional electronics via dedicated interface circuits [109].

### G. POSITION WITHIN PHYSICAL NEURAL COMPUTING

Spintronic and superconducting systems occupy the “high-performance” corner of the physical computing landscape. They offer a stark contrast to the slow, diffusive nature of wetware and the static nature of resistive crossbars. Spintronics offers a path to non-volatile, radiation-hard, and oscillatory computing at room temperature. Superconducting logic, conversely, represents the ultimate physical limit of speed and efficiency, theoretically capable of matching the connectivity of the human brain with the speed of a supercomputer, provided the cryogenic barrier is accepted. Together, they constitute the magnetic and quantum-electronic frontier of PNNs.

## VIII. PHOTONIC NEURAL COMPUTING

Photonic computing replaces electrons with photons as the information carrier, leveraging the intrinsic properties of light—ultra-high bandwidth, massive parallelism via wavelength division multiplexing (WDM), and propagation at the speed of light—to perform computation. While electronic processors are throttled by capacitive charging and interconnect latency, photonic processors operate in the “flight-time” regime, where the latency is determined solely by the refractive index. This makes photonics the premier candidate for accelerating linear operations at speeds and energy efficiencies unattainable by CMOS [110], [111].

### A. THE SUBSTRATE PRINCIPLE

The fundamental physical mechanism exploited is wave interference. Optical propagation naturally implements linear transformations; when coherent modes overlap, their amplitudes sum constructively or destructively based on phase. This allows matrix-vector multiplication (which arguably is the computational bottleneck of neural networks) to be performed passively as light traverses a medium. Information is encoded in amplitude, phase, polarization, or spatial mode. Unlike the charge-accumulation logic of electronics, the computation here is a transmission problem: the result is the scattering of the input field by the optical elements [112], [113]. However, since photons lack charge and interact only weakly (making “optical transistors” difficult), practical systems often adopt hybrid architectures: linear operations occur optically, while non-linear activations and control logic remain electronic [114], [115].

### B. MAPPING NEURAL PRIMITIVES

Mapping neural networks to photonics requires translating linear algebra into optical components while addressing the inherent difficulty of implementing non-linearity.

#### 1) WEIGHTS AND SUMMATION

In coherent architectures, a synaptic weight is physically represented by the transmission coefficient or phase shift of a component. Vector-Matrix Multiplication (VMM) is performed as light propagates through a programmable mesh of Mach-Zehnder Interferometers (MZIs) or diffractive surfaces. The summation is inherent to the physics of superposition at the detector or combiner ports [113], [116].

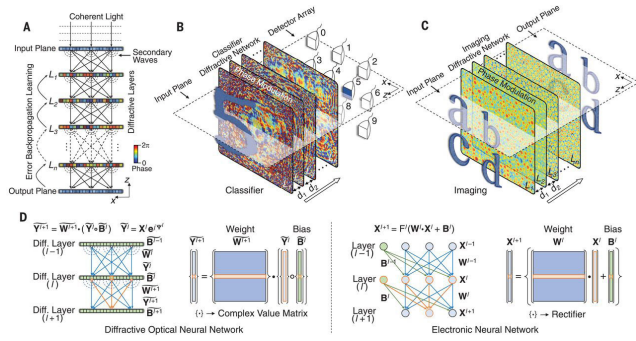
#### 2) ACTIVATION AND NONLINEARITY

Implementing non-linearity remains the primary engineering hurdle. While some approaches explore all-optical nonlinearities using saturable absorbers or bistable resonators, these often require prohibitively high power densities. The dominant paradigm relies on Opto-Electronic (O-E-O) conversion: the optical signal is detected, thresholded electronically, and re-modulated onto a carrier. This introduces a bandwidth bottleneck but ensures stable operation [114], [117].

### C. ARCHITECTURES AND I/O INTERFACES

Photonic architectures generally fall into three categories, trading off programmability against throughput.

Programmable photonic integrated circuits (PICs) utilize meshes of MZIs to implement arbitrary unitary matrices, following the triangular decomposition scheme originally proposed by Reck et al. [118]. In these “Optical TPUs,” thermo-optic or electro-optic phase shifters act as tunable weights. A landmark theoretical framework by Miller proved that such meshes can self-configure to perform any linear transformation, a concept later validated experimentally for high-speed inference [112], [113].



**FIGURE 7. Diffractive deep neural networks (D<sup>2</sup>NNs).**(A) A D<sup>2</sup>NN uses multiple trainable diffractive layers with complex-valued transmission or reflection coefficients to implement an optical function between input and output planes. Once the layer parameters have been optimized and the structure has been fabricated, inference is performed passively through wave propagation at the speed of light. (B) Example of a D<sup>2</sup>NN for handwritten digit classification. (C) D<sup>2</sup>NN functioning as an imaging lens. (D) Comparison with a conventional neural network, highlighting coherent complex-valued inputs and diffraction-based layer interconnections. The figure illustrates the key trade-off of this photonic regime: extremely fast and energy-efficient inference at the cost of limited post-fabrication programmability. Reprinted from [119] with permission from the American Association for the Advancement of Science.

Diffractive architectures operate in free space. A series of passive transmissive layers (engineered metasurfaces) modulate the phase and amplitude of an incoming wavefront. As the wave propagates through these layers, the diffraction pattern physically computes the inference. These systems operate at the speed of light with minimal energy consumption, but act as fixed, read-only classifiers once fabricated [119]. The approach is visualized in Figure 7.

For temporal processing, photonic reservoirs exploit the rich transient dynamics of optical feedback. Architectures using semiconductor lasers with delayed feedback or fiber-optic loops generate high-dimensional state spaces from time-series inputs. These systems excel at chaotic signal prediction and speech recognition, requiring training only at the electronic readout layer [120], [121].

### D. TRAINING AND DESIGN PARADIGMS

Training photonic circuits presents challenges due to phase noise and the “analog gap” between simulation and reality. Accordingly, photonic neural hardware is currently dominated by ex-situ digital training with calibrated deployment, while more recent work increasingly explores in situ gradient measurement directly on the physical chip.

The standard approach involves training a model digitally and mapping the weights to physical phase shifts. However, fabrication tolerances and thermal crosstalk require rigorous calibration routines to align the physical device with the digital model [112]. Algorithms specifically designed to robustly configure MZI meshes in the presence of imperfections are essential for scaling [113]. This paradigm is attractive because it leverages mature digital optimization while

preserving the high-speed, low-latency inference capabilities of photonic hardware. On the other hand, it remains limited by calibration overhead and residual mismatch between simulated and fabricated devices.

More recent advances utilize *in situ* training, where gradients are measured directly on the hardware. Techniques such as the “adjoint variable method” physically backpropagate light through the output ports to measure sensitivity, automatically accounting for the exact physical state of the chip, including defects [122]. Such approaches are particularly promising for reducing the analog gap and compensating for fabrication defects, although they still require stable optical control and precise measurement infrastructure to remain practical at scale.

### E. REPRESENTATIVE DEMONSTRATIONS

Representative demonstrations of photonic neural computing span three characteristic regimes: programmable coherent matrix processors, passive diffractive inference, and dynamical reservoir computing. Shen et al. demonstrated a programmable nanophotonic processor based on a mesh of Mach–Zehnder interferometers that performed vowel recognition, providing one of the first convincing experimental validations that coherent optical interference can implement trainable matrix operations for neural inference [112]. In the diffractive domain, Lin et al. realized a deep optical classifier using multiple 3D-printed diffractive layers that classified handwritten digits purely through wave propagation after fabrication, showing that a trained optical structure can perform passive inference at the speed of light with essentially no active computation during operation [119]. Figure 7 illustrates this principle and highlights both the layer-wise diffraction-based transformation and the fixed, read-only nature of the fabricated system. For temporal tasks, Appeltant et al. and Paquot et al. demonstrated delayed-feedback photonic reservoirs for speech recognition and chaotic time-series prediction, showing that optical dynamics can also support recurrent computation with fading memory rather than only static linear transforms [120], [121].

### F. ENGINEERING CONSTRAINTS

The adoption of photonic computing is constrained by scale and sensitivity. First, optical components are orders of magnitude larger than transistors (constrained by the wavelength of light), limiting the number of neurons per chip and thus the integration density. Second, analog optical computing is limited by laser noise, detector shot noise, and thermal drift, typically capping effective precision at 4-8 bits [114]. Third, the transduction tax has to be taken into account again: the energy and latency cost of converting data between the electrical and optical domains can negate the benefits of optical processing unless the computational depth is sufficiently large [110].

### G. POSITION WITHIN PHYSICAL NEURAL COMPUTING

Photonic neural computing occupies the “high-speed, low-latency” corner of the physical computing landscape. Unlike chemical or mechanical systems that excel at slow, adaptive dynamics, photonics excels at performing massive linear algebra operations at the bandwidth limit of the electromagnetic spectrum. Ideally suited for ultrafast signal processing and server-side acceleration, it complements substrates like memristors or wetware that offer higher density or plasticity but lower speed [114], [115].

### IX. MECHANICAL METAMATERIALS

Mechanical neural computing revisits mechanics, which is arguably the oldest substrate of calculation, through the modern lens of metamaterials. Unlike the rigid gears of the Babbage engine, contemporary mechanical computing relies on compliant mechanisms and the non-linear deformation of continuous media. By engineering the topology of lattice structures, information is processed not by the flow of electrons, but by the propagation of strain energy and the interactions of mechanical instabilities [123]. This substrate is particularly relevant for “Embodied Intelligence” and morphological computation, where the computation is physically integrated into the structural load-bearing elements of a robotic system [124], [125], [126], [127].

#### A. THE SUBSTRATE PRINCIPLE

The fundamental computational resource in mechanical systems is the elastic potential energy landscape. While linear elastic materials obey Hooke’s Law (force proportional to displacement), neural computation requires non-linearity. This is achieved through geometric instabilities, most notably “snap-through buckling.” A bistable buckled beam possesses two energy minima separated by an energy barrier. Transitioning between these states allows for binary storage (memory) and thresholding behavior [128]. In metamaterials, thousands of such unit cells are coupled in a lattice; the global deformation of the material is thus the aggregate result of local interactions, functioning effectively as a massive network of mechanically coupled non-linear springs [129].

#### B. MAPPING NEURAL PRIMITIVES

To act as a neural network, the mechanical substrate must exhibit weighted connections, non-linear activation functions, and mechanisms for plasticity.

##### 1) STIFFNESS AS WEIGHT

In a mechanical network, the “signal” is typically a force or displacement vector. The synaptic weight corresponds to the local stiffness ( $k$ ) of the beams connecting nodes. A stiffer beam transmits force more effectively, effectively amplifying the signal, while a compliant beam attenuates it. In programmable metamaterials, these weights can be adjusted by altering the beam thickness during fabrication

or, in active systems, by modulating the stiffness using piezoelectric or magnetic actuators [123], [129].

##### 2) INSTABILITY AS ACTIVATION

The non-linear activation function is realized through buckling phenomena and topological modes. A slender beam under compression initially acts linearly. However, once a critical load (Euler load) is exceeded, it snaps to a new configuration. This sudden change in stiffness acts as a physical sigmoid or Heaviside function, allowing the system to perform decision-making and logic operations purely through geometric deformation [130], [131].

##### 3) PLASTICITY AND MEMORY

While many architectures are static (read-only), adaptive materials enable in-materio learning. Memory can be encoded in the reconfigurable geometry of multistable elements and via time-dependent relaxation or aging processes. Recent work demonstrates supervised learning through physical parameter updates in creased sheets, and reveals how viscoelastic relaxation and hysteretic rearrangements in disordered (amorphous) media generate history-dependent responses and memory [132], [133], [134], [135].

### C. ARCHITECTURES AND I/O INTERFACES

Architectures in mechanical computing are defined by their unit cell topology and dimensionality.

2D and 3D lattices of interconnected beams form the bulk of mechanical neural networks. These structures function as “physical perceptrons,” where an input force pattern applied to one boundary propagates through the lattice to produce a target displacement at an output boundary [129].


Alternatively, plate-based systems utilize origami (folding) and kirigami (cutting) principles. Here, the “neurons” are the intersection vertices of folds. The non-linear rotational stiffness of the creases allows these sheets to function as reconfigurable logic gates and multi-input classifiers with high packing density [132], [136], [137].

Third, soft bodies and elastic solids with complex internal dynamics can be used as physical reservoirs. Similar to their fluidic counterparts, soft silicone arms or gels exploit their infinite degrees of freedom to project low-dimensional inputs (e.g., motor commands) into a high-dimensional dynamic state space, enabling temporal pattern recognition for soft robotic control [138].

### D. TRAINING AND DESIGN PARADIGMS

Training a mechanical material implies finding a configuration of stiffnesses or geometries that yields a desired force-displacement response. Accordingly, mechanical substrates are dominated by ex-situ inverse design, using either gradient-based differentiable mechanics or closed-box evolutionary optimization, whereas direct in-materio adaptation remains comparatively rare.

The classical engineering approach formulates the desired input-output mapping as a constraint optimization problem. Algorithms such as SIMP (Solid Isotropic Material with Penalization) optimize the density distribution of material within a design domain to satisfy these functional constraints [139]. This paradigm is well suited to mechanical systems because geometry and constitutive behavior can often be modeled explicitly, but it is limited by modeling assumptions and by the difficulty of capturing fabrication tolerances and material nonlinearities with high fidelity.



**FIGURE 8.** Mechanical neural network (MNN) fabricated for the experimental study of Lee et al. [123]. (A) Full-view photograph of the lattice-based mechanical network. (B) Close-up of the beam structure whose local stiffness values act as trainable mechanical weights. The figure illustrates the physical realization of a mechanical perceptron-like architecture in which desired behaviors, such as shape morphing, are achieved by optimizing the geometry and stiffness distribution of the lattice. Reprinted from [123] with permission from the American association for the advancement of science.

More recent *differentiable mechanics* approaches embed differentiable simulators or differentiable surrogate models into an optimization loop, enabling gradient-based inverse design of geometries and effective stiffness parameters [123], [140], [141]. Alternatively, evolutionary algorithms can mutate geometry representations under task-level fitness functions, which is attractive when accurate gradients are unavailable, or the design space is highly discrete [142], [143]. Gradient-based methods are typically more sample-efficient when accurate simulators are available, whereas evolutionary strategies are more robust to discontinuous design spaces and non-differentiable objectives, albeit usually at higher search cost.

### E. REPRESENTATIVE DEMONSTRATIONS

Representative demonstrations of mechanical neural computation span three characteristic regimes: trainable lattice-based networks, in-situ gradient-based learning, and geometry-driven logic and classification. A major milestone was the experimental realization of mechanical neural networks (MNNs) by Lee et al. They fabricated a lattice of tunable beams whose stiffnesses are adjusted analogously to ANN weights, enabling the structure to learn multiple mechanical behaviors (e.g., shape morphing) through repeated training/optimization and experimental validation [123]. Figure 8 shows the fabricated network and illustrates the physical realization of this trainable mechanical architecture.

Building on this, Li and Mao demonstrated that such networks can also be trained for regression and classification using in situ backpropagation, extracting gradient information directly from the mechanical system's physical response without relying on a digital twin [144].

A complementary line of work exploits geometry and multistability more directly. Coulais et al. realized combinatorial logic gates using bistable mechanical units arranged in planar lattices [129]. More broadly, origami- and kirigami-based metamaterials demonstrate that nonlinear classifier-like input-output mappings can emerge directly from geometry, folding mechanics, and physical adaptation rather than from electronically controlled processing [132], [137].

### F. ENGINEERING CONSTRAINTS

Mechanical computing faces limits imposed by material science and wave physics. First, signals travel at the speed of sound within the material (typically km/s), which is orders of magnitude slower than electromagnetic signals. This limits mechanical systems to low-frequency control tasks (Hz to kHz range). Second, repeated cycling of bistable elements can lead to plastic deformation or fatigue failure, altering the computational properties (weights) over time. Third, in viscoelastic materials, energy is dissipated as heat (damping). This limits the depth of the network, as the signal-to-noise ratio degrades with distance from the input source unless active amplification is introduced.

### G. POSITION WITHIN PHYSICAL NEURAL COMPUTING

Mechanical metamaterials occupy the niche of “structural intelligence.” They are unlikely to compete with electronic or photonic substrates in speed or accuracy, but they uniquely enable computation without electricity and intrinsic coupling to physical environments. By embedding intelligence directly into load-bearing structures, they represent a paradigm complementary to molecular or electronic approaches, ideal for applications requiring extreme robustness and morphological adaptation [124].

### X. MICROFLUIDIC NEURAL NETWORKS (PRESSURE-DRIVEN)

Microfluidic systems implement computation by routing liquids or gases through top-down engineered networks of micrometer-scale channels. Unlike chemical or DNA-based substrates, where the computational structure emerges bottom-up from molecular interactions, microfluidic neural systems constitute rigid devices whose topology is defined by lithographically fabricated geometries and boundary conditions [145]. Consequently, they function as deterministic hydraulic integrated circuits that process information through the controlled transport of mass rather than charge, enabling the direct manipulation of physical matter [146]. As detailed in a recent review by Law et al., this field is rapidly expanding towards fully fluidic neuromorphic architectures that integrate flow-based logic with fluidic memristive elements [147].

### A. THE SUBSTRATE PRINCIPLE

Microfluidic neural computing operates strictly in the low Reynolds number regime ( $Re \ll 1$ ), where viscous forces dominate over inertia. In this domain, fluid motion is governed by the Stokes equations, resulting in flow that is laminar, deterministic, and time-reversible in the absence of active switching elements [148]. This physical regime permits a rigorous analogy between hydraulic and electronic transport, enabling the systematic design of fluidic logic using established circuit theory [149].

In this analogy, long, narrow channels act as hydraulic resistors, inducing pressure drops linearly proportional to the volumetric flow rate. Elastic chambers function as fluidic capacitors, storing volume under pressure and introducing temporal integration constants. Asymmetric channel structures, such as Tesla valves, rectify flow similar to diodes, while pressure-actuated flexible membranes implement active gating, functionally replicating the behavior of transistors [150]. By composing these fundamental elements, microfluidic substrates can realize complex analog signal processing architectures completely within the fluidic domain.

### B. MAPPING NEURAL PRIMITIVES

To construct neural networks, the governing fluid dynamic equations must be mapped to the mathematical primitives of weighted summation, non-linear activation, and state retention.

#### 1) SYNAPTIC WEIGHTS AND SUMMATION

Synaptic weighting is physically encoded in the hydraulic resistance of the channel network. According to the Hagen-Poiseuille law, the resistance scales linearly with channel length and inversely with the fourth power of the hydraulic diameter. Therefore, weights can be statically defined during fabrication by optimizing channel geometries to restrict signal flow [151]. The summation operation emerges naturally from the conservation of mass at channel junctions, where merging fluid streams result in a combined pressure or flow rate proportional to the weighted inputs.

#### 2) ACTIVATION AND NONLINEARITY

While low-Reynolds-number flow is inherently linear, neural computation requires non-linearity. This is introduced through fluid-structure interactions or multiphase dynamics. Pressure-gain valves and flexible membranes exhibit thresholding behaviors where flow is blocked until a critical pressure is exceeded, effectively implementing rectified linear unit (ReLU) or sigmoidal activation functions [150]. Alternatively, “Bubble Logic” utilizes the non-linear interaction of immiscible fluids, where the presence or absence of a droplet in a channel acts as a discrete state variable capable of performing Boolean logic operations [146].

#### 3) MEMORY MECHANISMS

Temporal dynamics and memory are realized through fluidic capacitance and inertia. Elastic deformation of channel walls allows for short-term storage of signal state (integration). For long-term memory, systems utilize bistable mechanisms such as trapped droplets, hysteretic valve switching, or high-viscosity non-Newtonian fluids that retain their state in the absence of continuous power, functioning as fluidic flip-flops [152], [153].

### C. ARCHITECTURES AND I/O INTERFACES

Microfluidic architectures typically follow one of two paradigms, dictated by the trade-off between control complexity and dynamic richness.

Layered channel networks implement fixed-weight mappings where the computation is defined by the circuit topology. These systems function as analog feedforward networks for classification or regression tasks, or as digital logic circuits for Boolean operations [154]. While interpretable, they often require complex multilayer fabrication to handle signal routing (crossovers).

To bypass the need for precise component-level design, recent approaches utilize reservoir computing. Here, complex channel geometries, recirculation zones, or multiphase flows generate rich, high-dimensional transient dynamics. The fluidic body acts as a physical kernel, projecting inputs into a state space that is linearised by a simple readout layer. This approach is particularly robust against fabrication tolerances and has been successfully applied to temporal pattern recognition and soft robotic control [28], [138].

The interface between the macro- and microworld represents a significant engineering bottleneck. Inputs are encoded as continuous pressure levels, volumetric flow rates, or discrete droplet sequences, typically requiring external solenoid valves or syringe pumps. Readout is performed via integrated pressure transducers, flow sensors, or optical imaging of dye concentrations, often imposing the transduction tax discussed above that limits the overall system bandwidth [155].

### D. TRAINING AND DESIGN PARADIGMS

Unlike electronic neural networks, microfluidic systems generally do not provide native mechanisms for rapid, repeated in-situ parameter updates during operation. Consequently, training strategies have historically relied on in-silico optimization. Accordingly, microfluidic substrates are dominated by ex-situ digital-twin design, complemented by gradient-based inverse design and, where necessary, closed-box evolutionary optimization. Direct in-materio adaptation remains limited because channel geometries are typically fixed once fabricated and fluidic dynamics evolve on relatively slow timescales.

The dominant paradigm is offline training, where a differentiable model of the fluid dynamics (a digital twin) is trained digitally. The resulting optimal geometries (channel widths, lengths) are then transferred to the physical substrate via

lithography. This approach effectively treats the microfluidic chip as an inference-only ASIC. This strategy is attractive because fluid transport is often sufficiently structured to support simulation-driven optimization, but it remains limited by fabrication tolerances, model mismatch, and the absence of rapid post-fabrication reconfiguration.

Recent advances have integrated finite-element fluid solvers directly into automatic differentiation pipelines, enabling gradient-based optimization of channel geometries for complex flow control tasks [156]. Complementary approaches use evolutionary algorithms or machine learning to search the topological design space for networks that exhibit desired behaviors, such as oscillation or gait generation, without requiring an explicit gradient [157]. Gradient-based methods are typically more sample-efficient when accurate fluid models are available, whereas evolutionary search is better suited to highly discrete or topological design spaces, albeit usually at substantially higher search cost.

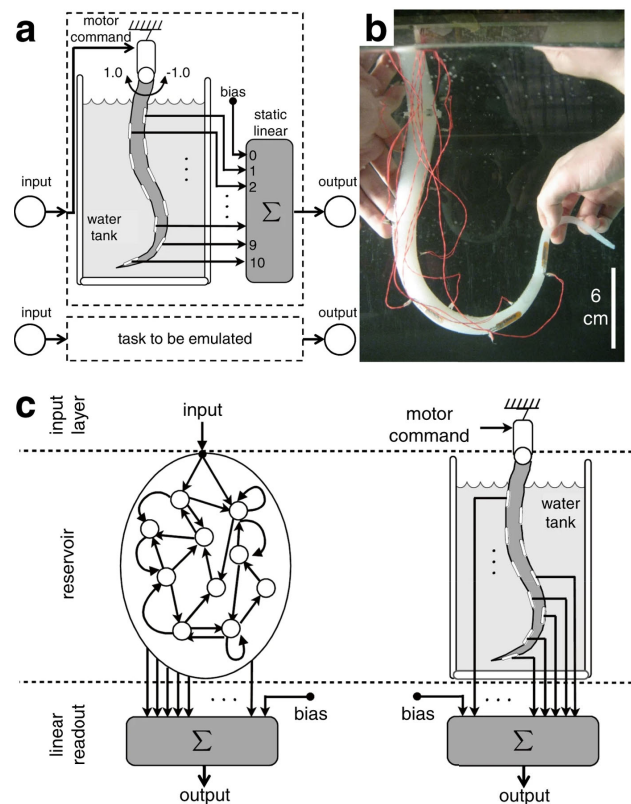
### E. REPRESENTATIVE DEMONSTRATIONS

Representative demonstrations of pressure-driven microfluidic neural computation span three characteristic regimes: digital fluidic logic, embodied reservoir computing, and fully fluidic control for autonomous soft robotics.

Prakash and Gershenfeld demonstrated that bubbles in a microfluidic channel can interact non-linearly to perform universal Boolean logic (AND, OR, NOT) and signal synchronization. This established the field of digital microfluidic computing, where the bit is a physical packet of matter [146]. This established a fluidic computing paradigm in which the bit is not an electrical state but a physical packet of matter moving through a designed hydraulic circuit.

A complementary regime is reservoir-style computation in compliant fluidic bodies. Nakajima et al. showed that the nonlinear dynamics of a soft silicone arm, as well as related fluidic networks, can be exploited as physical reservoirs for computation [28], [138]. Motor commands perturb the body, embedded bend sensors record the resulting transient dynamics, and only the readout layer is trained, enabling tasks such as chaotic time-series prediction and autonomous gait generation. Figure 9 illustrates this principle and highlights how the soft body itself replaces the abstract reservoir layer of a conventional Echo State Network.

A third line of work demonstrates fully fluidic embodied autonomy. Wehner et al. integrated microfluidic logic into the “Octobot”, a soft robot capable of self-powered movement without semiconductor electronics [158]. Building on such principles, Preston et al. realized pneumatic logic gates and ring oscillators that were embedded directly into soft robotic bodies to achieve closed-loop locomotion and obstacle navigation without electronic processors [153]. These systems show that fluidic substrates can implement not only isolated logic operations or reservoir dynamics, but also autonomous control loops in which sensing, computation, and actuation are tightly integrated in the material system itself.



**FIGURE 9.** Physical reservoir computing using a soft silicone arm. (a) Information-processing loop: motor commands actuate the arm, while embedded bend sensors capture the resulting body dynamics as reservoir states. (b) Experimental prototype of the soft arm. (c) Conceptual comparison showing how the physical soft body replaces the abstract reservoir layer of a conventional Echo State Network. The figure illustrates the central idea of embodied reservoir computing: computation emerges from the intrinsic transient dynamics of the compliant physical system, while only the readout is trained. Reprinted from [28], licensed under CC-BY 4.0.

### F. ENGINEERING CONSTRAINTS

The practical deployment of microfluidic neural networks is defined by a trade-off between integration and speed. First, due to viscosity and compressibility, signal propagation speeds are limited to the speed of sound in the fluid or slower. This restricts operation frequencies to the Hz to kHz range, making them unsuitable for high-speed data processing but ideal for biological timescales. Second, integration density is limited. While individual channels are small, the control infrastructure (valves, pumps) scales poorly. Large-scale integration requires multilayer soft lithography, which increases fabrication complexity. Unlike electronics, however, microfluidics is inherently robust to electromagnetic interference (EMI) and radiation, making it attractive for extreme environments [153].

### G. POSITION WITHIN PHYSICAL NEURAL COMPUTING

Microfluidic neural systems occupy an intermediate regime between chemical computing and solid-state hardware. They offer richer dynamics and stronger embodiment than

electronic substrates, while remaining deterministic and engineerable compared to biochemical systems. Their primary relevance lies in applications requiring tight integration of sensing, computation, and actuation in fluidic or biological environments, such as “lab-on-a-chip” diagnostics or soft robotics, rather than in high-speed inference [145].

## XI. IONTRONICS AND NANOFUIDIC NEURAL COMPUTING

While microfluidics mimics hydraulic machinery, iontronics mimics the nervous system. By transitioning from the microscale to the nanoscale ( $< 100$  nm), surface-to-volume ratios increase drastically, entering a regime where fluid transport is dominated by surface charges rather than bulk continuum mechanics. The information carrier shifts from the fluid mass to the specific ions ( $\text{Na}^+$ ,  $\text{K}^+$ ,  $\text{Cl}^-$ ) dissolved within it [159]. This substrate provides the technological bridge between solid-state electronics and biological wetware, enabling functional “artificial axons” and ionic transistors that operate using the brain’s native signaling medium [160].

### A. THE SUBSTRATE PRINCIPLE

The computational engine of iontronics is the Electrical Double Layer (EDL). When an electrolyte contacts a charged surface, counter-ions accumulate to screen the surface charge, forming a layer with a characteristic thickness known as the Debye length ( $\lambda_D$ ). In channels comparable in size to the Debye length ( $h \approx 2\lambda_D$ ), the EDLs from opposing walls overlap, rendering the entire channel unipolar [161].

Crucially, transport at this scale defies classical hydrodynamics. Kavokine et al. demonstrated that in carbon nanotubes, water flows with vanishingly low friction due to quantum mechanical coupling between fluid fluctuations and the electron gas of the wall. This “quantum friction” effect enables transport rates orders of magnitude higher than continuum predictions, providing the high-speed ionic throughput necessary for neuromorphic spiking [162].

### B. MAPPING NEURAL PRIMITIVES

Iontronics maps neural functions to the electro-kinetic equations (Poisson-Nernst-Planck) governing ion flux.

#### 1) SYNAPTIC WEIGHTS AS IONIC CONDUCTANCE

In an iontronic circuit, the synaptic weight  $w_{ij}$  is physically instantiated as the ionic conductance  $G$  of a nanopore. According to Ohm’s law for electrolytes ( $I = G \cdot V$ ), the signal is the product of conductance and voltage. Unlike in bulk fluids, this conductance is highly nonlinear and history-dependent. Recent work by Robin et al. demonstrated “Ion-Shuttling Memristors” using 2D nanofluidic channels. Here, the reversible intercalation of ions into graphite layers or the formation of ion clusters creates a hysteretic conductance response, enabling non-volatile analog weight storage similar to Long-Term Potentiation (LTP) in biological synapses [163]. Extending this capability, Ismail et al.

recently demonstrated fully programmable memristors based on two-dimensional nanofluidic channels, proving that stable analog weights can be precisely tuned in aqueous environments [164].

#### 2) ACTIVATION AND RECTIFICATION

The non-linear activation function required for deep learning is realized through geometric asymmetry. Conical nanopores act as Ion Current Rectifiers (ICR), where the interplay between electro-osmotic flow and the EDL creates a diode-like current-voltage characteristic. This provides a physical implementation of the Rectified Linear Unit (ReLU), permitting current flow in only one direction above a threshold voltage [161].

### C. ARCHITECTURES AND I/O INTERFACES

A key goal of neuromorphic engineering is the physical realization of the “Integrate-and-Fire” neuron. Robin, Kavokine, and Bocquet demonstrated that this behavior emerges naturally in nanofluidic slits. Their system couples two distinct transport channels, one as a model for the fast sodium and the other for the slow potassium channels of a biological axon, through an electronic feedback loop. Under a constant voltage input, ions accumulate (the “integration”) until a threshold is reached, triggering a negative differential resistance event that discharges the potential in a sharp spike (the “firing”). This device produces self-sustained trains of action potentials structurally identical to the Hodgkin-Huxley model, proving that spiking dynamics are an emergent property of confined ion transport [160].

The primary architectural advantage of iontronics lies in its interface. Since the computation uses the same ions ( $\text{Na}^+$ ,  $\text{K}^+$ ) as living cells, iontronic devices can communicate directly with biological tissue. This avoids the need for transduction layers required by electronic devices and minimizes the risk of toxic redox reactions, paving the way for “smart prosthetics” where the computational hardware is chemically indistinguishable from the biological neural network it regulates [159].

### D. TRAINING AND DESIGN PARADIGMS

Training and design in iontronic neural systems follow the general paradigms of physical learning, but are shaped by the fact that “weights” and state variables emerge from coupled electrokinetics rather than from explicitly addressable parameters. As consolidated in recent review literature, effective design therefore requires *physico-chemical co-design*: geometry (e.g., slit height/length, tapering), interfacial properties (surface charge, functionalization), and electrolyte conditions (ionic strength, pH, valence) jointly set electric-double-layer overlap, selectivity, and concentration polarization, which in turn determine nonlinearity and the volatility/retention trade-off [159], [161], [165]. Accordingly, iontronic systems instantiate all three major categories introduced in Section II-B, namely ex-situ digital-twin design,

hybrid digital–physical training, and in-materio adaptation. In practice, three strategies dominate. (i) *In-silico / digital-twin workflows* use reduced transport models to guide geometry and operating points and then rely on empirical calibration to compensate fabrication tolerances and surface-chemistry drift [159], [165]. (ii) *Hybrid training via reservoir computing* exploits transient ionic dynamics as a fixed high-dimensional projection, while only a simple readout is trained digitally; this is particularly attractive because it is robust to device variability and avoids requiring accurate gradients inside wet, nanoscale hardware [166], [167]. (iii) *In-materio adaptation* leverages history-dependent conductance arising from ionic rearrangements and interfacial state changes to realize synapse-like plasticity primitives; here, “training” is largely implemented through stimulation and programming protocols (pulse trains, duty cycles) and through electrolyte tuning [163], [164]. It should be noted that the emergence of *programmable* nanofluidic memristors strengthens this co-design perspective by showing that qualitatively different hysteresis regimes can be accessed within the same hardware through controlled changes in electrolyte composition and operating conditions, turning device programming into an explicit design handle for neuromorphic functionality [164]. Overall, digital-twin workflows benefit from model-guided design but remain sensitive to calibration drift and surface-chemistry variability; reservoir-style approaches are well matched to rich transient ion dynamics but sacrifice end-to-end optimization of internal states; and in-materio adaptation offers substrate-native plasticity, yet currently remains constrained by limited controllability, observability, and update speed.

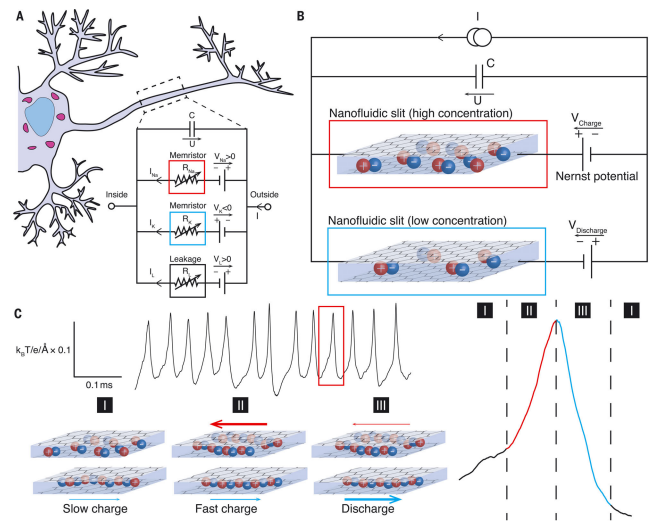
### E. REPRESENTATIVE DEMONSTRATIONS

Representative demonstrations in iontronic neural computing now span ionic spiking elements, synapse-like memory devices, reservoir-style processing, configurable memristive primitives, and biointerfaced spiking components. At the device-physics end, Hodgkin–Huxley-like spiking dynamics have been reproduced in angstrom-scale nanofluidic slit architectures, establishing an “artificial axon” concept where coupled ionic transport nonlinearities yield autonomous spike trains [160]. Figure 10 illustrates this principle and highlights how differential ionic selectivity and coupled transport give rise to quiescent, initiation, and discharge phases that closely resemble biological spiking.

At the learning-synapse end, two-dimensional nanofluidic channels have demonstrated long-term memory and synapse-like dynamics under electrochemical control, motivating the view of iontronics as a route to neuromorphic functionality that is directly rooted in aqueous transport physics [163]. On the systems side, a particularly clear engineering demonstration is fluidic reservoir computing with iontronic nanochannels, which achieves handwritten-digit classification with an experimentally reported accuracy of 81% under a simple readout protocol, highlighting how

transient concentration polarization can act as short-term memory without full in-situ weight programming [166].

Finally, programmability at the device level has been strengthened by nanofluidic memristors whose hysteresis type can be switched by tuning electrolyte chemistry and operating conditions, suggesting a practical route to configurable iontronic primitives for larger circuits [164]. Bio-interfacing directions are illustrated by nanofluidic spiking synapses based on PEDOT:PSS (a popular composite material in this field) confined in nanopores, providing spiking behavior in an aqueous environment compatible with bioelectronic integration [168].



**FIGURE 10.** Artificial Hodgkin-Huxley-like neuron based on 2D ionic memristors. (A) Circuit representation of the Hodgkin-Huxley model showing the ionic branches and corresponding Nernst potentials. (B) Brownian-dynamics simulation of the ionic prototype, consisting of two slits connected to reservoir pairs with different ionic concentrations. (C) Spontaneous voltage spikes generated by the prototype, with the right panel illustrating the quiescent, initiation, and discharge phases of a single spike. The figure exemplifies how nanofluidic transport asymmetry and ionic coupling can reproduce biologically familiar excitability in a fully physical device. Reprinted from [160] with permission from the American association for the advancement of science.

### F. ENGINEERING CONSTRAINTS

The practical application of iontronics is defined by a fundamental trade-off between bio-compatibility and computational speed. The primary limitation is the inherent mobility of the charge carriers; ions move approximately  $10^5$  times slower than electrons in silicon. Consequently, iontronic circuits operate in the kHz regime, which renders them unsuitable for general-purpose number crunching but perfectly matches the millisecond timescale of biological action potentials.

Furthermore, at the nanoscale, systems are subject to significant flicker noise and thermal fluctuations (Brownian motion). While traditionally viewed as a defect in deterministic logic, this stochasticity is increasingly exploited as a computational resource for probabilistic inference [160]. A significant manufacturing hurdle remains the fabrication

of circuits at the single-nanometer scale which currently limits large-scale integration compared to established CMOS processes [163].

### G. POSITION WITHIN PHYSICAL NEURAL COMPUTING

Iontronics occupies a unique niche as the “missing link” between abiotic electronics and biotic wetware. It is the only substrate that speaks the native language of biology. Its future lies not in competing with silicon for FLOPs, but in domains requiring high-fidelity bio-integration, acting as the computational interface between artificial intelligence and the human nervous system.

## XII. CROSS-SUBSTRATE BENCHMARKING

Comparing physical neural substrates is inherently difficult: reported performance figures depend not only on the substrate physics, but also on the experimental interface (encoding/decoding), the training protocol, and what is counted as system energy (core physics only vs. peripheral actuation and transduction). Rather than proposing a definitive benchmark, we introduce, in this section, a *first-order comparative framework*. This approach aligns with the methodology advocated by Stepney, whose recent tutorial emphasizes the need to rigorously define the “physical reservoir” properties to avoid false positives in computational capacity [167]. We therefore propose a framework that (i) uses a small set of standardized tasks to assess capability, and (ii) positions substrates in a common, physically interpretable landscape defined by characteristic speed, energy scale, and plasticity. The goal is to make trade-offs explicit and to identify where each substrate is a natural fit, not to crown a universal winner.

### A. STANDARDIZED TASKS: CAPABILITY TESTS INSTEAD OF A SINGLE SCOREBOARD

A single benchmark cannot fairly represent the diversity of physical substrates (e.g., a DNA network optimized for molecular diagnostics versus a photonic circuit optimized for GHz signal processing). We therefore propose two minimal capability domains that recur across the literature and collectively probe the key neural primitives discussed throughout this review: (1) *static classification*, requiring separability in a high-dimensional state space, and (2) *dynamic prediction/control*, based on fading memory and nonlinear temporal processing. These are not meant as rigid standards, but as *capability tests* that allow comparing whether a substrate can execute core neural functions at all. We do not claim that the proposed framework is complete; rather, we view it as a starting point that makes otherwise implicit trade-offs explicit and debatable.

The canonical reference for static classification is MNIST-style image classification ( $28 \times 28$ ) or compressed variants thereof, which stresses parallel input encoding and stable linear/nonlinear separation. Demonstrations exist for DNA strand-displacement classifiers on compressed MNIST-like inputs [42], for deep physical neural networks trained

with physics-aware backpropagation in wave-based platforms [16], and for integrated photonic neural networks based on coherent interference [112]. In electronic “computational memory”, memristive/PCM crossbars have demonstrated large-scale neural inference and learning with weights stored as conductances [87], [93], [94].

In the second domain, temporal tasks such as Mackey–Glass prediction or NARMA- $k$  emulation probe fading memory and nonlinearity; closed-loop control tasks additionally probe stability under feedback. Reservoir computing has long provided substrate-agnostic baselines (e.g., liquid/wave reservoirs) [27], [120]. Mechanical and soft robotic systems can embody dynamics directly (e.g., compliant bodies as reservoirs) [28], while microfluidic and pneumatic circuits have demonstrated autonomous oscillations and gait-like pattern generation driven by the system physics [153], [169]. Bio-hybrid systems (neuronal cultures) have shown closed-loop learning in real-time control environments [82]. Photonic reservoirs offer high-bandwidth temporal processing via delay dynamics [121].

On a practical note, the dominant limitation for many substrates is not the internal physics but the I/O mapping: providing 784 independent analog inputs (MNIST) is trivial in electronics, but a major engineering challenge for pressure-driven microfluidics; conversely, supplying molecular inputs to a silicon classifier often requires expensive transduction that wetware can avoid. These asymmetries motivate the second element of the framework: a physically grounded landscape view.

To make these capability tests more practically usable, Table 1 summarizes a minimal reference protocol. The protocol is intentionally modest: it is not meant to prescribe a single benchmark for all substrates, but to define a common reporting baseline. Substrate-specific adaptations are unavoidable, for example when a molecular system cannot expose hundreds of independent inputs or when a mechanical body naturally processes low-dimensional continuous signals. However, deviations from the reference protocol should be stated explicitly, so that differences in performance can be attributed to the substrate rather than to hidden choices in encoding, noise injection, or energy accounting.

This protocol also clarifies the interpretation of the landscape comparison below: speed, energy, and plasticity should not be read as isolated material constants, but as system-level quantities whose meaning depends on the stated task, input encoding, noise regime, and accounting boundary.

### B. QUANTITATIVE COMPARISON: THE PHYSICS–COMPUTE LANDSCAPE

Table 2 aggregates order-of-magnitude characteristics derived from representative demonstrations discussed across Parts 2 and 3. We emphasize that these values are *illustrative* rather than definitive: they indicate characteristic regimes, not optimized engineering limits. Where possible we cite

**TABLE 1. Minimal reference protocol for cross-substrate benchmarking of physical neural systems. The proposed values are intended as reporting defaults rather than rigid standards; substrate-specific deviations should be stated explicitly.**

Protocol element	Static classification task	Dynamic prediction/control task	Reporting requirement
Input dimensionality	Reference input: MNIST-style $28 \times 28$ images where feasible; compressed variants with $d \in \{16, 64, 256\}$ for substrates with limited I/O.	Low-dimensional time series with $d \in [1, 16]$ input channels; sequence length $T = 100$ – $1000$ samples or substrate-specific physical time horizon.	Report the physical encoding of each input dimension, compression method if used, and whether dimensions are applied in parallel or sequentially.
Noise levels	Evaluate clean inputs and additive input/readout noise at representative SNR levels, e.g., 40 dB, 20 dB, and 10 dB where experimentally meaningful.	Evaluate process or readout perturbations at matched SNR levels; for closed-loop tasks additionally report disturbance amplitude and duration.	State whether noise is injected digitally, physically at the input, internally by the substrate, or only at readout; report mean and variance over repeated trials.
Training regime	Report whether parameters are obtained by ex-situ digital training, device-aware training, in-situ adaptation, reservoir readout training, or closed box optimization.	Same categories; additionally report whether temporal memory is intrinsic to the substrate or provided by external delay/state variables.	Specify train/test split, number of training examples or episodes, number of physical evaluations, and whether the same physical instance is reused.
Performance metric	Classification accuracy, confusion matrix, and robustness degradation under noise/compression.	Prediction error (e.g., NRMSE) for time-series tasks; task reward, tracking error, or stability margin for control tasks.	Report mean, standard deviation, number of runs, and whether variability is dominated by device noise, fabrication spread, biological variability, or environmental drift.
Latency and throughput	Time from input presentation to stable or readable output, including settling or relaxation time.	Sampling interval, memory horizon, closed-loop reaction time, and maximum stable control/update rate.	Separate intrinsic substrate time constants from encoding, actuation, readout, and digital post-processing latency.
Energy accounting	Energy per inference or per classified sample.	Energy per prediction step, control step, or physical episode.	State the accounting boundary: core substrate only, substrate plus actuation/readout, or full system including ADC/DAC, pumps, lasers, cryogenics, incubation, or life support.
I/O and interface overhead	Number and type of physical input/output channels; transduction required for applying and reading inputs.	Same, with additional reporting of feedback-loop interfaces for control tasks.	Report whether the dominant cost lies in the substrate dynamics or in peripheral sensing, actuation, conversion, and stabilization hardware.

experimental demonstrations; where literature varies widely (e.g., optical energy per MAC depending on whether tuning and detection are included), we report broad ranges and state what is being counted.

Following our analysis, we can say that the following observations are robust across the surveyed literature.

#### 1) THERMODYNAMIC AND BIOCHEMICAL LIMITS ENABLE EXTREME ENERGY EFFICIENCY, BUT AT LOW BANDWIDTH

Molecular systems (DNA) and iontronic devices can approach very low physical energy scales per state transition, but their usable throughput is constrained by diffusion, reaction kinetics, and stochastic fluctuations [42], [160]. This makes them attractive when *transduction dominates* and sample rates are intrinsically slow (e.g., diagnostics), rather than for high-frequency control.

#### 2) WAVE AND FIELD SUBSTRATES EXCEL AT BANDWIDTH, BUT PAY FOR NONLINEARITY AND INTERFACING

Photonic and mechanical wave systems implement linear transformations “for free” via propagation and interference, supporting extremely high bandwidth for inference and signal processing [16], [112]. However, strong, compact nonlinearities and stable large-scale programmability remain system-level challenges, often pushing architectures toward hybrid electro-optic control [114], [115].

#### 3) IN-MEMORY ELECTRONIC SUBSTRATES SIT AT A FAVORABLE MIDPOINT IF CO-DESIGNED WITH ALGORITHMS

Memristive and PCM crossbars offer dense, scalable MAC at moderate-to-high bandwidth while reducing memory movement. Yet their advantages are inseparable from mixed-precision and device-aware algorithm design, because drift, noise, and nonlinear updates violate ideal arithmetic [85], [93], [94]. In this sense, they represent the most “engineering-ready” route to physical neural hardware at scale.

#### C. SYNTHESIS: COMPUTATIONAL IMPEDANCE MATCHING

The landscape view suggests that physical neural substrates are not drop-in replacements for silicon, but specialized accelerators whose value emerges when the physics of the substrate matches the physics of the problem. We term this *computational impedance matching*: choose the substrate that minimizes unavoidable transduction and that naturally implements the dominant dynamics of the task.

- **Chemical-to-decision pipelines.** If the input is already molecular (biomarkers, metabolites), DNA or biochemical networks can compute *in situ* without expensive sensing and digitization; long reaction times are acceptable when sample rates are low [42].
- **High-bandwidth inference.** If the task is ultra-fast signal processing (RF, optical communications), photonic inference layers can exploit propagation

TABLE 2. Comparative landscape of physical neural substrates.

Substrate	Info carrier	Speed (Hz)	Energy (J/op)	Plasticity	Demonstrated capability	Dominant bottleneck / natural niche
Silicon (GPU/TPU)	Electrons	$\sim 10^9$	$\sim 10^{-12}$	Simulated	Broadly solved	Excellent generality; heavy transduction for "non-digital" inputs.
Memristive / PCM in-memory	Electrons (conductance states)	$\sim 10^6$ – $10^9$	$\sim 10^{-15}$ – $10^{-12}$	Device / mixed-precision	Large-scale inference and learning demos	Array nonidealities, ADC/DAC overhead; best for dense in-memory MAC near sensors [85], [87], [93], [94].
Ferroelectric / In-Sensor	Polarization domains	$\sim 10^7$ – $10^9$	$\sim 10^{-15}$	Hysteresis / Domain walls	In-sensor tactile computing	Imprint effects; best for zero-latency edge sensing [7].
Spintronic / Superconducting	Spin / Flux quanta	$\sim 10^9$ – $10^{11}$	$\sim 10^{-19}$ – $10^{-16}$	Intrinsic / STDP	Skyrmion reservoirs; 50 GHz SNNs	Cryogenics (SC), thermal stability; best for ultra-low power HPC [102], [106].
Photonic neural hardware	Photons (phase/amplitude)	$\sim 10^9$ – $10^{11}$	(core $\ll 10^{-12}$ , system-dependent)	Mostly configured (hybrid)	High-speed classification; photonic reservoirs	Nonlinearity and opto-electronic interface; best for ultrafast inference and RF/optical signal processing [112], [114], [115], [121].
Mechanical / acoustic	Phonons / elastic waves	$\sim 10^3$ – $10^6$	$\sim 10^{-9}$ – $10^{-6}$	Coupled / emergent	Wave-based classification; embodied dynamics	Manufacturability, damping, sensing; best for smart structures and embedded vibration processing [16], [28].
Microfluidic / pneumatic	Fluid mass / pressure	$\sim 10^0$ – $10^2$	$\sim 10^{-6}$ – $10^{-3}$	In-situ / fixed logic	Oscillators, gait generation, soft control	I/O fan-in and actuation counted in energy; best for untethered soft robotics in harsh EM environments [153], [169].
Iontronic / nanofluidic	Ions (charge transport)	$\sim 10^2$ – $10^4$	$\sim 10^{-16}$ – $10^{-12}$	Memristive / STDP-like	Spiking primitives (axon-like)	Nanofabrication precision and noise; best for bio-interfaces and low-power spiking primitives [160], [163].
DNA (wet molecular)	Molecules / concentrations	$\sim 10^{-5}$ – $10^{-3}$	$\sim 10^{-19}$ – $10^{-16}$	Static / slow adaptive	Compressed MNIST-like classification	Latency and readout; best for in-sample diagnostics where transduction dominates silicon [42].
Chemical (reaction-diffusion, BZ)	Wavefronts / chemical potential	$\sim 10^{-2}$ – $10^0$	$\sim 10^{-8}$ – $10^{-5}$	Emergent	Spatial computing (maze, Voronoi)	Programmability and cascadability; best for analog spatial geometry problems [60], [68].
Living (cells / neurons)	Ions + biochemical regulation	$\sim 10^1$ – $10^3$	$\sim 10^{-15}$ – $10^{-12}$	True (STDP, adaptation)	Closed-loop learning/control (e.g., Pong)	Reproducibility and maintenance; best for adaptive bio-hybrid controllers and neuroscience-in-the-loop [82].

bandwidth, with electronics providing nonlinearity and control [112], [115]:

- **Dense edge inference and adaptive sensing.** If the bottleneck is moving weights and activations between memory and compute, memristive/PCM computational memory can provide energy-efficient MAC close to sensors, provided algorithms are co-designed for device nonidealities [93], [94].
- **Embodied control.** If the task is to generate or stabilize physical motion in soft systems, microfluidic/pneumatic or mechanical substrates can embody the control dynamics directly, often counting “computation” as mechanical work [28], [153], [169].
- **Bio-hybrid adaptivity.** If continuous adaptation under uncertainty is paramount, living neural tissue can learn in closed-loop settings, at the cost of reproducibility and maintenance overhead [82].

In short, a useful benchmark for physical intelligence must not only score accuracy on a task, but must also characterize the interface cost and the physics-algorithm fit. The framework above provides a starting point for that comparison and highlights where further standardized reporting (especially for end-to-end energy and I/O costs) is needed.

### XIII. CONCLUSION AND OUTLOOK

In this review, we have surveyed a broad spectrum of physical substrates for neural computation, ranging from molecular reaction networks and living cellular systems to solid-state electronic, photonic, mechanical, and fluidic hardware. Taken together, these developments suggest that neural computation is not tied to a particular material platform, but rather emerges whenever a physical system provides three ingredients: rich nonlinear dynamics, a mechanism for weighted interaction between degrees of freedom, and a process for adaptation or effective training [70], [85], [115]. The diversity of substrates explored in the preceding parts demonstrates that these ingredients can be realized through fundamentally different physical mechanisms: charge transport, wave interference, chemical kinetics, mechanical deformation, fluid flow, or biological plasticity [16], [28], [42], [82], [121].

#### A. FROM ALGORITHMS TO PHYSICAL DESIGN PRINCIPLES

A recurring theme across all parts of this article is that performance in physical neural systems does not arise from abstract learning algorithms alone, but from careful co-design of algorithms, architectures, and materials [93], [94], [112]. In digital accelerators, the physical layer is typically treated as an implementation detail. In contrast, physical neural computing elevates the substrate to a first-class design variable: device physics determines which operations are cheap, which sources of noise are unavoidable, and which forms of plasticity are accessible [79], [85].

We see the following general design principles emerge from the surveyed literature:

- **Exploit native dynamics.** Systems achieve their highest efficiency when the dominant computation (e.g., linear transformation, diffusion, wave propagation, or relaxation) coincides with the natural dynamics of the substrate, rather than being emulated through control circuitry [28], [121].
- **Locality of interaction.** Many substrates support only local coupling (chemical reactions, mechanical stress, nearest-neighbor electrical conduction). Architectures that embrace this locality are more scalable and robust than those attempting to impose global connectivity [42], [75].
- **Tolerance to imprecision.** Noise, drift, and variability are intrinsic to most physical substrates. Successful systems therefore rely on learning rules, redundancy, or reservoir-style computation that remain functional under significant uncertainty [92], [94], [121].
- **Interface-aware design.** In many demonstrations, the energetic and temporal cost of sensing, actuation, and signal conversion dominates the cost of the internal computation. End-to-end efficiency must therefore be evaluated at the level of the complete physical-digital loop [42], [112], [153].

These principles suggest that the traditional separation between “algorithm” and “hardware” becomes blurred in physical intelligence: the computational model must be chosen with the substrate in mind, and vice versa [85].

### B. THE ROLE OF BENCHMARKING AND COMPARATIVE EVALUATION

The benchmarking framework we introduced in Section XII provides a first step toward systematic comparison of heterogeneous substrates [27], [120]. While no single task or metric can capture the full richness of physical neural systems, standardized capability tests and physically interpretable performance dimensions (speed, energy scale, plasticity, and interface cost) enable more meaningful cross-domain discussion than isolated demonstrations [16], [82], [87], [112].

In the long term, progress in this field will benefit from community practices analogous to those in conventional machine learning: shared datasets adapted to physical constraints, reference tasks for static and dynamic processing, and transparent reporting of what is included in energy and latency measurements [85], [93]. In the absence of such conventions, performance comparisons between such disparate technologies like molecular classifiers, photonic inference engines, and in-memory electronic accelerators will lack the rigorous framework necessary for commensurate evaluation and cumulative progress.

### C. APPLICATION DOMAINS AND REALISTIC EXPECTATIONS

The preceding chapters make clear that physical neural systems are unlikely to displace digital computing as a universal platform. Their value instead lies in domains where their physical properties provide decisive advantages and where the required levels of reliability, stability, and environmental robustness remain compatible with the substrate:

- **In-sample molecular decision making**, where chemical or DNA-based networks can operate directly on biochemical inputs without costly transduction [42].
- **Ultrafast signal processing**, where photonic systems exploit wave propagation to perform linear inference at bandwidths unattainable electronically [112], [115].
- **Energy-constrained edge inference**, where memristive and phase-change memories reduce data movement by colocating storage and computation [93], [94].
- **Embodied control in soft and adaptive machines**, where mechanical or fluidic substrates naturally integrate computation with actuation [28], [153], [169].
- **Bio-hybrid systems**, where living neural tissue provides continuous adaptation in uncertain environments [70], [82].

At the same time, these niches differ strongly in their path to deployment. Some substrates are most promising for highly specialized laboratory or medical settings, where direct operation in the native physical domain outweighs limited speed or reusability; others are closer to technological translation because they can be fabricated, packaged, and interfaced with greater repeatability. Realistic expectations therefore require matching each substrate not only to a computational advantage, but also to its achievable operating stability over time.

We believe that recognizing these niches is essential to avoid unrealistic expectations and to guide investment toward problems where physical intelligence can offer genuine benefits over mature digital technologies. Equally important is to acknowledge that the most compelling application domains are those in which the reliability requirements are aligned with the physical operating regime of the substrate rather than inherited uncritically from conventional digital hardware.

### D. RELIABILITY, STABILITY, AND TRANSITION TO REAL-WORLD APPLICATIONS

Despite rapid progress, the transition of physical neural systems from laboratory demonstrations to practical applications depends not only on computational performance, but on whether reliable and stable operation can be maintained under realistic conditions. Across substrates, the central challenge is that computation is embodied in matter whose properties drift, age, fluctuate, or depend sensitively on the environment. As a result, deployment requires more than proof-of-principle functionality: it requires calibration strategies, tolerance

to variability, and application scenarios whose operational demands match the physical limits of the substrate.

These constraints appear in different forms across the PNN landscape. In molecular and chemical systems, reaction kinetics, cross-talk, reagent depletion, and resetability limit long-term reuse and favor applications such as in-sample diagnostics, programmable assays, or lab-on-chip decision making, where direct biochemical operation can justify lower throughput and stricter environmental control [42]. In living and bio-hybrid systems, biological variability, maintenance overhead, and limited interface bandwidth remain major barriers to standardized deployment, suggesting that near-term impact is more likely in closed-loop experimental platforms, biointerfaces, and scientific model systems than in robust mass-market computing technologies [70], [82].

By contrast, solid-state substrates such as memristive, phase-change, and ferroelectric systems are in several respects closer to practical edge-AI integration, but they remain constrained by device drift, stochastic switching, endurance limits, and the energy and area costs of peripheral interfacing [93], [94]. Photonic systems offer exceptional bandwidth and low-latency inference, yet their transition to real-world use depends critically on calibration stability, thermal robustness, and efficient electro-optical I/O [112], [115]. Mechanical, metamaterial, microfluidic, and iontronic approaches are particularly attractive where sensing, computation, and actuation or biochemical interfacing must be co-designed, but their prospects depend on controlling fatigue, fouling, packaging complexity, and environmental sensitivity over extended operating periods [28], [153], [163], [169].

A primary concern is therefore scalability in the broader engineering sense: not merely increasing system size, but doing so while preserving reproducibility, controllability, and acceptable failure rates [70], [85]. This issue is closely linked to programmability and observability; particularly in wetware and bio-hybrid systems, achieving identical behavior across nominally identical devices remains difficult to guarantee [70], [82]. Furthermore, the development of efficient learning rules for training at scale constitutes a major research gap, especially when dealing with non-ideal and only partially observable substrates [2], [79], [94]. Finally, long-term stability remains a critical bottleneck, as drift, fatigue, and aging effects inherent to almost all non-digital substrates complicate deployment beyond controlled laboratory settings [93], [163].

In our view, the most credible transition paths will therefore be substrate-specific rather than universal. PNN platforms are most likely to succeed first in application niches where their distinctive physics provides a decisive systems advantage and where reliability can be ensured either by environmental control, by hybrid digital supervision, or by a task formulation that tolerates limited precision and gradual drift. Overcoming these multifaceted obstacles will necessitate a concerted interdisciplinary effort across

machine learning, materials science, device physics, biology, and control theory.

### E. PERSPECTIVE

Physical neural computing reframes intelligence as a property that can be distributed across matter, rather than confined to abstract algorithms executed on standardized hardware [70], [115]. Whether implemented in molecules, photons, deformable bodies, or living cells, neural computation appears as a general organizational principle of dynamical systems far from equilibrium [79]. We expect that the most successful systems will be hybrid, combining multiple of these physical substrates rather than relying on a single universal platform. In such systems, digital components will often remain essential not only for orchestration and training, but also for calibration, error mitigation, and the stabilization of otherwise fragile physical dynamics.

From this perspective, current demonstrations should not be judged primarily by their raw performance relative to GPUs or CPUs, but by the new regimes of computation they make accessible: computation at the scale of chemistry, at the speed of light, or embedded directly into adaptive materials [28], [42], [112]. As these regimes mature, they may reshape how intelligent systems are engineered, not by replacing digital computers, but by complementing them with substrates whose physics performs part of the computation by default [85]. Their long-term impact, however, will depend on whether these gains can be translated into sufficiently reliable, stable, and maintainable operation outside carefully controlled laboratory settings.

In this sense, we believe that the study of physical neural networks is not only a search for alternative hardware, but also an exploration of the physical limits and manifestations of learning itself.

### ACKNOWLEDGMENT

The authors acknowledge the use of large language models, specifically Google Gemini and OpenAI ChatGPT, in the preparation of this manuscript. These tools were used to support the literature review on the state of the art for Sections III and XI. Wherever possible, the papers were accessed and checked individually (the vast majority). Otherwise, the abstracts were used in combination with AI summaries to extract the relevant information. In addition, AI was utilized to assist with the LaTeX formatting of tables and figures and to perform linguistic refinement and stylistic smoothing of the entire text. All AI-generated suggestions and research outputs were critically reviewed, verified, and edited by the authors, who maintain full responsibility for the technical accuracy and integrity of the final content.

### REFERENCES

- [1] J. Shalf, "The future of computing beyond Moore's law," *Phil. Trans. Roy. Soc. A, Math., Phys. Eng. Sci.*, vol. 378, no. 2166, 2020, Art. no. 20190061.
- [2] A. Momeni et al., "Training of physical neural networks," *Nature*, vol. 645, no. 8079, pp. 53–61, Sep. 2025.

- [3] J. Backus, "Can programming be liberated from the von Neumann style? A functional style and its algebra of programs," *Commun. ACM*, vol. 21, no. 8, pp. 613–641, 1978.
- [4] C. Mead, "Neuromorphic electronic systems," *Proc. IEEE*, vol. 78, no. 10, pp. 1629–1636, 1990.
- [5] D. B. Strukov, G. S. Snider, D. R. Stewart, and R. S. Williams, "The missing memristor found," *Nature*, vol. 453, no. 7191, pp. 80–83, 2008.
- [6] A. Adamatzky, B. D. L. Costello, and T. Asai, *Reaction-Diffusion Computers*. Amsterdam, The Netherlands: Elsevier, 2005. [Online]. Available: <https://www.sciencedirect.com/book/monograph/9780444520425/reaction-diffusion-computers>
- [7] Y. Chen, J. Cao, J. Qiu, D. Yang, M. Liu, M. Zhang, C. Li, Z. Wu, J. Yu, X. Zhang, X. Chen, Z. Huang, E. Song, M. Wang, Q. Liu, and M. Liu, "Capacitive in-sensor tactile computing," *Nature Commun.*, vol. 16, no. 1, p. 5691, Jul. 2025.
- [8] Y. Hu, J. Zhang, K. Shen, W. Shen, H. K. Lee, and S. Tang, "Intelligent molecular logic computing toolkits: Nucleic acid-based construction, functionality, and enhanced biosensing applications," *Chem. Sci.*, vol. 16, no. 43, pp. 20139–20180, 2025.
- [9] R. P. Feynman, "Simulating physics with computers," *Int. J. Theor. Phys.*, vol. 21, nos. 6–7, pp. 467–488, Jun. 1982.
- [10] J. J. Hopfield and D. W. Tank, "Neural computation of decisions in optimization problems," *Biol. Cybern.*, vol. 52, no. 3, pp. 141–152, Jul. 1985.
- [11] A. Rolandi, P. Abiuso, P. Lipka-Bartosik, M. Aifer, P. J. Coles, and M. Perarnau-Llobet, "Energy-time-accuracy tradeoffs in thermodynamic computing," 2026, *arXiv:2601.04358*.
- [12] L. Qian, E. Winfree, and J. Bruck, "Neural network computation with DNA strand displacement cascades," *Nature*, vol. 475, no. 7356, pp. 368–372, Jul. 2011.
- [13] H. Sohr, *The Navier-Stokes Equations: An Elementary Funct. Analytic Approach*. Cambridge, MA, USA: Birkhäuser, 2001.
- [14] X. Wang, H. Borrás, B. Klein, and H. Fröning, "Variance-aware noisy training: Hardening DNNs against unstable analog computations," in *Proc. Mach. Learn. Knowl. Discovery Databases. Res. Track*, 2025, pp. 147–163.
- [15] A. Adamatzky, J. Holley, L. Bull, and B. De Lacy Costello, "On computing in fine-grained compartmentalised belousov–zhabotinsky medium," *Chaos, Solitons Fractals*, vol. 44, no. 10, pp. 779–790, Oct. 2011.
- [16] L. G. Wright, T. Onodera, M. M. Stein, T. Wang, D. T. Schachter, Z. Hu, and P. L. McMahon, "Deep physical neural networks trained with backpropagation," *Nature*, vol. 601, no. 7894, pp. 549–555, 2022.
- [17] M. Davies et al., "Loihi: A neuromorphic manycore processor with on-chip learning," *IEEE Micro*, vol. 38, no. 1, pp. 82–99, Jan. 2018.
- [18] F. Akopyan, J. Sawada, A. Cassidy, R. Alvarez-Icaza, J. Arthur, P. Merolla, N. Imam, Y. Nakamura, P. Datta, G.-J. Nam, B. Taba, M. Beakes, B. Brezzo, J. B. Kuang, R. Manohar, W. P. Risk, B. Jackson, and D. S. Modha, "TrueNorth: Design and tool flow of a 65 mW 1 million neuron programmable neurosynaptic chip," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 34, no. 10, pp. 1537–1557, Oct. 2015.
- [19] E. M. Purcell, "Life at low Reynolds number," *Amer. J. Phys.*, vol. 45, no. 1, pp. 3–11, Jan. 1977.
- [20] C. Mead, *Analog VLSI and Neural Systems*. Reading, MA, USA: Addison-Wesley, 1989.
- [21] D. Soloveichik, G. Seelig, and E. Winfree, "DNA as a universal substrate for chemical kinetics," *Proc. Nat. Acad. Sci. USA*, vol. 107, no. 12, pp. 5393–5398, Mar. 2010. [Online]. Available: <https://www.pnas.org/doi/abs/10.1073/pnas.0909380107>
- [22] A. Hjelmfelt, E. D. Weinberger, and J. Ross, "Chemical implementation of neural networks and Turing machines," *Proc. Nat. Acad. Sci. USA*, vol. 88, no. 24, pp. 10983–10987, Dec. 1991. [Online]. Available: <https://www.pnas.org/doi/abs/10.1073/pnas.88.24.10983>
- [23] B. Florijn, C. Coullais, and M. van Hecke, "Programmable mechanical metamaterials," *Phys. Rev. Lett.*, vol. 113, no. 17, Oct. 2014, Art. no. 175503. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.113.175503>
- [24] L. M. Adleman, "Molecular computation of solutions to combinatorial problems," *Science*, vol. 266, no. 5187, pp. 1021–1024, Nov. 1994.
- [25] S. Kirkpatrick, C. D. Gelatt Jr., and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, no. 4598, pp. 671–680, May 1983.
- [26] T. W. Hughes, I. A. D. Williamson, M. Minkov, and S. Fan, "Wave physics as an analog recurrent neural network," *Sci. Adv.*, vol. 5, no. 12, p. 6946, Dec. 2019.
- [27] C. Fernando and S. Sojakka, "Pattern recognition in a bucket," in *Proc. Adv. Artif. Life*, 2003, pp. 588–597.
- [28] K. Nakajima, H. Hauser, T. Li, and R. Pfeifer, "Information processing via physical soft body," *Sci. Rep.*, vol. 5, no. 1, p. 10487, May 2015, doi: [10.1038/srep10487](https://doi.org/10.1038/srep10487).
- [29] K. Vandoorne, P. Mechet, T. Van Vaerenbergh, M. Fiers, G. Morthier, D. Verstraeten, B. Schrauwen, J. Dambre, and P. Bienstman, "Experimental demonstration of reservoir computing on a silicon photonics chip," *Nature Commun.*, vol. 5, no. 1, p. 3541, Mar. 2014, doi: [10.1038/ncomms4541](https://doi.org/10.1038/ncomms4541).
- [30] G. Tanaka, T. Yamane, J. B. Héroux, R. Nakane, N. Kanazawa, S. Takeda, H. Numata, D. Nakano, and A. Hirose, "Recent advances in physical reservoir computing: A review," *Neural Netw.*, vol. 115, pp. 100–123, Jul. 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0893608019300784>
- [31] B. Scellier and Y. Bengio, "Equilibrium propagation: Bridging the gap between energy-based models and backpropagation," *Frontiers Comput. Neurosci.*, vol. 11, p. 24, May 2017.
- [32] M. Stern, D. Hexner, J. W. Rocks, and A. J. Liu, "Supervised learning in physical networks: From machine learning to learning machines," *Phys. Rev. X*, vol. 11, no. 2, May 2021, Art. no. 021045. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevX.11.021045>
- [33] S. Dillavou, M. Stern, A. J. Liu, and D. J. Durian, "Demonstration of decentralized physics-driven learning," *Phys. Rev. Appl.*, vol. 18, no. 1, Jul. 2022, Art. no. 014040. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevApplied.18.014040>
- [34] J. J. Hopfield, "Neurons with graded response have collective computational properties like those of two-state neurons," *Proc. Nat. Acad. Sci. USA*, vol. 81, no. 10, pp. 3088–3092, May 1984. [Online]. Available: <https://www.pnas.org/doi/abs/10.1073/pnas.81.10.3088>
- [35] H. Jaeger, "The 'echo state' approach to analysing and training recurrent neural networks," GMD Forschungszentrum Informationstechnik, St. Augustin, Germany, Tech. Rep. GMD-Report 148, 2001.
- [36] F. C. Simmel, "Nucleic acid strand displacement—from DNA nanotechnology to translational regulation," *RNA Biol.*, vol. 20, no. 1, pp. 154–163, Dec. 2023.
- [37] C. Chen, J. Wen, Z. Wen, S. Song, and X. Shi, "DNA strand displacement based computational systems and their applications," *Frontiers Genet.*, vol. 14, Feb. 2023, Art. no. 1120791.
- [38] A. Zagitov, E. S. Korenkov, N. Bashirov, R. A. Maksimov, A. M. Vinogradov, A. Beznosikov, and M. P. Nikitin, "Backpropagation in molecular neural networks: Teaching DNA to solve machine learning tasks," *Preprints*, Jan. 2026, doi: [10.20944/preprints202601.0088.v1](https://doi.org/10.20944/preprints202601.0088.v1)
- [39] P. W. K. Rothmund, "Folding DNA to create nanoscale shapes and patterns," *Nature*, vol. 440, no. 7082, pp. 297–302, Mar. 2006.
- [40] P. Zhan, A. Peil, Q. Jiang, D. Wang, S. Mousavi, Q. Xiong, Q. Shen, Y. Shang, B. Ding, C. Lin, Y. Ke, and N. Liu, "Recent advances in DNA origami-engineered nanomaterials and applications," *Chem. Rev.*, vol. 123, no. 7, pp. 3976–4050, Apr. 2023.
- [41] C. Huang, J. Shao, B. Peng, Q. Guo, P. Li, J. Sun, and Y. Wang, "Design of a universal decoder model based on DNA winner-takes-all neural networks," *IEEE Trans. Comput.*, vol. 74, no. 4, pp. 1267–1277, Apr. 2025.
- [42] K. M. Cherry and L. Qian, "Scaling up molecular pattern recognition with DNA-based winner-take-all neural networks," *Nature*, vol. 559, no. 7714, pp. 370–376, Jul. 2018.
- [43] L. Qian and E. Winfree, "Scaling up digital circuit computation with DNA strand displacement cascades," *Science*, vol. 332, no. 6034, pp. 1196–1201, Jun. 2011.
- [44] X. Liu, Z. Zheng, P. Liu, Z. Yang, H. Hu, M. Wu, X. Lou, F. Xia, K. Tao, L. Li, J. Dai, and X. Xiao, "Deep convolutional and fully-connected DNA neural networks," *Nature Commun.*, vol. 16, no. 1, p. 10629, Nov. 2025.
- [45] J. J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proc. Nat. Acad. Sci. USA*, vol. 79, no. 8, pp. 2554–2558, 1982. [Online]. Available: <https://www.pnas.org/doi/abs/10.1073/pnas.79.8.2554>
- [46] J. Sun, H. Wang, Y. Yue, D. Ling, and Y. Wang, "Design of Hopfield neural network based on DNA strand displacement circuits and its application in sudoku conjecture," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 36, no. 10, pp. 18889–18899, Oct. 2025.
- [47] M. P. O'Hagan, Z. Duan, F. Huang, S. Laps, J. Dong, F. Xia, and I. Willner, "Photocleavable ortho-nitrobenzyl-protected DNA architectures and their applications," *Chem. Rev.*, vol. 123, no. 10, pp. 6839–6887, May 2023.

- [48] K. M. Cherry and L. Qian, "Supervised learning in DNA neural networks," *Nature*, vol. 645, no. 8081, pp. 639–647, Sep. 2025.
- [49] M. Sanjabi and A. Jahanian, "Multi-threshold and multi-input DNA logic design style for profiling the microRNA biomarkers of real cancers," *IET Nanobiotechnology*, vol. 13, no. 7, pp. 665–673, Sep. 2019. [Online]. Available: <https://ietresearch.onlinelibrary.wiley.com/doi/abs/10.1049/iet-nbt.2018.5275>
- [50] S. Okumura, G. Gines, N. Lobato-Dauzier, A. Baccouche, R. Deteix, T. Fujii, Y. Rondelez, and A. J. Genot, "Nonlinear decision-making with enzymatic neural networks," *Nature*, vol. 610, no. 7932, pp. 496–501, Oct. 2022.
- [51] S. Takiguchi, N. Takeuchi, V. Shenshin, G. Gines, A. J. Genot, J. Nivala, Y. Rondelez, and R. Kawano, "Harnessing DNA computing and nanopore decoding for practical applications: From informatics to microRNA-targeting diagnostics," *Chem. Soc. Rev.*, vol. 54, no. 1, pp. 8–32, 2025, doi: [10.1039/d3cs00396e](https://doi.org/10.1039/d3cs00396e).
- [52] L. Rizik, L. Danial, M. Habib, R. Weiss, and R. Daniel, "Synthetic neuromorphic computing in living cells," *Nature Commun.*, vol. 13, no. 1, p. 5602, Sep. 2022.
- [53] C. M. Leon, Y. Wang, F. B. Bisso, A. Mehta, and C. C. Samaniego, "CRISPR-based neuromorphic computing for solving regression and classification," *bioRxiv*, 2025. [Online]. Available: <https://www.biorxiv.org/content/early/2025/12/07/2025.12.03.692209>
- [54] C. Sun, X. Liu, J. Zhong, Q. Zhou, and J. Cheng, "Reusable noncomplementary DNA-based neural network," *J. Amer. Chem. Soc.*, vol. 147, no. 38, pp. 34339–34349, Aug. 2025.
- [55] T. Song and L. Qian, "Heat-rechargeable computation in DNA logic circuits and neural networks," *Nature*, vol. 646, no. 8084, pp. 315–322, Oct. 2025.
- [56] S. Agostini and V. Zarth, "Towards on-chip chemical computing: The oscillating Belousov-Zhabotinsky reaction in silicon microchips," in *Proc. SCS Fall Meeting Abstract Catalogue*, 2025. [Online]. Available: <https://research.ibm.com/publications/towards-on-chip-chemical-computing-the-oscillating-belousov-zhabotinsky-reaction-in-silicon-microchips>
- [57] M. G. Baltussen, T. J. de Jong, Q. Duez, W. E. Robinson, and W. T. S. Huck, "Chemical reservoir computing in a self-organizing reaction network," *Nature*, vol. 631, no. 8021, pp. 549–555, Jul. 2024.
- [58] C. Zhu, B.-Y. Liu, L.-Z. Zhang, and L. Xu, "Autonomous soft robots: Self-regulation, self-sustained, and recovery strategies," *Chin. J. Polym. Sci.*, vol. 43, no. 4, pp. 535–547, Apr. 2025.
- [59] X. Li and M. Lu, "Study on the anti-interference characteristics of neuronal networks: A comparative study of chemical synapses and electrical synapse," *Frontiers Neurosci.*, vol. 19, Jul. 2025, Art. no. 1581347. [Online]. Available: <https://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2025.1581347/full>
- [60] Á. Tóth and K. Showalter, "Logic gates in excitable media," *J. Chem. Phys.*, vol. 103, no. 6, pp. 2058–2066, Aug. 1995.
- [61] M. Cucchi, C. Gruener, L. Petrauskas, P. Steiner, H. Tseng, A. Fischer, B. Penkovsky, C. Matthus, P. Birkholz, H. Kleemann, and K. Leo, "Reservoir computing with biocompatible organic electrochemical networks for brain-inspired biosignal classification," *Sci. Adv.*, vol. 7, no. 34, Aug. 2021, Art. no. eabh0693. [Online]. Available: <https://www.science.org/doi/abs/10.1126/sciadv.abh0693>
- [62] G. Villar, A. D. Graham, and H. Bayley, "A tissue-like printed material," *Science*, vol. 340, no. 6128, pp. 48–52, Apr. 2013.
- [63] M. U. Gaimann and M. Klopotek, "Robustly optimal dynamics for active matter reservoir computing," 2025, *arXiv:2505.05420*.
- [64] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud, "Neural ordinary differential equations," 2018, *arXiv:1806.07366*.
- [65] W. Peng, Z.-S. Liu, and M. Boy, "SPIN-ODE: Stiff physics-informed neural ODE for chemical reaction rate estimation," in *Proc. 28th Eur. Conf. Artif. Intell.*, 2025, pp. 2033–2040, doi: [10.3233/faia251040](https://doi.org/10.3233/faia251040).
- [66] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *J. Comput. Phys.*, vol. 378, pp. 686–707, Feb. 2019.
- [67] J. M. Parrilla-Gutierrez, A. Sharma, S. Tsuda, G. J. T. Cooper, G. Aragon-Camarasa, K. Donkers, and L. Cronin, "A programmable chemical computer with memory and pattern recognition," *Nature Commun.*, vol. 11, no. 1, p. 1442, Mar. 2020.
- [68] O. Steinbock, Á. Tóth, and K. Showalter, "Navigating complex labyrinths: Optimal paths from chemical waves," *Science*, vol. 267, no. 5199, pp. 868–871, Feb. 1995.
- [69] I. Lagzi, S. Soh, P. J. Wesson, K. P. Browne, and B. A. Grzybowski, "Maze solving by chemotactic droplets," *J. Amer. Chem. Soc.*, vol. 132, no. 4, pp. 1198–1199, Feb. 2010.
- [70] L. Smirnova, "Organoid intelligence (OI)—The ultimate functionality of a brain microphysiological system," *ALTEX*, vol. 40, no. 2, pp. 191–203, 2023.
- [71] M. B. Elowitz and S. Leibler, "A synthetic oscillatory network of transcriptional regulators," *Nature*, vol. 403, no. 6767, pp. 335–338, Jan. 2000.
- [72] A. Tamsir, J. J. Tabor, and C. A. Voigt, "Robust multicellular computing using genetically encoded NOR gates and chemical 'wires,'" *Nature*, vol. 469, no. 7329, pp. 212–215, Jan. 2011.
- [73] T. Danino, O. Mondragón-Palomino, L. Tsimring, and J. Hasty, "A synchronized quorum of genetic clocks," *Nature*, vol. 463, no. 7279, pp. 326–330, Jan. 2010.
- [74] A. Tero, S. Takagi, T. Saigusa, K. Ito, D. P. Bebbler, M. D. Fricker, K. Yumiki, R. Kobayashi, and T. Nakagaki, "Rules for biologically inspired adaptive network design," *Science*, vol. 327, no. 5964, pp. 439–442, Jan. 2010.
- [75] A. Adamatzky, *Physarum Machines*. Singapore: World Scientific, 2010.
- [76] H. Cai, Z. Ao, C. Tian, Z. Wu, H. Liu, J. Tchieu, M. Gu, K. Mackie, and F. Guo, "Brain organoid reservoir computing for artificial intelligence," *Nature Electron.*, vol. 6, no. 12, pp. 1032–1039, Dec. 2023.
- [77] G.-Q. Bi and M.-M. Poo, "Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type," *J. Neurosci.*, vol. 18, no. 24, pp. 10464–10472, Dec. 1998.
- [78] R. Daniel, J. R. Rubens, R. Sarpeshkar, and T. K. Lu, "Synthetic analog computation in living cells," *Nature*, vol. 497, no. 7451, pp. 619–623, May 2013.
- [79] K. Friston, "The free-energy principle: A unified brain theory?" *Nature Rev. Neurosci.*, vol. 11, no. 2, pp. 127–138, Jan. 2010.
- [80] M. E. J. Obien, K. Deligkaris, T. Bullmann, D. J. Bakkum, and U. Frey, "Revealing neuronal function through microelectrode array recordings," *Frontiers Neurosci.*, vol. 8, p. 423, Jan. 2015.
- [81] F. D. Jordan, M. Kutter, J.-M. Comby, F. Brozzi, and E. Kurtys, "Open and remotely accessible neuroplatform for research in wetware computing," *Frontiers Artif. Intell.*, vol. 7, May 2024, Art. no. 1376042.
- [82] B. J. Kagan, A. C. Kitchen, N. T. Tran, F. Habibollahi, M. Khajehnejad, B. J. Parker, A. Bhat, B. Rollo, A. Razi, and K. J. Friston, "In vitro neurons learn and exhibit sentience when embodied in a simulated game-world," *Neuron*, vol. 110, no. 23, pp. 3952–3969.e8, Dec. 2022.
- [83] K. Brenner, L. You, and F. H. Arnold, "Engineering microbial consortia: A new frontier in synthetic biology," *Trends Biotechnol.*, vol. 26, no. 9, pp. 483–489, Sep. 2008.
- [84] A. A. K. Nielsen, B. S. Der, J. Shin, P. Vaidyanathan, V. Paralanov, E. A. Strychalski, D. Ross, D. Densmore, and C. A. Voigt, "Genetic circuit design automation," *Sci.*, vol. 352, no. 6281, Apr. 2016, Art. no. aac7341.
- [85] A. Sebastian, M. Le Gallo, R. Khaddam-Aljameh, and E. Eleftheriou, "Memory devices and applications for in-memory computing," *Nature Nanotechnol.*, vol. 15, no. 7, pp. 529–544, Jul. 2020.
- [86] H. Tan and S. van Dijken, "A universal neuromorphic vision processing system," *Nature Electron.*, vol. 7, no. 11, pp. 946–956, Nov. 2024.
- [87] G. W. Burr, R. M. Shelby, S. Sidler, C. di Nolfo, J. Jang, I. Boybat, R. S. Shenoy, P. Narayanan, K. Virwani, E. U. Giacometti, B. N. Kurdi, and H. Hwang, "Experimental demonstration and tolerancing of a large-scale neural network (165 000 synapses) using phase-change memory as the synaptic weight element," *IEEE Trans. Electron Devices*, vol. 62, no. 11, pp. 3498–3507, Nov. 2015.
- [88] J. Torrejon, M. Riou, F. A. Araujo, S. Tsunegi, G. Khalsa, D. Querlioz, P. Bortolotti, V. Cros, K. Yakushiji, A. Fukushima, H. Kubota, S. Yuasa, M. D. Stiles, and J. Grollier, "Neuromorphic computing with nanoscale spintronic oscillators," *Nature*, vol. 547, no. 7664, pp. 428–431, Jul. 2017.
- [89] A. I. Khan, A. Keshavarzi, and S. Datta, "The future of ferroelectric field-effect transistor technology," *Nature Electron.*, vol. 3, no. 10, pp. 588–597, Oct. 2020.
- [90] Q. Li, S. Wang, Z. Li, X. Hu, Y. Liu, J. Yu, Y. Yang, T. Wang, J. Meng, Q. Sun, D. W. Zhang, and L. Chen, "High-performance ferroelectric field-effect transistors with ultra-thin indium tin oxide channels for flexible and transparent electronics," *Nature Commun.*, vol. 15, no. 1, p. 2686, Mar. 2024.
- [91] K. Everschor-Sitte, A. Majumdar, K. Wolk, and D. Meier, "Topological magnetic and ferroelectric systems for reservoir computing," *Nature Rev. Phys.*, vol. 6, no. 7, pp. 455–462, Jun. 2024.

- [92] I. Boybat, M. Le Gallo, S. R. Nandakumar, T. Moraitis, T. Parnell, T. Tuma, B. Rajendran, Y. Leblebici, A. Sebastian, and E. Eleftheriou, "Neuromorphic computing with multi-memristive synapses," *Nature Commun.*, vol. 9, no. 1, p. 2514, Jun. 2018.
- [93] V. Joshi, M. Le Gallo, S. Haefeli, I. Boybat, S. R. Nandakumar, C. Piveteau, M. Dazzi, B. Rajendran, A. Sebastian, and E. Eleftheriou, "Accurate deep neural network inference using computational phase-change memory," *Nature Commun.*, vol. 11, no. 1, p. 2473, May 2020.
- [94] M. Le Gallo, A. Sebastian, R. Mathis, M. Manica, H. Giefers, T. Tuma, C. Bekas, A. Curioni, and E. Eleftheriou, "Mixed-precision in-memory computing," *Nature Electron.*, vol. 1, no. 4, pp. 246–253, Apr. 2018.
- [95] C. Du, F. Cai, M. A. Zidan, W. Ma, S. H. Lee, and W. D. Lu, "Reservoir computing using dynamic memristors for temporal information processing," *Nature Commun.*, vol. 8, no. 1, pp. 1–10, Dec. 2017.
- [96] V. Pistolesi, A. Ceni, G. Milano, C. Ricciardi, and C. Gallicchio, "A memristive computational neural network model for time-series processing," *APL Mach. Learn.*, vol. 3, no. 1, Mar. 2025, Art. no. 016117.
- [97] C. Can Horuz, A. Ceni, C. Gallicchio, and S. Otte, "Scalable memristive-friendly reservoir computing for time series classification," 2026, *arXiv:2604.19343*.
- [98] M. Prezioso, F. Merrikkh-Bayat, B. D. Hoskins, G. C. Adam, K. K. Likharev, and D. B. Strukov, "Training and operation of an integrated neuromorphic network based on metal-oxide memristors," *Nature*, vol. 521, no. 7550, pp. 61–64, May 2015.
- [99] J. Grollier, D. Querlioz, K. Y. Camsari, K. Everschor-Sitte, S. Fukami, and M. D. Stiles, "Neuromorphic spintronics," *Nature Electron.*, vol. 3, no. 7, pp. 360–370, 2020.
- [100] D. Marković and J. Grollier, "Quantum neuromorphic computing," *Appl. Phys. Lett.*, vol. 117, no. 15, Oct. 2020, Art. no. 150501.
- [101] C. H. Marrows, J. Barker, T. A. Moore, and T. Moorsom, "Neuromorphic computing with spintronics," *npj Spintronics*, vol. 2, no. 1, p. 12, Apr. 2024.
- [102] D. Pinna, G. Bourianoff, and K. Everschor-Sitte, "Reservoir computing with random skyrmion textures," *Phys. Rev. Appl.*, vol. 14, no. 5, Nov. 2020, Art. no. 054020.
- [103] T. Yokouchi, S. Sugimoto, B. Rana, S. Seki, N. Ogawa, Y. Shiomi, S. Kasai, and Y. Otani, "Pattern recognition with neuromorphic computing using magnetic field-induced dynamics of skyrmions," *Sci. Adv.*, vol. 8, no. 39, p. 5652, Sep. 2022. [Online]. Available: <https://www.science.org/doi/abs/10.1126/sciadv.abq5652>
- [104] M. L. Schneider, C. A. Donnelly, I. W. Haygood, A. Wynn, S. E. Russek, M. A. Castellanos-Beltran, P. D. Dresselhaus, P. F. Hopkins, M. R. Pufall, and W. H. Rippard, "Synaptic weighting in single flux quantum neuromorphic computing," *Sci. Rep.*, vol. 10, no. 1, p. 934, Jan. 2020.
- [105] S. E. Russek, C. A. Donnelly, M. L. Schneider, B. Baek, M. R. Pufall, W. H. Rippard, P. F. Hopkins, P. D. Dresselhaus, and S. P. Benz, "Stochastic single flux quantum neuromorphic computing using magnetically tunable Josephson junctions," in *Proc. IEEE Int. Conf. Rebooting Comput. (ICRC)*, Oct. 2016, pp. 1–5.
- [106] M. L. Schneider, E. M. Jué, M. R. Pufall, K. Segall, and C. W. Anderson, "A self-training spiking superconducting neuromorphic architecture," *npj Unconventional Comput.*, vol. 2, no. 1, Mar. 2025, Art. no. 5.
- [107] K. Raab, M. A. Brems, G. Beneke, T. Dohi, J. Rothörl, F. Kammerbauer, J. H. Mentink, and M. Kläui, "Brownian reservoir computing realized using geometrically confined skyrmion dynamics," in *Proc. Nature Commun.*, Nov. 2022, vol. 13, no. 1, p. 6982.
- [108] C. Niu, H. Zhang, C. Xu, W. Hu, Y. Wu, Y. Wu, Y. Wang, T. Wu, Y. Zhu, Y. Zhu, W. Wang, Y. Wu, L. Yin, J. Xiao, W. Yu, H. Guo, and J. Shen, "A self-learning magnetic Hopfield neural network with intrinsic gradient descent adaption," *Proc. Nat. Acad. Sci. USA*, vol. 121, no. 51, Dec. 2024, Art. no. 2416294121.
- [109] D. S. Holmes, A. L. Ripple, and M. A. Manheimer, "Energy-efficient superconducting computing—Power budgets and requirements," *IEEE Trans. Appl. Supercond.*, vol. 23, no. 3, Jun. 2013, Art. no. 1701610.
- [110] B. J. Shastri, A. N. Tait, T. Ferreira de Lima, W. H. P. Pernice, H. Bhaskaran, C. D. Wright, and P. R. Prucnal, "Photonics for artificial intelligence and neuromorphic computing," *Nature Photon.*, vol. 15, no. 2, pp. 102–114, Feb. 2021.
- [111] G. Wetzstein, A. Ozcan, S. Gigan, S. Fan, D. Englund, M. Soljačić, C. Denz, D. A. B. Miller, and D. Psaltis, "Inference in artificial intelligence with deep optics and photonics," *Nature*, vol. 588, no. 7836, pp. 39–47, Dec. 2020.
- [112] Y. Shen, N. C. Harris, S. Skirlo, M. Prabhu, T. Baehr-Jones, M. Hochberg, X. Sun, S. Zhao, H. Larochelle, D. Englund, and M. Soljačić, "Deep learning with coherent nanophotonic circuits," *Nature Photon.*, vol. 11, no. 7, pp. 441–446, Jul. 2017.
- [113] D. A. B. Miller, "Self-configuring universal linear optical component," *Photon. Res.*, vol. 1, no. 1, pp. 1–15, 2013.
- [114] A. N. Tait, T. F. de Lima, E. Zhou, A. X. Wu, M. A. Nahmias, B. J. Shastri, and P. R. Prucnal, "Neuromorphic photonic networks using silicon photonic weight banks," *Sci. Rep.*, vol. 7, no. 1, p. 7430, Aug. 2017.
- [115] P. R. Prucnal, B. J. Shastri, T. Ferreira de Lima, M. A. Nahmias, and A. N. Tait, "Recent progress in semiconductor excitable lasers for photonic spike processing," *Adv. Opt. Photon.*, vol. 8, no. 2, pp. 228–299, 2016.
- [116] W. R. Clements, P. C. Humphreys, B. J. Metcalf, W. S. Kolthammer, and I. A. Walsnley, "Optimal design for universal multiport interferometers," *Optica*, vol. 3, no. 12, pp. 1460–1465, 2016.
- [117] I. A. D. Williamson, T. W. Hughes, M. Minkov, B. Bartlett, S. Pai, and S. Fan, "Reprogrammable electro-optic nonlinear activation functions for optical neural networks," *IEEE J. Sel. Topics Quantum Electron.*, vol. 26, no. 1, pp. 1–12, Jan. 2020.
- [118] M. Reck, A. Zeilinger, H. J. Bernstein, and P. Bertani, "Experimental realization of any discrete unitary operator," *Phys. Rev. Lett.*, vol. 73, no. 1, pp. 58–61, Jul. 1994.
- [119] X. Lin, Y. Rivenson, N. T. Yardimci, M. Veli, Y. Luo, M. Jarrahi, and A. Ozcan, "All-optical machine learning using diffractive deep neural networks," *Sci.*, vol. 361, no. 6406, pp. 1004–1008, Jul. 2018.
- [120] L. Appeltant, M. C. Soriano, G. Van der Sande, J. Danckaert, S. Massar, J. Dambre, B. Schrauwen, C. R. Mirasso, and I. Fischer, "Information processing using a single dynamical node as complex system," *Nature Commun.*, vol. 2, no. 1, p. 468, Sep. 2011.
- [121] Y. Paquot, F. Duport, A. Smerieri, J. Dambre, B. Schrauwen, M. Haelterman, and S. Massar, "Optoelectronic reservoir computing," *Sci. Rep.*, vol. 2, no. 1, p. 287, Feb. 2012.
- [122] T. W. Hughes, M. Minkov, Y. Shi, and S. Fan, "Training of photonic neural networks through in situ backpropagation and gradient measurement," *Optica*, vol. 5, no. 7, pp. 864–871, 2018.
- [123] R. H. Lee, E. A. B. Mulder, and J. B. Hopkins, "Mechanical neural networks: Architected materials that learn behaviors," *Sci. Robot.*, vol. 7, no. 71, p. 7278, Oct. 2022.
- [124] R. Pfeifer and J. Bongard, *How the Body Shapes the Way We Think*. Cambridge, MA, USA: MIT Press, 2007.
- [125] S. Kim, C. Laschi, and B. Trimmer, "Soft robotics: A bioinspired evolution in robotics," *Trends Biotechnol.*, vol. 31, no. 5, pp. 287–294, May 2013.
- [126] K. Zahedi and N. Ay, "Quantifying morphological computation," *Entropy*, vol. 15, no. 5, pp. 1887–1915, May 2013. [Online]. Available: <https://www.mdpi.com/1099-4300/15/5/1887>
- [127] K. Ghazi-Zahedi, C. Langer, and N. Ay, "Morphological computation: Synergy of body and brain," *Entropy*, vol. 19, no. 9, p. 456, Aug. 2017. [Online]. Available: <https://www.mdpi.com/1099-4300/19/9/456>
- [128] K. Bertoldi, V. Vitelli, J. Christensen, and M. van Hecke, "Flexible mechanical metamaterials," *Nature Rev. Mater.*, vol. 2, no. 11, p. 17066, Oct. 2017.
- [129] C. Coullais, E. Teomy, K. de Reus, Y. Shokef, and M. van Hecke, "Combinatorial design of textured mechanical metamaterials," *Nature*, vol. 535, no. 7613, pp. 529–532, Jul. 2016.
- [130] C. L. Kane and T. C. Lubensky, "Topological boundary modes in isostatic lattices," *Nature Phys.*, vol. 10, no. 1, pp. 39–45, Jan. 2014.
- [131] A. Ion, L. Wall, R. Kovács, and P. Baudisch, "Digital mechanical metamaterials," in *Proc. Conf. Human Factors Comput. Syst.*, pp. 977–988.
- [132] M. Stern, C. Arinze, L. Perez, S. E. Palmer, and A. Murugan, "Supervised learning through physical changes in a mechanical system," *Proc. Nat. Acad. Sci. USA*, vol. 117, no. 26, pp. 14843–14850, Jun. 2020.
- [133] D. M. J. Dykstra, J. Busink, B. Ennis, and C. Coullais, "Viscoelastic snapping metamaterials," *J. Appl. Mech.*, vol. 86, no. 11, Nov. 2019, Art. no. 111012.
- [134] N. C. Keim, J. Hass, B. Kroger, and D. Wiekler, "Global memory from local hysteresis in an amorphous solid," *Phys. Rev. Res.*, vol. 2, no. 1, Jan. 2020, Art. no. 012004.
- [135] N. Pashine, D. Hexner, A. J. Liu, and S. R. Nagel, "Directed aging, memory, and nature's greed," *Sci. Adv.*, vol. 5, no. 12, p. 4215, Dec. 2019.

- [136] S. Waitukaitis, R. Menaut, B. G.-G. Chen, and M. van Hecke, "Origami multistability: From single vertices to metasheets," *Phys. Rev. Lett.*, vol. 114, no. 5, Feb. 2015, Art. no. 055503.
- [137] Z. Zhai, L. Wu, and H. Jiang, "Mechanical metamaterials based on origami and kirigami," *Appl. Phys. Rev.*, vol. 8, no. 4, Dec. 2021, Art. no. 041319.
- [138] K. Nakajima, T. Li, H. Hauser, and R. Pfeifer, "Exploiting short-term memory in soft body dynamics as a computational resource," *J. Roy. Soc. Interface*, vol. 11, no. 100, Nov. 2014, Art. no. 20140437.
- [139] O. Sigmund and K. Maute, "Topology optimization approaches," *Structural Multidisciplinary Optim.*, vol. 48, no. 6, pp. 1031–1055, Dec. 2013.
- [140] T. Du, K. Wu, P. Ma, S. Wah, A. Spielberg, D. Rus, and W. Matusik, "DiffPD: Differentiable projective dynamics," *ACM Trans. Graph.*, vol. 41, no. 2, pp. 1–21, Apr. 2022.
- [141] D. Dold and D. Aranguren van Egmond, "Differentiable graph-structured models for inverse design of lattice materials," *Cell Rep. Phys. Sci.*, vol. 4, no. 10, Oct. 2023, Art. no. 101586.
- [142] M. Z. Miskin and H. M. Jaeger, "Adapting granular materials through artificial evolution," *Nature Mater.*, vol. 12, no. 4, pp. 326–331, Apr. 2013.
- [143] L. Dong and D. Wang, "Optimal design of three-dimensional voxel printed multimaterial lattice metamaterials via machine learning and evolutionary algorithm," *Phys. Rev. Appl.*, vol. 18, no. 5, Nov. 2022, Art. no. 054050.
- [144] S. Li and X. Mao, "Training all-mechanical neural networks for task learning through in situ backpropagation," *Nature Commun.*, vol. 15, no. 1, p. 10528, Dec. 2024.
- [145] G. M. Whitesides, "The origins and the future of microfluidics," *Nature*, vol. 442, no. 7101, pp. 368–373, Jul. 2006.
- [146] M. Prakash and N. Gershenfeld, "Microfluidic bubble logic," *Science*, vol. 315, no. 5813, pp. 832–835, Feb. 2007. [Online]. Available: <https://www.science.org/doi/abs/10.1126/science.1136907>
- [147] C. S. Law, J. Wang, K. Nielsch, A. D. Abell, J. Bisquert, and A. Santos, "Recent advances in fluidic neuromorphic computing," *Appl. Phys. Rev.*, vol. 12, no. 2, Jun. 2025, Art. no. 021309.
- [148] T. M. Squires and S. R. Quake, "Microfluidics: Fluid physics at the nanoliter scale," *Rev. Modern Phys.*, vol. 77, no. 3, pp. 977–1026, Oct. 2005.
- [149] K. W. Oh, K. Lee, B. Ahn, and E. P. Furlani, "Design of pressure-driven microfluidic networks using electric circuit analogy," *Lab Chip*, vol. 12, no. 3, pp. 515–545, 2012.
- [150] M. A. Unger, H.-P. Chou, T. Thorsen, A. Scherer, and S. R. Quake, "Monolithic microfabricated valves and pumps by multilayer soft lithography," *Science*, vol. 288, no. 5463, pp. 113–116, Apr. 2000. [Online]. Available: <https://www.science.org/doi/abs/10.1126/science.288.5463.113>
- [151] H. Bruus, *Theoretical Microfluidics* (Oxford Master Series in Condensed Matter Physics), vol. 18. London, U.K.: Oxford Univ. Press, 2008.
- [152] J. A. Weaver, J. Melin, D. Stark, S. R. Quake, and M. A. Horowitz, "Static control logic for microfluidic devices using pressure-gain valves," *Nature Phys.*, vol. 6, no. 3, pp. 218–223, Mar. 2010.
- [153] D. J. Preston, P. Rothmund, H. J. Jiang, M. P. Nemitz, J. Rawson, Z. Suo, and G. M. Whitesides, "Digital logic for soft devices," *Proc. Nat. Acad. Sci. USA*, vol. 116, no. 16, pp. 7750–7759, 2019. [Online]. Available: <https://www.pnas.org/doi/abs/10.1073/pnas.1820672116>
- [154] S. K. W. Dertinger, D. T. Chiu, N. L. Jeon, and G. M. Whitesides, "Generation of gradients having complex shapes using microfluidic networks," *Anal. Chem.*, vol. 73, no. 6, pp. 1240–1246, Mar. 2001.
- [155] E. K. Sackmann, A. L. Fulton, and D. J. Beebe, "The present and future role of microfluidics in biomedical research," *Nature*, vol. 507, no. 7491, pp. 181–189, Mar. 2014.
- [156] P. Holl, V. Koltun, and N. Thuerey, "Learning to control PDEs with differentiable physics," 2020, *arXiv:2001.07457*.
- [157] D. McIntyre, A. Lashkaripour, P. Fordyce, and D. Densmore, "Machine learning for microfluidic design and control," *Lab a Chip*, vol. 22, no. 16, pp. 2925–2937, 2022, doi: [10.1039/d2lc00254j](https://doi.org/10.1039/d2lc00254j).
- [158] M. Wehner, R. L. Truby, D. J. Fitzgerald, B. Mosadegh, G. M. Whitesides, J. A. Lewis, and R. J. Wood, "An integrated design and fabrication strategy for entirely soft, autonomous robots," *Nature*, vol. 536, no. 7617, pp. 451–455, Aug. 2016.
- [159] L. Bocquet, "Nanofluidics coming of age," *Nature Mater.*, vol. 19, no. 3, pp. 254–256, Mar. 2020.
- [160] P. Robin, N. Kavokine, and L. Bocquet, "Modeling of emergent memory and voltage spiking in ionic transport through angstrom-scale slits," *Science*, vol. 373, no. 6555, pp. 687–691, Aug. 2021. [Online]. Available: <https://www.science.org/doi/abs/10.1126/science.abf7923>
- [161] H. Daiguji, Y. Oka, and K. Shirono, "Nanofluidic diode and bipolar transistor," *Nano Lett.*, vol. 5, no. 11, pp. 2274–2280, Nov. 2005.
- [162] N. Kavokine, M.-L. Bocquet, and L. Bocquet, "Fluctuation-induced quantum friction in nanoscale water flows," *Nature*, vol. 602, no. 7895, pp. 84–90, Feb. 2022.
- [163] P. Robin, T. Emmerich, A. Ismail, A. Niguès, Y. You, G.-H. Nam, A. Keerthi, A. Siria, A. K. Geim, B. Radha, and L. Bocquet, "Long-term memory and synapse-like dynamics in two-dimensional nanofluidic channels," *Science*, vol. 379, no. 6628, pp. 161–167, Jan. 2023. [Online]. Available: <https://www.science.org/doi/abs/10.1126/science.adc9931>
- [164] A. Ismail, G.-H. Nam, A. Lokhandwala, S. V. Pandey, K. V. Saurav, Y. You, H. Jyothilal, S. Goutham, R. Sajja, A. Keerthi, and B. Radha, "Programmable memristors with two-dimensional nanofluidic channels," *Nature Commun.*, vol. 16, no. 1, p. 7008, Jul. 2025.
- [165] J. Luo, A. Remy, and Y. Zhang, "Iontronic devices from biological nanopores to artificial systems: Emerging applications and future perspectives," *Chem. Rev.*, vol. 125, no. 24, pp. 11840–11877, Dec. 2025. [Online]. Available: <https://pubs.acs.org/doi/10.1021/acs.chemrev.5c00579>
- [166] T. M. Kamsma, J. Kim, K. Kim, W. Q. Boon, C. Spitoni, J. Park, and R. van Roij, "Brain-inspired computing with fluidic iontronic nanochannels," *Proc. Nat. Acad. Sci. USA*, vol. 121, no. 18, Apr. 2024, Art. no. 2320242121.
- [167] S. Stepney, "Physical reservoir computing: A tutorial," *Natural Comput.*, vol. 23, no. 4, pp. 665–685, Dec. 2024.
- [168] Y.-T. Xu, S.-Y. Yu, Z. Li, B.-H. Kou, J.-X. Pang, W.-W. Zhao, H.-Y. Chen, and J.-J. Xu, "A nanofluidic spiking synapse," *Proc. Nat. Acad. Sci. USA*, vol. 121, no. 28, Jul. 2024, Art. no. 2403143121. [Online]. Available: <https://www.pnas.org/doi/abs/10.1073/pnas.2403143121>
- [169] D. Drotman, S. Jadhav, D. Sharp, C. Chan, and M. T. Tolley, "Electronics-free pneumatic circuits for controlling soft-legged robots," *Sci. Robot.*, vol. 6, no. 51, Feb. 2021, Art. no. eaay2627.



**STEFAN FISCHER** (Member, IEEE) received the Diploma degree in information systems and the Ph.D. degree in computer science from the University of Mannheim, Germany, in 1992 and 1996, respectively. After a postdoctoral year, he joined the University of Montreal, Canada. He held positions at International University, Germany, as an Assistant Professor, and the Technical University of Braunschweig as an Associate Professor, until he joined the University of Lübeck, Germany, in 2004. He is currently a Full Professor of computer science with the University of Lübeck and the Director of the Institute for Telematics. His current research interests include network and distributed system structures, such as ad hoc and sensor networks, the Internet of Things, smart cities, and especially molecular communications on the nano level based on DNA structures. He has (co-)authored more than 220 scientific books and articles.



**NIHAT AY** received the Ph.D. degree in mathematics from Leipzig University, in 2001. In 2003 and 2004, he was a Postdoctoral Fellow with the Santa Fe Institute, NM, USA, and the Redwood Neuroscience Institute (now the Redwood Center for Theoretical Neuroscience at UC Berkeley). From 2005 to 2021, he was a Max Planck Research Group Leader with the Max Planck Institute for Mathematics in the Sciences, Leipzig, where he headed the Information Theory of Cognitive Systems Group. He is currently a Full Professor with Hamburg University of Technology, Germany, where he heads the Institute for Data Science Foundations. He is an External Faculty Member of the Santa Fe Institute and the Complexity Science Hub, Vienna, Austria. He works on information geometry and its applications to embodied intelligence and machine learning. He serves as the Editor-in-Chief for *Information Geometry* (Springer).



**OLAF LANDSIEDEL** received the Ph.D. degree from RWTH Aachen University, in 2010. He held a postdoctoral position with the KTH Royal Institute of Technology, Stockholm, Sweden. He was a Fulbright Scholar in USA, from 2002 to 2004. He is currently a Full Professor of computer science with Hamburg University of Technology (TUHH), Germany, where he heads the Institute for Networked Cyber-Physical Systems (NCPS). Before joining TUHH, he held faculty positions with the Chalmers University of Technology, Gothenburg, Sweden, and was a Full Professor with the Distributed Systems Group, Kiel University, Germany. His research focuses on distributed and networked systems, with particular emphasis on the Internet of Things (IoT), cyber-physical systems, TinyML, edge AI, and edge and fog computing.



**ESFANDIAR MOHAMMADI** received the Ph.D. degree from Saarland University, in 2015. He was with ETH Zürich as a ZISC Fellow, from 2016 to 2019, and joined the University of Lübeck (UzL), Germany, in 2019. Since 2022, he has been the Director of the AnoMed competence center, which focuses on data privacy and secure computation for medical applications and is one of the nation-wide three competence centers on data privacy. He is currently a Professor with UzL, where he heads the Privacy and Security Group. His research focuses on privacy-preserving machine learning, distributed learning, data privacy, backdoor injection in machine learning, and anonymous communication networks. Since 2018, he has served for the technical PCs of various international conferences, such as ACM CCS, USENIX Security, and PoPETS.



**SEBASTIAN OTTE** received the Ph.D. degree in AI and computer science from the University of Tübingen, in 2017. Subsequently, he conducted postdoctoral research in the Neuro-Cognitive Modeling Group, University of Tübingen, from 2017 to 2023. In 2020, he was an Acting Professor with the Chair of Distributed Intelligence. He was a Feodor Lynen Fellow and a Visiting Scientist with CWI, Amsterdam, from 2022 to 2023. He assumed his current professorship in Lübeck, in September 2023. He is currently a Professor with the University of Lübeck, Germany, where he heads the Research Group for Adaptive Artificial Intelligence. His research focuses on flexible, continuously learning AI systems that operate efficiently under constraints, such as limited energy, memory, and computational resources. The work combines artificial intelligence, cognitive science, and robotics in an interdisciplinary framework.



**BERND-CHRISTIAN RENNER** was an Assistant Professor with Hamburg University of Technology (TUHH), from 2016 to 2020. From 2020 to 2022, he was an Associate Professor with the University of Koblenz and Landau. He is currently a Full Professor and the Head of the Institute for Autonomous Cyber-Physical Systems (ACPS), TUHH. His research interests include applications of networked low-power sensing and cyber-physical systems running on energy harvested from the environment. He is also active in networking protocols for networks of embedded sensors, active and passive backscatter acoustic communication, localization, and navigation in mobile, low-power underwater networks. He has published papers in several international journals and conferences, and reviewed for and served on several TPCs. He is a member of its steering committee. He has served the ENSsys Workshop at ACM SenSys in several roles, since 2014.



**NELE RUSSWINKEL** received the Ph.D. degree in framework from the DFG-funded Research Training Group “Prometei” (Prospective Human Machine Interaction), TU Berlin. She is currently a Full Professor with the University of Lübeck, Germany, for human-aware AI, and the Head of the Institute for Information Systems. Her research focuses on cognitive systems and intelligent agents interacting in dynamic environments. She is also working on cognitive robotics and the question of how to provide a sense of control for such systems in dynamic and partly unknown environments. She was holding a Junior Professorship with TU Berlin and got the call to Lübeck, in 2022. She is the President of the International Society for Cognitive Modeling and the Vice President of German Cognitive Science Community (GK e.V.).

• • •