

Oncomorphic computing: a bio-derived paradigm for decentralized adaptation across software and hardware

Rosalia Moreddu^{1*}, Philip Greulich², Christian Baumgartner³, Francesca Borghi⁴, Michael Levin⁵

¹School of Electronics and Computer Science, University of Southampton, Southampton, UK

²School of Mathematical Sciences, University of Southampton, Southampton, UK

³Institute of Health Care Engineering with European Testing Center of Medical Devices, Graz University of Technology, Graz, Austria

⁴Department of Physics and Interdisciplinary Centre for Nanostructured Materials and Interfaces (CIMAINE), University of Milan, Milan, Italy

⁵Department of Biology, Tufts University, Medford (MA), USA

Correspondence: rosalia.moreddu@eng.ox.ac.uk

Keywords: bioinspired computing; decentralization; adaptation; nonlinear systems; heterogeneous specialization; edge AI; digital health; autonomous systems; embodied robotics; multi-agent systems; *in vitro* bioelectronics.

Abstract

Contemporary artificial intelligence excels in stable and resource-rich settings, but many of the environments in which autonomy and adaptive restructuring are most needed are defined by noise, damage, scarcity, and unpredictability. Spacecraft, distributed sensors, medical implants, digital cell twins, and robotic swarms must often continue functioning despite partial hardware failure and network fragmentation. Here, we present oncomorphic computing as a paradigm for designing resilient machine intelligence across software and hardware. We abstract selected survival principles seen in cancer biology, particularly decentralized behavior, phenotypic plasticity, heterogeneous specialization, resource recruitment, migration of function, and tissue-level adaptation, and repurpose them bidirectionally to complement standard approaches: using cancer dynamics in computing and engineering, and, conversely, bringing biological decentralization to the center of digital cancer research. In this way, we shift the approach from a highly organized and predictable ideal of machine intelligence to computational systems defined as groups of agents continuously negotiating access to memory, power and processing bandwidth through independent restructuring under scarcity. We discuss how this paradigm could address current computing bottlenecks and outline its implications for future hardware applications.

1. Introduction

Artificial intelligence is increasingly being deployed in environments where the assumptions behind its success do not hold^{1, 2}. Large models and autonomous systems are developed on stable hardware, reliable communication, large power, and the possibility of supervision or recovery^{3, 4}. Yet many of the settings in which autonomy matters, including deep-space platforms, implantable medical devices, digital medicine, distributed sensor networks, and robotics, often operate under the opposite conditions⁵⁻⁸. They must tolerate radiation-induced faults, intermittent connectivity, and adversarial interference while continuing to act in real time^{9, 10}. In these settings, accuracy is only one component of intelligence. The more basic requirement is the ability to preserve reliability and longevity when infrastructure, data flow, and even the computational substrate begin to fail⁹. Existing approaches have addressed parts of this problem. Fault-tolerant computing provides redundancy, rollback, and error correction at the cost of computational overhead and resource consumption¹¹. Federated and edge learning distribute data collection and inference, yet they still rely on complex network architectures and resource-constrained edge devices¹². Swarm robotics exploits local coordination, and reconfigurable hardware enables remapping after damage¹³. These strategies are valuable, but they often presume that failures or failure types are known in advance, isolable, or recoverable through predefined mechanisms¹⁴. They are less suited to conditions in which degradation is sudden, resources are competitive, communication is delayed, and adaptation must emerge locally before any diagnosis is possible.

Biology offers models of persistence under stress. Most bioinspired computing has drawn from neural, developmental, and evolutionary systems^{15, 16}. Cancer provides a different and underused example of adaptation^{1, 17}. Tumors persist without central control, maintain heterogeneous subpopulations, allocate local resources, and reorganize function under scarcity¹⁸⁻²⁰. Cancer has been also described as a loss of connections into a large-scale collective that acquires new capabilities when it disbands²¹. From a computing and engineering perspective, abstracting decentralized survival principles from cancer biology, for instance phenotypic plasticity, selectivity, and migration of function, might enable to address specific bottlenecks encountered in machine learning and, consequently, impact hardware deployment in a range of use cases, from computational biology to humanoid robot interaction. We recently demonstrated how basic cancer properties can be integrated into computational neural agents for resource-limited sequential learning¹⁷. At the same time, the presence of quantifiable electric behavior in cancer²² and non-excitable cell populations²³, epigenetic states,²⁴ and mechanical adaptation²⁵ in metastatic cancer cells makes a standardization of decentralized dynamics, and a hardware implementation frontier, thinkable. We call this framework oncomorphic computing (**Figure 1**). In an oncomorphic system, computation is distributed across local agents that continuously assess their own state and negotiate access to memory, power, communication, and bandwidth. Under normal conditions, agents cooperate to maximize performance. Under severe stress, they can isolate compromised regions, reassign roles and consolidate critical functions. Therefore, oncomorphic computing represents a resilience-oriented design lens for software, hardware, and embodied machines whose defining challenge is continued usefulness in uncertainty^{1, 17}.




Figure 1. The oncomorphic computing concept. (a) Baseline architectural model, featuring a fully decentralized population of autonomous agents operating within a finite local resource ecosystem. (b) Dynamic adaptation response under catastrophic stress, demonstrating asymmetric node recruitment where surviving agents absorb memory, power, and compute pools from damaged neighbors to maintain critical system function.

1.1. High-level formalization of the paradigm

We provide a high-level formalization of this paradigm to distinguish it from standard fault tolerance, in which redundancy and recovery are usually predefined¹¹, and to provide a common language across the very different systems discussed later in this perspective. From a general standpoint, an oncomorphic system can be represented as a constrained decentralized control problem on a time-varying graph²⁶

$$G(t) = (V, E(t))$$

Each vertex $i \in V$ is a local unit, i.e. a software process, a hardware tile, a sensor node, a robot subsystem, or, in a computational biology setting, an individual cell agent. Each edge $i, j \in E(t)$ indicates a pathway whose availability may itself vary in time because of congestion, damage or noise²⁷. This representation is adopted because topology is not fixed in the systems of interest, and the ability to continue functioning depends on the capacity to update internal state. A minimal state description for agent i is

$$z_i(t) = [x_i(t), h_i(t), \tau_i(t), r_i(t), \sigma_i(t)]$$

where $x_i(t)$ is the internal task state of the agent (i.e., model weights, cached data, control variables, local memory of recent events, or, in digital twin applications, biological state variables such as proliferation), where the future response depends on local history. The scalar $h_i(t) \in [0,1]$ represents health, i.e. the remaining ability of a unit to contribute useful task. A continuous health variable is preferable to an alive/dead label because the failure modes are progressive²⁸. The scalar $\tau_i(t) \in [0,1]$ represents confidence, which differs from $h_i(t)$ because a unit may still be operational in a physical sense when its outputs become unreliable because of uncertainty²⁹. The vector

$$i(t) = [m_i(t), p(t), b_i(t)]^T$$

collects local allocable resources, indicated as memory m_i , power p_i , and bandwidth b_i , where bandwidth includes communication throughput and compute capacity. These were chosen because they recur across all the presented application domains. Finally, $\sigma_i(t)$ is a discrete mode, such as active, relay, dormant, or migrating. This variable is the formal analogue of phenotypic plasticity. Because no agent has perfect global knowledge, the system is naturally partially observable³⁰. Each unit acts on a local observation:

$$o_i(t) = H_i(z_i(t), \{z_j(t): j \in N_i(t)\}) + \nu_i(t),$$

where $N_i(t)$ is the current neighborhood and $\nu_i(t)$ represents observation noise or uncertainty. A local control policy π_i then selects an action $u_i(t) = \pi_i(o_i(t))$, yielding the update

$$z_i(t + \Delta t) = F_i(z_i(t), u_i(t), \{z_j(t): j \in N_i(t)\}, \xi_i(t)),$$

Where $\xi_i(t)$ indicates external perturbation. The equation is meant to signify that oncomorphic adaptation does not require global optimization at every instant, rather that each unit estimate enough of its own state, local stress, and neighborhood condition to update in a manner relevant to the system goal. The variable $h_i(t)$ must itself evolve under damage. An update law is

$$h_i(t + \Delta t) = \Pi_{[0,1]}(h_i(t) - \sigma_D D_i(t) - \alpha_L L_i(t) + \alpha_S S_i(t)),$$

Where $\Pi_{[0,1]}$ projects onto the interval $[0,1]$, $D_i(t)$ is accumulated damage, $L_i(t)$ is induced stress, and $S_i(t)$ is compensatory support. This form was chosen because it separates three different mechanisms: damage $D_i(t)$ captures direct physical or logical damage, load $L_i(t)$ captures the workload, support $S_i(t)$ captures the core oncomorphic claim that neighboring units can partially offset local decline by lending capacity, assuming load, rerouting traffic, or inheriting state⁸. The meaning of $S_i(t)$ depends on the application: in software it may be checkpoint migration, in hardware it may be rerouted logic, in robotics it may be redistributed control, in digital twins it may correspond to parameter reassignment within the model.

Resource recruitment is modeled as a core differentiator between oncomorphic computing from classical degradation. For each resource class $\ell \in \{m, p, b\}$, $T_{ij}^{(\ell)}(t) \geq 0$ indicates the amount transferred from agent i to agent j during $[t, t + \Delta t]$. Resource dynamics can be written as

$$r_i^{(\ell)}(t + \Delta t) = r_i^{(\ell)}(t) + \sum_{j \in N_i(t)} T_{ji}^{(\ell)}(t) - \sum_{j \in N_i(t)} T_{ij}^{(\ell)}(t) - c_i^{(\ell)}(u_i(t)) - \delta_i^{(\ell)}(t),$$

subject to transfer-capacity constraints

$$0 \leq T_{ij}^{(\ell)}(t) \leq C_{ij}^{(\ell)}(t),$$

where $C_{ij}^{(\ell)}(t)$ is the instantaneous capacity of the link and $\delta_i^{(\ell)}(t)$ represents irreversible loss. This equation encodes that adaptation reallocates resources but does not create them. Specific domains may add other terms including thermal, actuation, and storage.

The system must also represent tasks and their unequal importance, which also varies over time in the oncomorphic approach²⁸. Let $K = \{1, \dots, K\}$ be the set of system functions, with $a_{ik}(t) \in \{0,1\}$ indicating whether agent i currently supports task k . A local contribution score is

$$q_{ik}(t) = a_{ik}(t)h_i(t)\tau_i(t) \min_{\ell \in \{m,p,b\}} \left(1, \frac{r_i^{(\ell)}}{\bar{r}_k^{(\ell)}} \right),$$

Where $\bar{r}_k^{(\ell)}$ is the nominal resource requirement of task k . This definition couples function delivery to the ingredients that later applications depend on. The delivered quality of task k can then be written as

$$y_k(t) = \min \left(1, \sum_{i \in V} q_{ik}(t) \right),$$

and the overall system viability can be defined as

$$V(t) = \sum_{k=1}^K w_k y_k(t), \quad \sum_{k=1}^K w_k = 1,$$

Where w_k is the mission weight of task k . This weighted viability formulation is preferable because oncomorphic systems are designed for unequal functions. Examples are a rover dropping non-essential imaging to preserve localization, or a wearable sensor reducing sampling density to maintain arrhythmia detection. For a set of critical functions K_{crit} , one can impose lower bounds

$$y_k(t) \geq \underline{y}_k, \quad k \in K_{crit},$$

to formalize the idea of protected mission backbone introduced later. Local plasticity must remain bounded, otherwise adaptation becomes indistinguishable from uncontrolled self-modification. A basic quarantine rule is

$$\tau_i(t) < \theta_\tau \quad \vee \quad h_i(t) < \theta_h \quad \rightarrow \quad \sigma_i(t + \Delta t) = \textit{quarantine},$$

meaning that units falling below integrity thresholds may be isolated from shared state propagation, actuation, or write operations³¹. This is especially important in cybersecurity, medical robotics, and closed-loop bioelectronic systems, where a degraded but still connected component may do more harm than a silent one¹⁰. Conversely, mode changes such as relay, reserve, or dormant are the controlled forms of functional plasticity that allow the system to preserve critical work while reducing burden elsewhere. Because later sections involve safety-critical settings, an external safe-mode condition can be defined as

$$s_{ext}(t) = 1 \quad \rightarrow \quad \sigma_i(t + \Delta t) = \textit{safe} \quad \forall i \in V.$$

At the system level, the design objective is maximal retained utility. This is represented as

$$\max_{\{\pi_i\}} \mathbb{E} \left[\sum_{t=0}^T \gamma^t \left(V(t) - \lambda_E E_{adapt}(t) - \lambda_M M_{mig}(t) \right) \right],$$

where $E_{adapt}(t)$ is the energy cost of adaptation, $M_{mig}(t)$ comprehends migration, replication, and rerouting, and $\lambda \in (0,1]$ is a discount factor. Implementations may approximate this objective through model predictive control or safe reinforcement learning under safety monitor³².

This formalization introduces when the paradigm is deployable in the domains discussed later. In distributed software, the approach is implementable as an orchestration layer. In hardware, the same variables can be instantiated on field-programmable gate array (FPGA). In computational oncology and digital twins, the agents can be deployed in simulation, calibration, and prediction. For testing, three metrics are particularly useful: the retained mission utility relative to nominal baseline V_0 , the time required to recover a fraction ρ of pre-fault utility after a fault at time t_f , and the energy overhead of adaptation relative to baseline operation. These are represented as

$$RMU = \frac{1}{T} \int_0^T \frac{V(t)}{V_0} dt,$$

$$T_{rec} = \{t > t_f : V(t) \geq \rho V_{pre}\},$$

$$\eta_{adapt} = \frac{\int_0^T E_{adapt}(t) dt}{\int_0^T E_{base}(t) dt}.$$

A further safety metric is defined as

$$N_{unsafe} = \sum_{t=0}^T 1[x(t) \notin S],$$

where S is the safe operating set³³. These measures support fault-injection experiments in the presented application areas³⁴. This representation identifies the variables that matter, how damage and support alter local viability, makes resource competition explicit, and defines when adaptation is allowed.

1.2. Comparison with adjacent computing paradigms

Decentralization, local adaptation, reconfiguration, and embodied intelligence are all established ideas³⁵⁻³⁷. The novelty of oncomorphic computing lies in redefining what they are designed to preserve: under uncertainty, the system utility precedes task accuracy to prioritize longevity, adaptation, and continued functionality. In **Section 1.1**, we have shown how adjacent frameworks optimize task state x_i under stable infrastructure assumptions³⁸. Oncomorphic computing, on the other hand, treats health h_i , integrity τ_i , allocable resources r_i , and functional role σ_i as first-class state variables, and evaluates success through the viability $V(t)$ retained under attrition, complementarily to classical digital computing, a.k.a. von Neumann architectures (**Table 1**). Conventional processor-memory systems are built around exact execution on hardware assumed to be stable, with faults handled as exceptions³⁹. Large-scale machine learning builds on the same assumptions. These are appropriate in data-center environments, but become brittle when communication is delayed, the substrate itself degrades, or intervention is not possible⁴⁰. Here, useful operation may require model simplification, priority of tasks, or relocation of function. The comparison with distributed computing, edge AI, and federated learning is closer, because these already move away from centralization^{12, 38}.

However, they usually assume nodes that are broadly stable and failures that fit known protocols. Extending the principles of classical swarm intelligence, oncomorphic computing explicitly models heterogeneous agent health, resource competition, and functional inheritance during progressive degradation. Swarm intelligence and decentralized multi-agent control provide another similarity. Classical swarms achieve robustness through local rules and homogeneous agents¹⁴.

Neuromorphic computing⁴¹ offers a conceptual precedent that mimics the operating principles of neurons and synapses, such as spiking, event-driven communication, local plasticity, co-located memory and computation, and import them in conventional technology to achieve energy-efficient, adaptive information processing. However, it also implicitly assumes that the underlying silicon substrate itself remains physically stable. Memristive devices support in-memory computing⁴², addressing the processor-memory bottleneck of von Neumann architectures. Unconventional devices, such as those based on electrochemical organic materials⁴³ and self-assembled systems⁴⁴, have been demonstrated as analog memories and in-material processing units. Oncomorphic computing, instead, imports the operating principles of tissue and cellular biology under stress and damage to manage degradation of the substrate itself, without a central control unit. Where neuromorphic systems optimize how information is learned and represented, oncomorphic systems ask how a population of physically compromised units can remain viable to accomplish core mission goals. Oncomorphic agents, thus, must retain memory of prior stress, differ in role, and participate in a population-level economy. The degradation of one unit may trigger inheritance of key state, reweighting of neighboring agents, or isolation of a node that remains physically present but has lost temporary priority.

A different comparison is with continual learning and evolutionary computation. Continual learning addresses how to update a model without erasing established capabilities, and evolutionary algorithms adapt populations of candidate solutions across generations⁴⁵. Oncomorphic computing is compatible with both, but it looks at reorganization of a system as resources shrink. In this setting, catastrophic forgetting, i.e. the forgetting of previous tasks after learning successive ones, becomes an architectural question on (i) which modules contain essential function and remain plastic, (ii) when a useful local adaptation should be replicated, and (iii) when a new capability should be transferred^{17, 45}. Competitive learning offers a further point of comparison, since it also produces role differentiation among initially homogeneous units through local, decentralized dynamics rather than central control. In winner-take-all schemes such as self-organizing maps, units compete to represent a given input, and the winner updates its weights while other units remain largely unchanged, so that different units come to specialize in different regions of input space. Neuromorphic and adaptive model of classification, based on non-linear elements, has been recently introduced in this learning context to address catastrophic forgetting on simple architectures. Oncomorphic agents undergo an analogous process of differentiation, but the competition is waged over physical survival rather than representational accuracy: agents compete for allocable resources r_i under conditions of degradation. Other close concepts are autonomic computing and organic computing, which also propose self-configuration³⁶. On the hardware side, the most relevant

comparison is with reconfigurable computing. Reconfigurable hardware already supports logic remapping through FPGAs, reconfigurable arrays, and network-on-chip⁵. Neuromorphic hardware relies on event-driven operation, local memory-compute, and distributed plasticity, making it attractive for low-power systems, but its objective is efficient sensing and predictability¹⁵. Finally, distributed working subregions that preserve critical function when the substrate becomes unreliable, proposed in the oncomorphic paradigm, also align with principles of morphological computation and embodied AI, both of which established that intelligence cannot be separated from its physical adaptive substrate¹⁶.

Table 1. Comparison of oncomorphic computing with adjacent computing paradigms.

<i>Paradigm</i>	<i>Primary design objective</i>	<i>Principal unit/state</i>	<i>Control and adaptation structure</i>	<i>Uncertainty model</i>	<i>Resource treatment</i>	<i>Preserved quantity</i>	<i>Relation to OC</i>
Oncomorphic computing	Mission utility under uncertainty, subject to safety bounds	Local agent with task state, health, resources, and role	Decentralized local policies: migration of task, and self-restructuring	Progressive damage, node loss, drift, and resource depletion	Memory, power, and bandwidth are redistributable	Viability of critical functions over state preservation	Proposed framework
Classical digital computing (von Neumann)	Execute programs efficiently on a stable processor-memory	Instruction stream, process, thread, address space	Centralized scheduling; adaptation at design, compile, or restart time	Faults treated as exceptions, recovery through restart or replacement	CPU, memory, storage, and I/O scheduled under stable hardware assumptions	Program correctness, state consistency, and exact output	OC shifts target to preservation of valuable functions
Distributed, edge, and federated AI	Scale computation, reduce data movement, preserve availability	Node, service, shard, local model, replica, or edge device	Protocol-based coordination with partial local autonomy	Crash faults, partitions, delay; less emphasis on progressive shift	Bandwidth, latency, storage, replication, and local compute	Service continuity, utility, convergence, data availability	OC adds health-aware node valuation and function inheritance
Evolutionary computation	Search for solutions by variation, selection, and retention	Genome, candidate program, controller, or rule set	Population-level adaptation across generations, often offline	Low-fitness variants are discarded. Not a model of damage	Budget, diversity, mutation rate, and selection are main resources	Fitness progress, novelty, or a quality-diversity poll	OC adapts within the lifetime of a deployed system under stress
Swarm intelligence and decentralized multi-agent control	Generate collective behavior from local decentralized interaction	Agent with local state, rules, and limited memory	Fully decentralized coordination, built from homogeneous agents	Individual loss tolerated statistically. Integrity and criticality less explicit	Local energy, position, and neighborhood information dominate	Coverage, transport, exploration, formation or consensus	OC has richer internal state, heterogeneity, competition, function preservation
Neuromorphic computing	Low-power sensing, inference, control via distributed dynamics	Neuron, synapse, event core, or spiking behavior	Asynchronous parallel local dynamics with fast temporal response and plasticity	Noise and variability tolerated, but not organized around uncertainty	Sparse signaling and local memory-compute minimize energy use	Low-power temporal inference, perception, and control	Substrate for oncomorphic policies, with goal shift to persistence of key function
Autonomic, self-adaptive, and organic computing	Runtime self-management/ configuration/ healing, and adaptation	Managed component, monitors, models, knowledge	Monitor-analyze-plan-execute loops, often hierarchical	Faults, policy violations, environmental shifts trigger adaptation	Metrics, budgets, and fault statistics guide adaptation	Availability, quality of service, and policy compliance	OC pushes adaptation deeper into local agents
Fault-tolerant computing	Preserve correctness and availability through redundancy	Replica, protected memory region, or spare module	Recovery against predefined fault models via error correction	Discrete faults and known fault classes are handled explicitly	Spare capacity, timing, and redundancy reserved in advance	Correctness, availability, continuity of exact service where possible	OC builds on fault tolerance but permits function migration

Reconfigurable computing	Remap logic to changing workloads or defective substrate	Logic block, routing segment, memory tile	Remapping can occur from runtime to deployment lifetime	Hardware defects, routing, and performance motivate remapping	Logic area, bandwidth, power and non-volatile state are key resources	Capability and throughput on an altered substrate	OC couples remapping to integrity, resource scarcity, and broad mission
Morphological computation and embodied AI	Exploit body mechanics, compliance, environment as part of computation	Body segment, joint, geometry, material field	Distributed through physical dynamics plus local control	Perturbations absorbed via compliance and body-environment interaction	Geometry, compliance, actuator limits, and passive dynamics are resources	Useful embodied behavior	OC adds state, health, trust, and resource redistribution, under damage

2. Addressing Bottlenecks in Machine Learning

Recent progress in machine learning (ML) has been driven largely by scale: larger datasets, larger models, longer training runs, and increasingly centralized compute infrastructure³⁴⁻³⁶. This strategy has delivered impressive results, but it also reflects the core assumption that training and deployment happen on hardware that is stable, well connected, and energy-rich, and that failures can be handled by restart, resynchronization, or human intervention in an unlimited resource scenario^{2, 10, 36, 38, 39}. Many of the systems discussed in this perspective do not satisfy those assumptions. These include space platforms, distributed sensor networks, implantable devices, humanoid robots, and autonomous systems which must often continue to perform while communication is intermittent, energy demand is massive and its resources are limited, or parts of the substrate are degraded^{5, 9, 11}. This mismatch highlights several bottlenecks in current machine learning. First, many models remain operationally dependent on centralized coordination for training and updating³⁸. Second, continual adaptation still struggles with interference between old and new knowledge, particularly when learning must occur online⁴⁵. Third, large populations of agents remain difficult to coordinate without introducing control layers⁴⁶. Oncomorphic computing conceptually complements classical ML in these problems by potentially offering a runtime architecture for preserving useful function when the computational environment becomes unstable. **Figure 2** summarizes this idea at the software level. A local disruption creates a degraded computational region, triggering nearby agents to reallocate available memory and communication capacity toward the affected area or toward functions that have become newly critical. If the damaged region cannot be recovered, the system shrinks to a critical subset of tasks in healthier parts of the network. In the next subsections we introduce relevant areas of potential development.

2.1. Centralized computational overhead

A major ML bottleneck is the dependence on centralized infrastructure². Even when inference is pushed to the edge, the surrounding system often still relies on remote coordination for model updates, task scheduling, and fallback⁴⁰. This becomes problematic in systems operating under long latency, sparse bandwidth, or intermittent contact⁴⁷. In those cases, moving data to a central processor, synchronizing state across nodes, or recovering from a failed server may cost more time and energy than the task itself⁴⁸. The primary value of an oncomorphic architecture emerges at deployment. It suggests breaking monolithic software stacks into interacting

modules or agents with partial state, local health estimates, and bounded autonomy over specific functions. In practice, this could be implemented through health-aware task placement, parameter-efficient local adaptation, selective migration, event-triggered synchronization, and temporary clustering of specialized modules around a critical task⁴⁹. Existing machine learning tools already support parts of this approach. Mixture-of-experts routing, modular policies, adapter-based fine-tuning, compressed model exchange, and hierarchical service orchestration all make it possible to shift computation without fully retraining or moving an entire model^{47, 48}. For many applications centralized training remains the most effective way to build high-performing models. The issue is what happens after deployment, when communication degrades or hardware fails. In those settings, a resilient system should be able to simplify its own computation. A robot could switch from a large perception stack to a smaller verified model handled temporarily, a wearable device may suspend nonessential analytics, a spacecraft may concentrate remaining compute on navigation and fault diagnosis. From this perspective, resource recruitment means reallocating scarce compute time, memory budget, and communication opportunities toward the functions with the highest current mission value. This is less efficient under ideal conditions, but far more robust under degraded or adversarial settings. For this reason, this concept is meant to complement standard methods to enhance system capabilities.




Figure 2. Software-level expression of oncomorphic adaptation under localized failure. (a) A local disruption reduces the health and connectivity of a subset of agents, creating a degraded computational region. (b) Neighboring agents reallocate available memory, compute, and communication capacity toward the affected area or toward tasks of higher current priority. (c) If local recovery is not feasible, the system isolates low-integrity components and consolidates essential state and functions within a smaller protected cluster, preserving system-level continuity through bounded reconfiguration.

2.2. Catastrophic forgetting

Catastrophic forgetting remains one of the key limitations of continual learning⁴⁵. When a model is updated on new data or new tasks, previously acquired capabilities can be overwritten, especially when the same parameters must meet conflicting objectives over time^{45, 50}. Existing methods address this problem at the level of parameter dynamics: replay buffers preserve representative experiences, regularization methods constrain movement away from previously

important weights, and modular approaches try to isolate task-specific updates^{51, 52}. In oncomorphic computing, instead of forcing all adaptation into a single shared parameter space, heterogeneity is treated as a computational resource. Some modules can be designated as relatively stable repositories of high-value capabilities, while others can remain more plastic and exploratory. New information is then not required to overwrite established function immediately. It can first be encoded in local adapters, and only later promoted if it proves useful and safe. In this sense, forgetting becomes not only a problem of optimization, but also one of allocation dealing with where new information should be written, which components should remain protected, and under what conditions should local adaptations be replicated. This view is most relevant in safety-critical or long-lived systems. A robot, implant, or autonomous spacecraft should not have to sacrifice validated core behavior every time it encounters a new operating condition⁹.

Receptron intrinsic input-dependent selectivity⁵³ offers an example of plausible computational primitive for the "local adapter" role described in the oncomorphic computing paradigm. Because a single receptron unit can carve out selective activation only within a bounded region of the input space, exhibiting intrinsic selective activation properties for analog inputs when the input vector lies within cubic domains in a 3D space, it can, in principle, absorb new information in a spatially and functionally circumscribed way, rather than perturbing the entire shared parameter space. A receptron localized selectivity is the kind of mechanism that would let a system write new information into a contained unit without immediately overwriting the function of other units, a prerequisite for treating heterogeneity as a resource rather than a liability. A more resilient design would separate a protected backbone of essential competencies from a peripheral layer that can adapt online under tighter local constraints. Mechanisms already familiar in machine learning, such as modular networks, sparse expert routing, episodic memory, fine-tuning, and validation-gated model updates, fit this scheme. This reduces the degree to which all learning pressure is concentrated on the same vulnerable parameters.

2.3. Scalability and coordination

As machine learning systems become more distributed, coordination becomes a bottleneck⁵⁴. Multi-agent systems, edge inference networks, and autonomous sensor populations all face the same tension: local components must retain enough independence to remain robust, but enough coherence must be preserved for system-level behavior to remain useful^{54, 55}. Classical solutions rely on consensus protocols, parameter servers, central aggregators, or frequent state exchange⁵⁵. These mechanisms become expensive or fragile when nodes differ substantially in capability. In this context, the oncomorphic paradigm treats coordination as something that should depend on local condition, besides task structure. Agents do not need full agreement on global state at every step if they can respond to a shared set of bounded signals⁵⁶. This favors event-triggered communication over continuous synchronization, health-aware routing over uniform load balancing, and selective inheritance of critical state over blanket replication. Under mild stress, such a system may look like conventional decentralized computing. Under severe stress, however, it can shift into triage: low-priority functions are shed, damaged or low-integrity regions are quarantined, and surviving agents concentrate resources around a smaller

protected backbone of services⁵⁷. This is relevant to machine learning systems that must operate under partial compromise^{55, 58}. The main technical challenge is governance. A distributed learning system that can reassign roles can also become opaque if those transitions are not logged. For that reason, oncomorphic coordination should always be paired with explicit supervisory constraints, such as immutable safety rules, limits on replication and migration, and an external safe mode.

3. Adaptive Hardware and Co-Design

If machine learning software based on oncomorphic paradigm is expected to remain useful under damage, scarcity, and uncertain connectivity, the prerequisite of its corresponding hardware system is to possess a dynamic and adaptive internal state, whether the physical system is a single device or an interconnected network. Although demonstrators of pure hardware able to proficiently adapt their internal states by external stimuli modifying have been provided, the dualism between software and hardware cannot be easily overcome for non-living systems aimed at performing complex tasks. The oncomorphic paradigm does not simply bypass the idea of a centralized control system decoupled from the physical state of an expendable substrate; rather, it proposes a control system that serves that substrate, efficiently exploiting and reorganizing a heterogeneous hardware that is itself dynamic (**Figure 3**). In this sense, the concept of programmability of self-reorganizing data processing systems is a virtuous example⁵⁹. In most deployed systems today, resilience is still implemented primarily above the device level: software is used to mask faults occurring in processors and memories⁶⁰. These methods assume that the underlying platform remains sufficiently intact to host the recovery logic^{60, 61}. In harsh environments, this flips. Faults may accumulate in-hardware as routing degradation, thermal hotspots, or sensor drift, turning resilience into a cross-layer problem where hardware–software co-design becomes paramount⁶². In the formalization introduced earlier, the hardware must have at least partial estimates of local h_i , τ_i , and r_i , while the software stack must be able to react by implementing a change (e.g., in routing or timing). This logic is already implemented in some techniques⁶¹. In SRAM-based FPGAs, for example, radiation-induced configuration upsets can be mitigated through error detection and correction and partial dynamic reconfiguration, which allows faulty configuration regions to be refreshed while unaffected logic continues to run⁶⁰. Recent work on aerospace SRAM FPGA mitigation highlights that dynamic partial reconfiguration is useful to correct only the affected configuration, reducing interruption relative to full reconfiguration⁶³.

NASA describes SpaceCube as an in-flight reconfigurable hybrid science data processor that couples commercial radiation-tolerant FPGA technology with integrated upset detection and correction to achieve large improvements in onboard processing relative to traditional rad-hard processors⁶⁴. Likewise, the RadSat concept used a commercial FPGA with real-time partial reconfiguration as a route to radiation-tolerant small spacecraft computing, again illustrating that reconfiguration under fault is already in practice⁶⁵. What remains underdeveloped is how to turn these mechanisms into a more autonomous architecture. Oncomorphic hardware can propose a different control objective for existing and emerging substrates. A useful way to organize this is to think in terms of three hardware layers: (i) reconfigurable digital hardware

with fault-aware management⁶⁶, (ii) memory-centric and near-memory acceleration, where the cost of moving data between memory and compute becomes dominant and can itself be managed as a constrained resource⁴², (iii) emerging non-volatile analog or mixed-signal fabrics, such as memristive crossbars, where adaptation may occur at finer spatial scales but under much stronger non-idealities⁶⁷. We identify the hardware horizons of this paradigm into reconfigurable data processing devices in the short run, and on-device learning and in-memory computing in the longer run.

3.1. Short-term horizon: reconfigurable data processing devices

Reconfigurable digital electronics is the near-term core of oncomorphic hardware. FPGAs, reconfigurable non-linear threshold logic gates⁶⁸, and related devices, already support logic remapping and partial update after faults^{61, 69}. Mission systems rely on combinations of configuration scrubbing, error-correcting memory, and partial reconfiguration to maintain operation⁶⁰. These systems address several mechanisms discussed in *Section 1*. A local region of the fabric can be assigned a lower health score if its upset rate, timing failures, or thermal excursions rise⁶¹. Scrubbing can be interpreted as a form of local repair, spare regions as recruitable substrate, and partial reconfiguration as migration of function⁶⁰. The novel step is the decision policy over that mechanism. Instead of waiting for a complete failure and then restoring the original mapping, an oncomorphic controller could perform graded remapping⁶⁹. In this sense, adaptation means constrained changes in placement, routing, timing, replication level, and service allocation⁷⁰. In edge platforms, industrial autonomy, and biomedical devices, the most realistic path to hardware adaptation is localized observability plus bounded remapping^{70, 71}. A system may detect rising error rates in a sensor interface, increased write failures in non-volatile memory, or excessive power consumption in an accelerator tile, and respond by reducing duty cycle, migrating a task, changing the scheduling order, or switching to a lower-complexity model⁷². In other words, deployable oncomorphic electronics are likely to look less like continuously mutating circuits and more like health-aware reconfigurable systems with explicit supervisory boundaries⁷⁰. Oncomorphic hardware proposes a resilience-oriented design strategy that combines runtime observability, reconfigurable substrates, cross-layer fault mitigation, and mission-aware control of finite hardware resources. Existing FPGA-based flight processors provide several of the enabling ingredients. The challenge is to unify them into systems that continuously reorganize computational function as substrate condition and energy availability evolve.

3.2. Long-term horizon: on-device learning and in-memory computing

Memristive and resistive-switching devices are frequently mentioned in discussions of adaptive hardware because they can combine non-volatile storage with matrix operations performed directly in memory arrays⁷³. The motivation is well established: conventional computing is increasingly limited by the movement of data between physically separate memory and processing units, often described as the von Neumann bottleneck⁴². In-memory computing with resistive switching devices has been presented as a route to area- and energy-efficient computation⁷⁴. Crossbar arrays of resistive devices can implement analog matrix–vector

multiplication, and recent hardware demonstrations have shown increasingly capable compute-in-memory chips based on RRAM⁷⁵.

For oncomorphic computing, however, the significance of memristors is not only energy efficiency, but the fact that computation, storage, and adaptation can be set more locally than in conventional digital systems⁷⁶. This matters because local adaptation becomes expensive if every update requires movement of state across fragile interconnects. A crossbar-based subsystem could, in principle, support local inference, local weight storage, and local remapping around damaged cells⁷⁷. However, contemporary memristive accelerators are not yet straightforward hardware embodiments. Their limitations around device-to-device and cycle-to-cycle variability, nonlinear and asymmetric programming, and substantial peripheral overhead in DAC/ADC impact precision, area, and energy trade-offs at the system level⁷⁸. In some implementations, the peripheral converters can occupy most of the area and energy of the compute unit, sharply reducing the apparent gains of the array itself. This is why memristive hardware is relevant to oncomorphic computing: it shows both the promise and the difficulty of building substrates where physical imperfections must be managed. Recent work has begun to show how algorithmic and architectural adaptation can compensate for imperfect devices. For example, recent work demonstrated an indirect matrix representation that could bypass faulty memristive devices and maintain high functional fidelity⁷⁹. Neuromorphic self-assembled materials with heterogeneous networks, non-volatile memory effects and nonlinear behaviours have been demonstrated, with reprogrammable connectivity and reconfigurable functionality exploitable for data processing tasks^{59, 80}, even on soft materials, paving the way for embodied computation⁸¹. These results shows that fault-aware remapping at the representation level can restore useful computation.




Figure 3. Adaptive hardware horizon. (a) Localized physical degradation and pathway severing of the silicon substrate under severe environmental or cosmic radiation stress. (b) Real-time hardware-level adaptation and autonomous gate recruitment via polymorphic silicon fabrics and memristor crossbars, allowing the coarse-

grained reconfigurable architecture (CGRA) to rewrite physical circuit topologies within a single clock cycle to completely bypass dead silicon zones.

4. Application Frontiers in Computational Biology and Digital Medicine

The oncomorphic paradigm can play a role in computational biology and digital medicine by targeting adaptive modeling and control under shifting information. Depending on the application, an agent i may represent an individual cell, a clonal compartment, a spatial patch of tissue, a specific submodel, an electrode site, a biosensor, or a local state estimator. Its health h_i measures the integrity of the local substrate or the validity of the local model, its confidence τ_i measures how much trust should be placed in its current output, its resources r_i represent finite compute, memory, power, communication, or sampling budget, and its mode σ_i determines whether it contributes at full fidelity, operates in a reduced form, relays information, or is quarantined (see *Section 1.1*). Here, we discuss potential applications in computational oncology, digital cell twins, and closed-loop bioelectronics (**Figure 4**).

4.1. Digital twins: agent-based adaptive modeling of cells

The main obstacle to useful cancer digital twins is state identifiability under sparse representation of a heterogeneous system. Biopsies sample a tiny part of the lesion. Radiology resolves morphology, perfusion, diffusion, but not the local distribution of proliferation, dormancy, metabolic stress, or resistant states. Any patient-specific twin that ignores heterogeneity risks fitting the visible compartment while missing the subregions most likely to drive recurrence. Within oncomorphic computing, digital twins can be constructed at the level of interacting cell agents to create calibrated predictive replicas of specific biological samples, such as patient biopsies or organoids⁸². The purpose is to provide personalized predictions of disease trajectory under various intervention scenarios⁸³. Current twin architectures often remain too coarse because they operate at lesion or organ scale and are optimized for forecasting cumulative response instead of inferring hidden local failure modes. A more useful twin could be formulated as a probabilistic multicellular model with local state variables, updated as new imaging, pathology, molecular, or ex vivo data arrive. The process begins by reconstructing a spatial scaffold of the tumor architecture from patient imaging data, which is then populated with oncomorphic agents initialized with states derived from molecular profiling, such as proliferation rates and therapy sensitivity⁸⁴. A crucial differentiator of the oncomorphic approach is the potential capability for the twins to continuously assimilate data. As new patient data becomes available, it is ingested using advanced probabilistic techniques like Bayesian calibration to update local agent parameters and interaction rules. This may allow the twin to evolve in parallel with the disease, learning its specific adaptive pathways (**Figure 4a**). At that point, the twin then could estimate which subregions are predicted to preserve viability under the next perturbation. Several recent technical developments make this approach viable^{84, 85}. Representation-learning methods for calibrating agent-based tumor models to images show that local morphology can be linked directly to executable model parameters. High-throughput cell-resolved analysis of 3D organoid dynamics now makes it possible to validate not just endpoint predictions but intermediate local trajectories, such as delayed outgrowth, protected niche survival, or morphology-dependent treatment resistant.

Closed-loop vascularized tumor-on-chip systems with integrated molecular readout provide temporally resolved measurements during intervention instead of only endpoint assays, which is exactly what sequential twin updating requires. Such systems provide the substrate to which the oncomorphic twin adds the control architecture.

4.2. Computational oncology: studying cancer adaptation dynamics *in silico*

The major issue in computational oncology is the currently weak representation of history-dependent adaptive competence at the level of single cancer cells or microregions. Computational oncology is rapidly evolving toward integrative, multiscale frameworks combining mechanistic modeling, artificial intelligence, digital twins, and *in silico* trials. Despite these advances, accurately representing history-dependent cellular decision-making under continuously changing environmental conditions remains a major challenge⁸⁶⁻⁸⁸. Cancer is intrinsically a distributed adaptive process, making it an ideal candidate in which local decision-making under stress can be formalized^{1, 17}. Oncomorphic computing addresses this limitation by representing tumors as distributed adaptive systems composed of interacting computational agents. Traditional models often simplify tumors into homogenous populations, whereas oncomorphic computing treats cancer as a complex system of individual agents (cells) whose local interactions drive tissue-scale outcomes like invasion and resistance¹⁷. In this framework, each oncomorphic agent possesses an evolving internal state that includes dynamic variables for metabolic reserves, biophysical properties like membrane potential, and a memory of recent environmental exposures. Agents can be embedded within explicit spatial environments defined by fluctuating fields of oxygen, nutrients, and drugs. In addition to biochemical cues, these agents can explicitly incorporate bioelectrical state variables. Recent work indicates that the membrane potential, intracellular calcium dynamics, ion channel activity, and localized endoplasmic reticulum (ER)-plasma membrane calcium signaling form an interconnected regulatory layer that governs proliferation and cell-cycle progression. Incorporating these variables into the internal state of each agent allows for adaptive simulations that go beyond purely genetic or metabolic descriptions, providing mechanistic representations of electrophysiological regulation^{87, 88}.

In each update step, an agent can integrate its internal state with its local context to determine its next action. Existing models can now capture growth, transport and competition, but many do not explicitly represent the transient adaptive states that allow a small subpopulation to survive therapy and enter a reversible drug-tolerant state. Ordinary differential equation (ODE) models, for example, are useful for bulk population kinetics but fail to capture the heterogeneity and subpopulation dynamics that often cause treatment failure⁸⁹. Continuum models are also adopted to describe transport phenomena, but they average out individual cellular behaviors and stochastic events⁹⁰. Standard agent-based models and evolutionary game theory consider spatial interactions and competition, but they frequently lack representation of the state-dependent plasticity and metabolism⁹¹. Oncomorphic agents could help bridge these gaps by allowing for the simulation of transient phenomena⁹². Importantly, this framework can simulate not only tumor growth trajectories, but also transient adaptive states that are difficult to access experimentally. These states include reversible drug tolerance, bioelectrical reprogramming,

localized calcium microdomain dynamics, and history-dependent shifts in proliferative competence. Transient internal states are attractive targets for mechanistic digital twins because they link molecular perturbations to tissue-level behavior and treatment response⁸⁷. Recent work on non-excitable bioelectrical interfaces further supports the view that membrane potential, coupling, and related electrical variables are active regulators of cancer behavior¹⁹. Evolutionary oncology studies show how geometry and local constraints impact progression. Image-calibrated agent-based methods demonstrate that spatial tumor simulations can be fitted to tissue images, allowing local transitions to be constrained by real architecture. Similarly, recent phenotyping techniques show that local state can be extracted in forms compatible with mechanistic modeling. An oncomorphic oncology model should therefore represent each cell or microregion by an internal state encompassing proliferative competence, metabolic reserves, accumulated sublethal damage, drug exposure, biomechanical cues, and a bioelectric state comprising membrane potential, ion-channel activity, intracellular calcium homeostasis, and local bioelectric coupling. This multidimensional state representation enables the model to capture adaptive competence as an emergent systems property resulting from the dynamic interplay between metabolism, mechanics, electrophysiology, the surrounding microenvironment, and genetic state, thereby providing a mechanistic framework for simulating cancer adaptation over multiple spatial and temporal scales.

4.3. Closed-Loop *In Vitro* Bioelectronics: dynamic automated biophysics

The main reason current *in vitro* platforms are limited for adaptive disease modelling is that they remain mostly open loop⁹³. Cancer-relevant transitions unfold over hours to days and are only indirectly reflected in measured signals. A fixed stimulation/readout protocol cannot resolve such transitions well because it does not respond to the evolving state of the culture. The technical bottleneck is therefore not better sensing but rather closed-loop state estimation under uncertainty. In standard *in vitro* assays, stimulation is controlled by the user or programmed non-adaptively. Oncomorphic closed-loop systems might enable automated event-triggered or state-triggered perturbation, and hence causal analysis of biophysical data, moving past classical action potentials recording to oncomorphic bioelectronics tailored for the bioelectric changes across diverse cell types and their spatial patterns that guide long-term morphological and phenotypic outcomes. This application area can leverage oncomorphic computing as a real-time control model for *in vitro* cell systems (**Figure 4c**). Incorporating state inference in such systems, i.e. mapping indirect electrical signals onto distributed cellular states, such as collective migration probability, coupling strength, metabolic stress levels, altered mechanics, can then trigger adaptive actuation through electrical stimulation or targeted drug delivery. Technical engineering advancements allow to realistically think about this model. Conformal 3D electrodes now allow intimate *in vitro* access to cells and tissues^{94, 95}, potentially improving long-term coupling as tissue grows or deforms. Bespoke materials interfaces enable tight coupling to record interacellular signals from flat surfaces⁹⁶. Multimodal systems combining electrophysiology, impedance, mechanical sensing, imaging, microfluidics, and molecular readout provide complementary readouts of a same phenomenon. At the same time, multimodality makes the inference problem harder. An oncomorphic controller would address this challenge by treating sensing regions as local agents with their own health and

confidence, intervening only when confidence is sufficient. Intervention may happen in the form of electrical stimulation, perfusion change, altered sampling density, or a decision to not perturb the system. In this context, oncomorphic computing represents the coordination layer linking multimodal measurement, perturbation, and sequential updating for automated biophysics.




Figure 4. Frontier applications in computational biology and digital medicine. (a) Oncomorphic digital twin of a heterogeneous tumor, represented as a population of interacting local agents whose internal states are updated as new imaging, molecular, pathology, or ex vivo data become available. The twin continuously recalibrates local predictions of proliferation, dormancy, resistance, and stress, enabling simulation of individualized disease trajectories under alternative interventions. (b) Computational oncology model based on decentralized adaptive cancer agents embedded in a spatially structured microenvironment. Each agent integrates internal history with local cues such as oxygen, nutrients, drugs, mechanics, and bioelectric state to determine transitions between proliferative, quiescent, migratory, or invasive modes. (c) Closed-loop in vitro bioelectronic platform in which multimodal sensing regions act as local agents with their own health and confidence estimates. Real-time state inference links distributed measurements to adaptive actuation, such as electrical stimulation, perfusion control, altered sampling density, or targeted drug delivery, enabling automated perturbation of evolving multicellular systems.

5. Application Frontiers in Robotic Systems

Here, the oncomorphic paradigm is applied to seemingly different contexts compared with those addressed in Section 4, namely robotic systems. We discuss how oncomorphic policies might enable continued function under local damage, uncertain sensing, or intermittent communication. An agent i may represent a space rover node, a joint, a soft actuator, a skin patch, a surgical tool, a sensing spot, or a local state estimator. Its health h_i measures the integrity of the local hardware or the validity of the local model, its confidence τ_i measures how much trust should be placed in its present estimate, its resources r_i represent compute,

memory, power, communication, actuation, or sensing budget, and its mode σ_i determines whether it operates at full capability, as a relay, in a compensatory role, or in quarantine (see Section 1.1.). Here, we discuss potential applications in deep-space exploration, humanoid robotics, soft robotics, surgical robotics, and wearable sensing systems. This focus shift compared with Section 4 is made possible by looking at all these applications as fundamentally sharing a control and state inference problem, aligning with broader bioengineering frameworks of teleonomy and goal-directed behavior in synthetic substrates⁹⁷.

5.1. Deep-Space Exploration

Deep-space autonomy has advanced by teaching spacecraft and rovers to plan, observe, and react without waiting for Earth, yet these systems still assume that the machine underneath the software remains functional^{9, 11}. The classic planning architecture flown on Deep Space 1 proved that onboard reasoning could turn mission goals into activity sequences while also performing diagnosis and recovery when expected behavior drifted from the model⁶⁴. Earth Observing-1 then allowed onboard event detection to trigger replanning, so observation was no longer a passive output of a fixed schedule but a cause of the next decision⁹⁸. On planetary surfaces, this logic was extended further by letting rovers select targets from imaging and redirect instruments toward them under compute limits^{99, 100}. Distributed mission concepts later widened the scope from a single vehicle to coordinated groups, where communication delay, and local resource asymmetry shape what autonomy must achieve¹⁰¹. These studies solved real problems, but they mostly treated degradation as an exception inside a stable architecture. From the oncomorphic perspective, control should be reorganized around local viability (**Figure 5a**): in a deep-space robot, perception, locomotion, memory, radio, would each carry a compact internal state derived from quantities such as energy, thermal margin, and link quality. A rover whose drive system becomes energetically expensive would reduce terrain ambition and preserve relay instead. In a multi-agent mission, a unit with poor communication could move from explorer to navigation reference because the system would evaluate mission goals rather than rigid local role. The result is a machine that continuously decides which computations remain worth protecting when the original architecture can no longer be fully trusted.

5.2. Humanoid Robots

Humanoid robotics is no longer limited by the ability to coordinate the body. This capacity became practical in the Atlas, where locomotion, state estimation, and whole-body control were integrated into a working system¹⁰². More recently, reinforcement learning has shown that humanoids can transfer walking policies from simulation to hardware without per-terrain retuning, including outdoors and under disturbances, while newer visuomotor and loco-manipulation systems have begun to treat knees, hands, and trunk posture as active resources^{103, 104}. What still remains underdeveloped is function preservation under partial bodily decline. The field already knows two important things. First, a robot can maintain or rebuild an internal model of its own body closely enough to detect and localize damage, then adjust behavior around the altered morphology¹⁰⁵. Second, it can recover quickly after damage if it searches the remaining behavioral space intelligently instead of relying on fixed contingency¹⁰⁶.

Oncomorphic computing turns these ideas into the normal grammar of control. In a humanoid implementation, each joint, tactile region, perception branch, would carry a local viability state computed from quantities the machine already measures: temperature rise, covariance growth, delay. Those signals would continuously reshape the controller itself. A thermally stressed knee would not simply trigger a warning, the optimizer would reduce its authority in push-off and reassign work toward safer joints. If leg confidence drops during recovery, the hands could be promoted from manipulation organs to braces or probing contacts. The humanoid remains one agent, but internally it stops assuming that all parts deserve equal trust at all times. Control migrates toward the most competent structures available in the moment, and task identity deforms with the viable body instead of with the nominal one. The result is a controller organized around graded self-preservation, where locomotion, manipulation, and balance survive by reallocating trust, energetic budget, and decision rights across the body as local competence changes.

5.3. Soft Robotics

In a soft body, mechanics, sensing, and control fail all together. A cut changes compliance, signal quality, fluid transport, and actuation at once. That is why the most important advances in the field have come from systems that treat damage as a change in the operating logic. Liquid-metal composites showed this clearly by maintaining conductive function after puncture through autonomous rerouting of electrical pathways, so the circuit did not need to remain anatomically intact to remain useful¹⁰⁷. Soft fluidic robots then pushed the same principle further by combining actuation, damage detection, self-heating, and repair in one body, allowing the machine to recognize a damaged state and shift temporarily from action to self-protection before resuming operation¹⁰⁸. At the sensory layer, self-healing skins and tactile structures have shown that fracture does not mean informational collapse, provided the material can restore electrical continuity quickly enough and the controller can interpret the repaired region as provisional¹⁰⁹. Mergeable robotic systems demonstrated that coordinated behavior can persist even when the body itself splits or replaces malfunctioning units, which matters because morphology in such machines is not a fixed container for control but part of the control problem¹¹⁰. Recent octopus-inspired systems deepen this line by embedding distributed sensing and local response into the arm itself, so grasping competence emerges partly from decentralization¹¹¹. The oncomorphic approach gives these a common architecture by letting every component maintain its own viability estimate, then assemble these values into a body-wide viability graph, which control reads continuously. If the body can merge or reconfigure, the region or module with the strongest global viability can temporarily become the coordination root^{108, 110}. This middle condition is where soft robots are most difficult to trust. They often remain active, but it becomes unclear which signals are credible and which actions are safe to scale. The value of an oncomorphic design is that it supplies an arbitration layer for this phase that keeps deciding which regions are reliable enough to govern behavior, which are only conditionally usable, and which functions deserve protection while the morphology renegotiates itself in real time.




Figure 5. Frontier applications of oncomorphic computing in robotic systems. (a) Deep-space exploration. In a distributed exploratory mission, partial probe loss and hardware degradation trigger local reorganization rather than global failure. Surviving units tighten formation, exchange state through peer-to-peer transfer, and inherit residual mission metadata from compromised neighbors to preserve navigation, sensing, and relay capability. **(b)**

Humanoid robot interaction. A humanoid platform experiencing a local mechanical failure, such as an internal joint lock, maintains task continuity by drawing on peer-to-peer mesh communication and shared environmental representations. Functions such as mapping, coordination, and task execution are redistributed across the network to sustain operational fulfillment despite partial bodily impairment. **(c)** Self-healing robotics. Damage to a mechanical actuator tendon induces a transient control reorganization phase, shown by the inset force-versus-time trajectory, in which grip behavior becomes temporarily unstable before converging to a stabilized asymmetric control mode. This illustrates bounded adaptation in response to structural injury while preserving residual manipulative function. **(d)** Adaptive surgical robotics. Failure of one robotic arm, here represented by an internal tension-cable snap, is compensated by workload recruitment from neighboring manipulators, spatial realignment, and kinematic remapping of the planned motion. The system preserves procedural continuity through adaptive trajectory completion under supervised constraints. **(e)** Smart wearable sensors. A severe physical tear in an electronic garment disrupts part of the sensing substrate, but local nodes continue operation through rerouted signal pathways and peer-to-peer metadata exchange. This allows uninterrupted physiological data logging while damaged regions are isolated or compensated by surviving sensing channels.

5.4. Surgical Robotics

In surgical robotics, during a live procedure, the machine rarely knows all parts equally well, yet most architectures still behave as if sensing confidence were globally uniform until a threshold is crossed. Recent work already points beyond that assumption. Virtual fixtures and active constraints established the idea that safety can be embedded directly into motion by shaping tool trajectories around anatomy, task geometry, and forbidden regions rather than leaving restraint entirely to the operator^{112, 113}. More autonomous systems such as STAR showed that, under structured conditions, tissue tracking, imaging, and specialized instrumentation can be assembled into closed loops that perform suturing and anastomosis with high consistency¹¹⁴. In parallel, continuum and minimally invasive platforms developed optical shape-sensing methods, especially fiber Bragg grating approaches, to recover tool curvature and configuration when line of sight is poor, and newer systems now combine geometric reconstruction with interaction-aware sensing in the same fiber-based instrument¹¹⁵. These advances suggest that safe autonomy in surgery depends mostly on managing several imperfect channels. From an oncomorphic standpoint, in a surgical robot, every sensing and inference route would maintain its own local state. Endoscopic vision would track blur, occlusion, blood masking, and registration drift. Shape-sensing fibers would track calibration quality, strain consistency, and whether the current bending regime remains observable. Force surrogates would compare expected and measured interaction to estimate model mismatch. Tissue classifiers would not merely output labels; they would expose temporal stability and local confidence. Even procedural policies would carry a viability estimate, reflecting whether the available evidence is actually sufficient for the next substep. These signals would form a live trust field across perception, state estimation, planning, and execution. A low-level safety kernel would continue to enforce hard anatomical and kinematic bounds, but the higher controller would allocate authority according to the present trust field rather than according to nominal capability alone.

5.5. Wearable Sensors

The central problem in wearable sensing is whether inference remains trustworthy when interfaces drift. Skin contact changes with motion, sweat arrives intermittently, microneedles

shift, biofouling accumulates, temperature perturbs chemistry. The field has made real progress on the hardware side. Epidermal electronics and microfluidics established soft and conformal sensors¹¹⁶⁻¹¹⁸. Integrated sweat platforms then showed how electrochemistry, temperature compensation, fluid handling, and wireless communication can be combined in a single on-body system¹¹⁹. More recent multimodal devices such as epidermal patches now couple hemodynamic and metabolic signals, e-skins track physicochemical correlates of stress, microneedle platforms support continuous access to interstitial biomarkers, and stretchable transistor architectures increasingly push drift suppression into the sensing substrate itself¹²⁰. Battery-free soft sensor arrays have achieved long-term monitoring, for instance pressure-injury prevention in wheelchair settings where local loading patterns matter clinically¹²¹. Newer sweat platforms have introduced automated in-sensor regeneration¹²². What these systems still need is a control architecture that treats reliability as distributed and unequal. The oncomorphic approach looks at giving each sensing node its own viability state, turning the wearable from a collection of smart patches into a distributed network that protects inference by negotiating trust, energy, and relevance across the network as conditions change on the body. The state would include interface integrity, impedance stability, fluid availability, microneedle quality, motion artifact. These values would feed a common trust graph, and the network would sample and fuse data according to information value under viability constraints. Meanwhile, neighboring thermal, motion, or hemodynamic channels could provide contextual bounds that stabilize state estimation while a biochemical sensor recovers, recalibrates, or regenerates^{120, 122}. At regular intervals, the wearable network solves a small allocation problem: which nodes should sample densely, which should transmit, which should spend time recalibrating, and how much confidence each stream carries, targeting continuity of trustworthy physiological information.

6. Challenges

6.1. Integration and roadmap

Monitoring local health, estimating trust, migrating tasks, and validating changed behavior consume time, power, area, and design effort^{123, 124}. This is well established in the literature on edge AI, spacecraft electronics, fault-tolerant computing, and implantable systems. A resilient architecture that spends too much energy preserving itself can easily defeat its own purpose¹²⁵. For oncomorphic computing, the question is therefore on when adaptation yields more retained utility per cost than other strategies^{17, 125}. This issue is immediate in reconfigurable hardware such as FPGAs, which already support partial reconfiguration, but the mechanisms carry overhead^{61, 63, 65}. The same caution applies to memristive and in-memory substrates, facing variability, non-linear programming, retention loss^{77, 78}. In other words, the hardware pieces that seem most attractive from an oncomorphic perspective are also the ones that demand cross-layer compensation. This is also needed in robotics and biomedicine. A wearable or implant cannot afford continuous high rate monitoring and adaptation if its battery or wireless budget is small. A digital twin updated continuously from noisy clinical streams may become useless if its latency exceeds the timescale of intervention. The engineering consequence is that oncomorphic systems will almost certainly need hierarchical adaptation: very fast local

responses for containment and fallback, slower supervisory processes for remapping and role reassignment, and even slower learning loops for policy improvement. Techniques already familiar from low-power systems are likely to be part of the solution. The deeper integration problem is that health estimation itself must be reliable. If the system misreads drift as damage, it will reorganize around the wrong information. The field will advance when it can show that health-aware adaptation improves mission-level performance under realistic resource limits. Another challenge is preventing the emergence of maladaptive behavior. If supervisory constraints are insufficient, distributed local adaptation can generate oscillatory resource allocation, instability, or pathological feedback loops. Thus, there is a need for hierarchical supervision and bounded adaptation policies instead of unconstrained self-organization.

Figure 6 shows a potential roadmap. The first implementation in neural agent populations shows the central idea of this paradigm: local agents preserve function under resource pressure. The immediate next steps are resource-aware continual learning, health-aware distributed inference, and explicit mechanisms for state migration and quarantine. These can be tested on edge clusters, robotic middleware, and distributed sensing platforms using fault injection, intermittent communication, and bounded online adaptation. The next stage links software policies to substrates that already exist. Here the relevant milestones are fault-aware FPGA or CGRA fabrics, partial reconfiguration, and hybrid digital-in-memory accelerators used under digital supervision. In parallel, the same formalism should mature in computational biology, where Bayesian multicellular digital twins and closed-loop bioelectronic control offer a tractable way to test local adaptation and intervention in systems that are distributed but observable. Upon success in the previous two, the third stage is embodied deployment. Damage-aware embodied autonomy in robotics, surgical systems, prosthetics, and wearable sensing is the long-term target. By that point, the aim is a machine that can narrow its behavior intelligently, protect its most valuable functions, and hand over control safely when necessary. Progress toward it will depend on the convergence of software and hardware.




Figure 6. Tentative roadmap for the development of oncomorphic computing. Conceptual timeline illustrating a staged progression from early software demonstrations to long-term embodied adaptive systems. Near-term milestones include oncomorphic neural agents, resource-aware continual learning, health-aware distributed inference, and state migration and quarantine. The intermediate stage links these policies to adaptive substrates through fault-aware FPGA/CGRA fabrics and hybrid digital-in-memory accelerators, while extending the formalism to biology through Bayesian multicellular digital twins and closed-loop bioelectronic control. The longer-term horizon targets damage-aware embodied autonomy and, ultimately, uncertainty-aware adaptive systems that preserve critical function across software, hardware, and living-machine interfaces under progressive damage, scarcity, and partial observability.

6.2. Ethics, safety, and validation

A system designed to preserve function can easily become difficult to control¹²⁶. If agents are allowed to change role, inherit state, down-rank other functions, or refuse low-integrity inputs, then the architecture must define in advance which of those transitions are safe, who can authorize them, and how they are recorded¹²⁷. The most realistic path is extensive training and supervised adaptation. Existing work already points in this direction. A trusted safety layer should remain external to the adaptive substrate and retain the ability to force safe mode and preserve immutable constraints on actuation, communication, and data access. This means that changes in role σ_i , resource transfer, or service migration should be interpretable¹²⁸. This challenge becomes sharper in biomedicine, where digital twins and wearable sensors may process personal data and may adapt in response to physiological drift^{129, 130}. In those cases, governance is not only about safety, but also about data provenance, user consent, and clinical accountability¹³¹. That concern already appears in the literature on adaptive medical AI, where post-deployment modification and change control are regulated. Similar tensions arise in cybersecurity: a system that autonomously reconfigures routes or trust relations may improve resilience, yet weak boundaries can make adaptive behavior difficult to distinguish from hostile behavior¹³². The ethical requirement is adaptability regulated by stop conditions and a fixed goalset.

Self-restructuring systems are difficult to verify even when their individual components are well understood, and formal verification is demanding because correctness depends on observability and environment assumptions¹²³. Oncomorphic systems broaden this already-existing problem by allowing policies to change placement, function allocation, and communication topology during operation. There is already a strong technical base on which to build in this space. Autonomous systems offer practical tools for bounding adaptive behavior, which could be translated to oncomorphic models by handling mode changes and resource transfers as supervised transitions¹²⁸. An example is showing the route by which a robot stayed upright instead of falling, over simply registering the outcome. In software, benchmarks should include node loss, delayed synchronization, adversarial participants, and contested resource pools. In hardware, they should include single event upset emulation, routing cuts, memory corruption, and latency. In robotics, they should include actuator lock, tendon damage, sensor drift, occlusion, bandwidth collapse, and controlled loss of redundancy. In biological models and digital twins, the equivalent challenges are uncertain state inference, incomplete measurements, evolving stress fields, and sequential intervention under uncertainty. The metrics proposed in Section 1.1 provide a starting point, to be completed with application-specific calibration of uncertainty and benchmarking, and the cost of unnecessary adaptation.

7. Conclusion

Oncomorphic computing formalizes a new approach to system resilience by shifting the core objective of autonomous machines from task accuracy to continuous mission viability under scarcity. It abstracts survival mechanisms from cancer biology to provide a mathematical framework that bridges software dynamics with adaptive hardware design. In safety-critical robotics and wearable diagnostics, treating structural degradation as a dynamic state variable

allows distributed networks to maintain operational continuity long after individual components suffer physical or logical failure. Conversely, in computational oncology and bioelectronics, implementing decentralized agent-based loops offers a pathway to calibrate digital twins against the non-stationary trajectories of metabolic and bioelectric dynamics. Realizing the full potential of this paradigm requires addressing substantial challenges. The computational and energetic overhead of localized state estimation, the physical non-idealities of emerging compute-in-memory substrates, and the need for non-adaptive safety boundaries to prevent resource loops remain primary technical bottlenecks. Resolving these limitations requires a structured co-design strategy. Ultimately, integrating decentralized biological survival strategies into digital and analog architectures establishes the theoretical foundations to engineer uncertainty-aware systems capable of autonomous structural and functional adaptation.

Funding Declaration

RM acknowledges support from the University of Southampton.

References

1. Moreddu, R. Bioinspired Engineering beyond Homeostasis. *Advanced Intelligent Systems* **2025**, 8 (1). DOI: 10.1002/aisy.202500435.
2. Singh, R.; Gill, S. S. Edge AI: A survey. *Internet of Things and Cyber-Physical Systems* **2023**, 3, 71-92. DOI: 10.1016/j.iotcps.2023.02.004.
3. Acharya, K.; Song, H. A Comprehensive Review of Neuro-symbolic AI for Robustness, Uncertainty Quantification, and Intervenability. *Arabian Journal for Science and Engineering* **2025**, 51 (1), 35-67. DOI: 10.1007/s13369-025-10887-3.
4. Du, Y.; Tang, G. Breaking context inertia: Adaptive context management for time series foundation models in non-stationary environments. *Neurocomputing* **2026**, 696. DOI: 10.1016/j.neucom.2026.134129.
5. van Harten, L. D.; Mousavi, M.; Jordans, R.; Pourshaghagh, H. R. Determining the necessity of fault tolerance techniques in FPGA devices for space missions. *Microprocessors and Microsystems* **2018**, 63, 1-10. DOI: 10.1016/j.micpro.2018.08.001.
6. Demosthenous, A. Advances in Microelectronics for Implantable Medical Devices. *Advances in Electronics* **2014**, 2014, 1-21. DOI: 10.1155/2014/981295.
7. Sayed, M. E.; Roberts, J. O.; Donaldson, K.; Mahon, S. T.; Iqbal, F.; Li, B.; Franco Aixela, S.; Mastorakis, G.; Jonasson, E. T.; Nemitz, M. P.; et al. Modular Robots for Enabling Operations in Unstructured Extreme Environments. *Advanced Intelligent Systems* **2022**, 4 (5). DOI: 10.1002/aisy.202000227.
8. Jiang, M.; Liu, Y.; Cui, H.; Tian, X.; Zheng, Q.; Shen, Y.; Xiong, Z.; Li, N.; Li, J.; Feng, C.; et al. Towards extreme application scenarios: perspectives on artificial intelligence-driven smart energy management systems. *Energy and AI* **2026**, 24. DOI: 10.1016/j.egyai.2026.100725.
9. Carbone, M. A.; Loparo, K. A. Fault Detection and Diagnosis in Spacecraft Electrical Power Systems. *Journal of Aerospace Information Systems* **2023**, 20 (6), 308-318. DOI: 10.2514/1.1011136.
10. Kwarteng, E.; Cebe, M. A survey on security issues in modern Implantable Devices: Solutions and future issues. *Smart Health* **2022**, 25. DOI: 10.1016/j.smhl.2022.100295.

11. Rech, P. Artificial Neural Networks for Space and Safety-Critical Applications: Reliability Issues and Potential Solutions. *IEEE Transactions on Nuclear Science* **2024**, *71* (4), 377-404. DOI: 10.1109/tns.2024.3349956.
12. Alatawi, M. N. Edge computing and federated learning for privacy-preserving IoT analytics. *Journal on Wireless Communications and Networking* **2025**, *2026* (1). DOI: 10.1186/s13638-025-02545-x.
13. Hernández, H. A.; Mondragón, I. F.; González, S. R.; Pedraza, L. F. Reconfigurable agricultural robotics: Control strategies, communication, and applications. *Computers and Electronics in Agriculture* **2025**, *234*. DOI: 10.1016/j.compag.2025.110161.
14. O'Keeffe, J. Detecting and diagnosing faults in autonomous robot swarms with an artificial antibody population model. *R Soc Open Sci* **2025**, *12* (10), 251252. DOI: 10.1098/rsos.251252.
15. Ahmed, T. Bio-inspired artificial synapses: Neuromorphic computing chip engineering with soft biomaterials. *Memories - Materials, Devices, Circuits and Systems* **2023**, *6*. DOI: 10.1016/j.memori.2023.100088.
16. Doursat, R.; Sayama, H.; Michel, O. A review of morphogenetic engineering. *Natural Computing* **2013**, *12* (4), 517-535. DOI: 10.1007/s11047-013-9398-1.
17. Greulich, P.; Levin, M.; Moreddu, R. Oncomorphic neural agent populations for resource-limited sequential learning. *arXiv e-prints* **2025**, arXiv: 2503.12743.
18. Meacham, C. E.; Morrison, S. J. Tumour heterogeneity and cancer cell plasticity. *Nature* **2013**, *501* (7467), 328-337. DOI: 10.1038/nature12624.
19. Cadinu, P.; Burgess, M. K.; Franco Jones, C.; Iarossi, M.; Schröter, M.; Nakatsuka, N.; Djamgoz, M. B. A.; Gonçalves, G.; Abayzeed, S.; Sanjuán-Alberte, P.; et al. Bioelectrical Interfaces Beyond Excitable Cells: Cancer, Aging, and Gene Expression Modulation. *Advanced Materials Interfaces* **2026**, *13* (10). DOI: 10.1002/admi.202500999.
20. Franco Jones, C.; Dias, D.; Moreira, A. C.; Goncalves, G.; Cinti, S.; Djamgoz, M. B. A.; Castelo Ferreira, F.; Sanjuan-Alberte, P.; Moreddu, R. Multilevel classification framework for breast cancer cell selection and its integration with advanced disease models. *iScience* **2025**, *28* (10), 113579. DOI: 10.1016/j.isci.2025.113579.
21. Levin, M. Bioelectrical approaches to cancer as a problem of the scaling of the cellular self. *Prog Biophys Mol Biol* **2021**, *165*, 102-113. DOI: 10.1016/j.pbiomolbio.2021.04.007.
22. Melikov, R.; De Angelis, F.; Moreddu, R. High-frequency extracellular spiking in electrically-active cancer cells. *bioRxiv* **2024**, 2024.2003. 2016.585162.
23. Abayzeed, S.; Galvis, D.; Medel, K. R.; Jones, C. F.; Gonzalez, O. B.; Setchfield, K.; Moreddu, R.; Somekh, M. G.; Wedgwood, K. C. A.; Smith, P. Plasmonic imaging of living pancreatic beta-cell networks. *Sci Rep* **2026**, *16* (1), 3993. DOI: 10.1038/s41598-025-34094-0.
24. Franca, G. S.; Baron, M.; King, B. R.; Bossowski, J. P.; Bjornberg, A.; Pour, M.; Rao, A.; Patel, A. S.; Misirlioglu, S.; Barkley, D.; et al. Cellular adaptation to cancer therapy along a resistance continuum. *Nature* **2024**, *631* (8022), 876-883. DOI: 10.1038/s41586-024-07690-9.
25. Agrawal, A.; Javanmardi, Y.; Watson, S. A.; Serwinski, B.; Djordjevic, B.; Li, W.; Aref, A. R.; Jenkins, R. W.; Moeendarbary, E. Mechanical signatures in cancer metastasis. *NPJ Biol Phys Mech* **2025**, *2* (1), 3. DOI: 10.1038/s44341-024-00007-x.
26. Olfati-Saber, R.; Fax, J. A.; Murray, R. M. Consensus and Cooperation in Networked Multi-Agent Systems. *Proceedings of the IEEE* **2007**, *95* (1), 215-233. DOI: 10.1109/jproc.2006.887293.

27. Sebastián, E.; Duong, T.; Atanasov, N.; Montijano, E.; Sagüés, C. Physics-Informed Multiagent Reinforcement Learning for Distributed Multirobot Problems. *IEEE Transactions on Robotics* **2025**, *41*, 4499-4517. DOI: 10.1109/tro.2025.3582836.
28. Soro, I. W.; Noureldath, M.; Aït-Kadi, D. Performance evaluation of multi-state degraded systems with minimal repairs and imperfect preventive maintenance. *Reliability Engineering & System Safety* **2010**, *95* (2), 65-69. DOI: 10.1016/j.res.2009.08.004.
29. Li, Y.; Chen, Y.; Zhang, Q.; Kang, R. Belief reliability analysis of multi-state deteriorating systems under epistemic uncertainty. *Information Sciences* **2022**, *604*, 249-266. DOI: 10.1016/j.ins.2022.05.022.
30. Lee, H.-R.; Lee, T. Multi-agent reinforcement learning algorithm to solve a partially-observable multi-agent problem in disaster response. *European Journal of Operational Research* **2021**, *291* (1), 296-308. DOI: 10.1016/j.ejor.2020.09.018.
31. Wang, Y.; Xu, X. Adaptive safety-critical control for a class of nonlinear systems with parametric uncertainties: A control barrier function approach. *Systems & Control Letters* **2024**, *188*. DOI: 10.1016/j.sysconle.2024.105798.
32. Liu, S.; Wang, S.; Dai, M.; Lian, G.; Xu, J.; Kong, X. Safe multi-agent reinforcement learning based on adversarial strategy and control barrier function feasible set. *Neurocomputing* **2026**, *683*. DOI: 10.1016/j.neucom.2026.133488.
33. Hu, W.; Nishimura, Y.; Zhu, Q. Control barrier functions and safety-critical control with input constraints for stochastic systems. *Systems & Control Letters* **2026**, *216*. DOI: 10.1016/j.sysconle.2026.106530.
34. Madeira, H.; Quadros, G.; Silva, J. G. Experimental evaluation of a set of simple error detection mechanisms. *Microprocessing and Microprogramming* **1990**, *30* (1-5), 513-520. DOI: 10.1016/0165-6074(90)90292-h.
35. Kephart, J. O.; Chess, D. M. The vision of autonomic computing. *Computer* **2003**, *36* (1), 41-50. DOI: 10.1109/mc.2003.1160055.
36. Weyns, D.; Gerostathopoulos, I.; Abbas, N.; Andersson, J.; Biffi, S.; Brada, P.; Bures, T.; Di Salle, A.; Galster, M.; Lago, P.; et al. Self-Adaptation in Industry: A Survey. *ACM Transactions on Autonomous and Adaptive Systems* **2023**, *18* (2), 1-44. DOI: 10.1145/3589227.
37. Bartolozzi, C.; Indiveri, G.; Donati, E. Embodied neuromorphic intelligence. *Nat Commun* **2022**, *13* (1), 1024. DOI: 10.1038/s41467-022-28487-2.
38. Zhang, Y.; Zeng, D.; Luo, J.; Fu, X.; Chen, G.; Xu, Z.; King, I. A Survey of Trustworthy Federated Learning: Issues, Solutions, and Challenges. *ACM Transactions on Intelligent Systems and Technology* **2024**, *15* (6), 1-47. DOI: 10.1145/3678181.
39. Fiolhais, L.; Sousa, L. Transient-Execution Attacks: A Computer Architect Perspective. *ACM Computing Surveys* **2023**, *56* (3), 1-38. DOI: 10.1145/3603619.
40. Merenda, M.; Porcaro, C.; Iero, D. Edge Machine Learning for AI-Enabled IoT Devices: A Review. *Sensors (Basel)* **2020**, *20* (9). DOI: 10.3390/s20092533.
41. Mead, C. How we created neuromorphic engineering. *Nature Electronics* **2020**, *3* (7), 434-435. DOI: 10.1038/s41928-020-0448-2.
42. Ielmini, D.; Wong, H. S. P. In-memory computing with resistive switching devices. *Nature Electronics* **2018**, *1* (6), 333-343. DOI: 10.1038/s41928-018-0092-2.
43. van de Burgt, Y.; Lubberman, E.; Fuller, E. J.; Keene, S. T.; Faria, G. C.; Agarwal, S.; Marinella, M. J.; Alec Talin, A.; Salleo, A. A non-volatile organic electrochemical device as a low-voltage artificial synapse for neuromorphic computing. *Nat Mater* **2017**, *16* (4), 414-418. DOI: 10.1038/nmat4856.

44. Vahl, A.; Milano, G.; Kuncic, Z.; Brown, S. A.; Milani, P. Brain-inspired computing with self-assembled networks of nano-objects. *Journal of Physics D: Applied Physics* **2024**, *57* (50). DOI: 10.1088/1361-6463/ad7a82.
45. Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. Overcoming catastrophic forgetting in neural networks. *Proc Natl Acad Sci U S A* **2017**, *114* (13), 3521-3526. DOI: 10.1073/pnas.1611835114.
46. Sajko, G.; Babič, J. Evidence accumulation with adaptive weighting of social and personal information for collective perception. *Swarm Intelligence* **2025**, *19* (4), 295-316. DOI: 10.1007/s11721-025-00253-2.
47. Wu, J.; Dong, F.; Leung, H.; Zhu, Z.; Zhou, J.; Drew, S. Topology-aware Federated Learning in Edge Computing: A Comprehensive Survey. *ACM Computing Surveys* **2024**, *56* (10), 1-41. DOI: 10.1145/3659205.
48. Abreha, H. G.; Hayajneh, M.; Serhani, M. A. Federated Learning in Edge Computing: A Systematic Survey. *Sensors (Basel)* **2022**, *22* (2). DOI: 10.3390/s22020450.
49. Gurmu, M. D.; Hu, W. Fully distributed event-triggered consensus of linear multiagent systems with multiplicative faults. *Transactions of the Institute of Measurement and Control* **2025**, *48* (10), 2090-2102. DOI: 10.1177/01423312251334190.
50. Luo, Y.; Yang, Z.; Meng, F.; Li, Y.; Zhou, J.; Zhang, Y. An Empirical Study of Catastrophic Forgetting in Large Language Models During Continual Fine-Tuning. *IEEE Transactions on Audio, Speech and Language Processing* **2025**, *33*, 3776-3786. DOI: 10.1109/taspro.2025.3606231.
51. van de Ven, G. M.; Siegelmann, H. T.; Tolias, A. S. Brain-inspired replay for continual learning with artificial neural networks. *Nat Commun* **2020**, *11* (1), 4069. DOI: 10.1038/s41467-020-17866-2.
52. Perkonigg, M.; Hofmanninger, J.; Herold, C. J.; Brink, J. A.; Panykh, O.; Prosch, H.; Langs, G. Dynamic memory to alleviate catastrophic forgetting in continual learning with medical imaging. *Nat Commun* **2021**, *12* (1), 5678. DOI: 10.1038/s41467-021-25858-z.
53. Paroli, B.; Borghi, F.; Potenza, M.; Milani, P. The receptron is a nonlinear threshold logic gate with intrinsic multi-dimensional selective capabilities for analog inputs. *arXiv preprint arXiv:2506.19642* **2025**.
54. Ma, Q.; Xu, Y.; Xu, H.; Jiang, Z.; Huang, L.; Huang, H. FedSA: A Semi-Asynchronous Federated Learning Mechanism in Heterogeneous Edge Computing. *IEEE Journal on Selected Areas in Communications* **2021**, *39* (12), 3654-3672. DOI: 10.1109/jsac.2021.3118435.
55. Sun, Y. C.; Yang, G. H. Event-Triggered Distributed State Estimation for Multiagent Systems Under DoS Attacks. *IEEE Trans Cybern* **2022**, *52* (7), 6901-6910. DOI: 10.1109/TCYB.2020.3034456.
56. Xiao, F.; Wang, X.; Wei, B.; Feng, Q. Synchronization of asynchronous networked systems via hybrid edge-event-triggered control with model uncertainties. *Systems & Control Letters* **2026**, *209*. DOI: 10.1016/j.sysconle.2025.106329.
57. Matsume, H.; Wang, Y.; Ishii, H.; Défago, X. Mixed self/event-triggered ternary control for resilient consensus against mobile adversarial agents. *Nonlinear Analysis: Hybrid Systems* **2024**, *52*. DOI: 10.1016/j.nahs.2024.101473.
58. Cao, H.; Hong, M.; Wang, Z.; Feng, Y.; Xiao, M. Event-triggered bipartite output control of heterogeneous multi-agent systems with input saturation and DoS attacks. *International Journal of Systems Science* **2026**, 1-24. DOI: 10.1080/00207721.2026.2672079.

59. Decastri, D.; Nieuw, T.; Zuccali, C.; Giacomozzi, F.; Lorenzelli, L.; Borghi, F.; Milani, P. Programmable Functional Connectivity and Synchronous Activity in Resistive Switching Self-Assembled Nanostructured Networks. *Small Structures* **2025**, *6* (11). DOI: 10.1002/ssstr.202500330.
60. Austin, H.; Major, C.; Barney, C.; Williams, J.; Becker, Z.; Smith, M.; LaMeres, B. Single-event upset simulation and detection in configuration memory. *Frontiers in Space Technologies* **2025**, *6*. DOI: 10.3389/frspt.2025.1610424.
61. Siecha, R. T.; Alemu, G.; Prinzie, J.; Leroux, P. Analysis of the SEU Tolerance of an FPGA-Based Time-to-Digital Converter Using Emulation-Based Fault Injection. *Electronics* **2025**, *14* (11). DOI: 10.3390/electronics14112176.
62. Aranda, L. A.; Wessman, N.-J.; Santos, L.; Sánchez-Macián, A.; Andersson, J.; Weigand, R.; Maestro, J. A. Analysis of the Critical Bits of a RISC-V Processor Implemented in an SRAM-Based FPGA for Space Applications. *Electronics* **2020**, *9* (1). DOI: 10.3390/electronics9010175.
63. Cano-Páez, J.; Entrena, L.; García-Valderas, M.; Lindoso, A. Self-healing architecture for RISC-V soft core processor in SRAM-based FPGAs using Dynamic Partial Reconfiguration. *IEEE Transactions on Nuclear Science* **2026**, 1-1. DOI: 10.1109/tns.2026.3678753.
64. Petrick, D.; Geist, A.; Albaijes, D.; Davis, M.; Sparacino, P.; Crum, G.; Ripley, R.; Boblitt, J.; Flatley, T. SpaceCube v2.0 space flight hybrid reconfigurable data processing system. In 2014 IEEE Aerospace Conference, 2014.
65. Julien, C. R.; LaMeres, B. J.; Weber, R. J. An FPGA-based radiation tolerant SmallSat Computer System. In 2017 IEEE Aerospace Conference, 2017.
66. Tonfat, J.; Lima Kastensmidt, F.; Artola, L.; Hubert, G.; Medina, N. H.; Added, N.; Aguiar, V. A. P.; Aguirre, F.; Macchione, E. L. A.; Silveira, M. A. G. Analyzing the Influence of the Angles of Incidence and Rotation on MBU Events Induced by Low LET Heavy Ions in a 28-nm SRAM-Based FPGA. *IEEE Transactions on Nuclear Science* **2017**, *64* (8), 2161-2168. DOI: 10.1109/tns.2017.2727479.
67. Li, C.; Hu, M.; Li, Y.; Jiang, H.; Ge, N.; Montgomery, E.; Zhang, J.; Song, W.; Dávila, N.; Graves, C. E.; et al. Analogue signal and image processing with large memristor crossbars. *Nature Electronics* **2017**, *1* (1), 52-59. DOI: 10.1038/s41928-017-0002-z.
68. Martini, G.; Paroli, B.; Milani, P. Non-linear reconfigurable threshold logic gates based on nanostructured metallic films. *npj Unconventional Computing* **2025**, *2* (1). DOI: 10.1038/s44335-025-00037-1.
69. Herrera-Alzu, I.; Lopez-Vallejo, M. Design Techniques for Xilinx Virtex FPGA Configuration Memory Scrubbers. *IEEE Transactions on Nuclear Science* **2013**, *60* (1), 376-385. DOI: 10.1109/tns.2012.2231881.
70. Beckhoff, C.; Koch, D.; Torresen, J. Design Tools for Implementing Self-Aware and Fault-Tolerant Systems on FPGAs. *ACM Transactions on Reconfigurable Technology and Systems* **2014**, *7* (2), 1-23. DOI: 10.1145/2617597.
71. Sahoo, S. S.; Ranjbar, B.; Kumar, A. Reliability-Aware Resource Management in Multi-/Many-Core Systems: A Perspective Paper. *Journal of Low Power Electronics and Applications* **2021**, *11* (1). DOI: 10.3390/jlpea11010007.
72. Zhang, B.; Li, Y.; Zhang, S.; Zhang, Y.; Zhu, B.; Wang, P. An Adaptive Task Migration Scheduling Approach for Edge-Cloud Collaborative Inference. *Wireless Communications and Mobile Computing* **2022**, *2022* (1). DOI: 10.1155/2022/8804530.
73. Cai, F.; Correll, J. M.; Lee, S. H.; Lim, Y.; Bothra, V.; Zhang, Z.; Flynn, M. P.; Lu, W. D. A fully integrated reprogrammable memristor-CMOS system for efficient multiply-accumulate operations. *Nature Electronics* **2019**, *2* (7), 290-299. DOI: 10.1038/s41928-019-0270-x.

74. Chen, W.-H.; Dou, C.; Li, K.-X.; Lin, W.-Y.; Li, P.-Y.; Huang, J.-H.; Wang, J.-H.; Wei, W.-C.; Xue, C.-X.; Chiu, Y.-C.; et al. CMOS-integrated memristive non-volatile computing-in-memory for AI edge processors. *Nature Electronics* **2019**, *2* (9), 420-428. DOI: 10.1038/s41928-019-0288-0.
75. Wan, W.; Kubendran, R.; Schaefer, C.; Eryilmaz, S. B.; Zhang, W.; Wu, D.; Deiss, S.; Raina, P.; Qian, H.; Gao, B.; et al. A compute-in-memory chip based on resistive random-access memory. *Nature* **2022**, *608* (7923), 504-512. DOI: 10.1038/s41586-022-04992-8.
76. Yang, J. J.; Strukov, D. B.; Stewart, D. R. Memristive devices for computing. *Nat Nanotechnol* **2013**, *8* (1), 13-24. DOI: 10.1038/nnano.2012.240.
77. Yao, P.; Wu, H.; Gao, B.; Tang, J.; Zhang, Q.; Zhang, W.; Yang, J. J.; Qian, H. Fully hardware-implemented memristor convolutional neural network. *Nature* **2020**, *577* (7792), 641-646. DOI: 10.1038/s41586-020-1942-4.
78. Wang, L.; Li, W.; Zhou, Z.; An, J.; Ye, W.; Li, Z.; Gao, H.; Hu, H.; Liu, J.; Chen, X.; et al. A near-threshold memristive computing-in-memory engine for edge intelligence. *Nat Commun* **2025**, *16* (1), 5897. DOI: 10.1038/s41467-025-61025-4.
79. Li, Y.; Wang, S.; Zhao, Y.; Wang, S.; Wang, B.; Zhang, W.; He, Y.; Lin, N.; Cui, B.; Chen, X.; et al. Pruning random resistive memory for optimizing analog AI. *Nat Commun* **2026**, *17* (1), 1190. DOI: 10.1038/s41467-025-67960-6.
80. Bressan, S.; Camillini, L.; Borghi, F.; Galafassi, G.; Milani, P. A Polymorphic Reconfigurable Multi-Electrode Device Based on Electrically Bistable Nanostructured Metallic Films. *Advanced Electronic Materials* **2026**, *12* (8). DOI: 10.1002/aelm.202500636.
81. Nadalini, G.; Dallinger, A.; Sottocorno, D.; Greco, F.; Borghi, F.; Milani, P. Soft Electronic Switches and Adaptive Logic Gates Based on Nanostructured Gold Networks. *Advanced Electronic Materials* **2024**, *11* (7). DOI: 10.1002/aelm.202400717.
82. Laubenbacher, R.; Adler, F.; An, G.; Castiglione, F.; Eubank, S.; Fonseca, L. L.; Glazier, J.; Helikar, T.; Jett-Tilton, M.; Kirschner, D.; et al. Forum on immune digital twins: a meeting report. *NPJ Syst Biol Appl* **2024**, *10* (1), 19. DOI: 10.1038/s41540-024-00345-5.
83. Cess, C. G.; Finley, S. D. Calibrating agent-based models to tumor images using representation learning. *PLoS Comput Biol* **2023**, *19* (4), e1011070. DOI: 10.1371/journal.pcbi.1011070.
84. Bera, K.; Braman, N.; Gupta, A.; Velcheti, V.; Madabhushi, A. Predicting cancer outcomes with radiomics and artificial intelligence in radiology. *Nat Rev Clin Oncol* **2022**, *19* (2), 132-146. DOI: 10.1038/s41571-021-00560-7.
85. Mukashyaka, P.; Kumar, P.; Mellert, D. J.; Nicholas, S.; Noorbakhsh, J.; Brugiolo, M.; Courtois, E. T.; Anczukow, O.; Liu, E. T.; Chuang, J. H. High-throughput deconvolution of 3D organoid dynamics at cellular resolution for cancer pharmacology with Cellos. *Nat Commun* **2023**, *14* (1), 8406. DOI: 10.1038/s41467-023-44162-6.
86. Baumgartner, C. Computational modeling and simulation in oncology. *Clin Transl Med* **2025**, *15* (9), e70456. DOI: 10.1002/ctm2.70456.
87. Desoyer, C.; Supe, R.; Ghorbanpour, S.; Goswami, D.; Khalaj, M.; Karami, N.; Brislinger, D.; Hauser, C. A. E.; Baumgartner, C. Cancer: A bioelectric disease? *Clin Transl Med* **2026**, *16* (5), e70701. DOI: 10.1002/ctm2.70701.
88. Langthaler, S.; Zumpf, C.; Rienmuller, T.; Shrestha, N.; Fuchs, J.; Zhou, R.; Pelzmann, B.; Zorn-Pauly, K.; Frohlich, E.; Weinberg, S. H.; et al. The bioelectric mechanisms of local calcium dynamics in cancer cell proliferation: an extension of the A549 in silico cell model. *Front Mol Biosci* **2024**, *11*, 1394398. DOI: 10.3389/fmolb.2024.1394398.

89. Dagogo-Jack, I.; Shaw, A. T. Tumour heterogeneity and resistance to cancer therapies. *Nat Rev Clin Oncol* **2018**, *15* (2), 81-94. DOI: 10.1038/nrclinonc.2017.166.
90. Braun, R. D.; Beatty, A. L. Modeling of oxygen transport across tumor multicellular layers. *Microvasc Res* **2007**, *73* (2), 113-123. DOI: 10.1016/j.mvr.2006.11.001.
91. Dingli, D.; Chalub, F. A.; Santos, F. C.; Van Segbroeck, S.; Pacheco, J. M. Cancer phenotype as the outcome of an evolutionary game between normal and malignant cells. *Br J Cancer* **2009**, *101* (7), 1130-1136. DOI: 10.1038/sj.bjc.6605288.
92. Parker, J. J.; Canoll, P.; Niswander, L.; Kleinschmidt-DeMasters, B. K.; Foshay, K.; Waziri, A. Intratumoral heterogeneity of endogenous tumor cell invasive behavior in human glioblastoma. *Sci Rep* **2018**, *8* (1), 18002. DOI: 10.1038/s41598-018-36280-9.
93. Asefifeyzabadi, N.; Nguyen, T.; Li, H.; Zhu, K.; Yang, H. Y.; Baniya, P.; Medina Lopez, A.; Gallegos, A.; Hsieh, H. C.; Dechiraju, H.; et al. A pro-reparative bioelectronic device for controlled delivery of ions and biomolecules. *Wound Repair Regen* **2024**, *32* (5), 709-719. DOI: 10.1111/wrr.13191.
94. Moreddu, R.; Boschi, A.; d'Amora, M.; Hubarevich, A.; Dipalo, M.; De Angelis, F. Passive Recording of Bioelectrical Signals from Non-Excitable Cells by Fluorescent Mirroring. *Nano Lett* **2023**, *23* (8), 3217-3223. DOI: 10.1021/acs.nanolett.2c05053.
95. Boschi, A.; Iachetta, G.; Buonocore, S.; Hubarevich, A.; Hurtaud, J.; Moreddu, R.; Marta, d. A.; Formoso, M. B.; Tantussi, F.; Dipalo, M.; et al. Interferometric Biosensor for High Sensitive Label-Free Recording of HiPS Cardiomyocytes Contraction in Vitro. *Nano Lett* **2024**, *24* (22), 6451-6458. DOI: 10.1021/acs.nanolett.3c04291.
96. Kyndiah, A.; Zemignani, G. Z.; Ronchi, C.; Tullii, G.; Khudiakov, A.; Iachetta, G.; Chiodini, S.; Moreddu, R.; Viola, F. A.; Schwartz, P. J.; et al. Non-invasive action potential recordings using printed electrolyte-gated polymer field-effect transistors. *Nat Commun* **2025**, *16* (1), 8143. DOI: 10.1038/s41467-025-63484-1.
97. Clawson, W. P.; Levin, M. Endless forms most beautiful 2.0: teleonomy and the bioengineering of chimaeric and synthetic organisms. *Biological Journal of the Linnean Society* **2023**, *139* (4), 457-486. DOI: 10.1093/biolinnean/blac073.
98. Middleton, E. M.; Ungar, S. G.; Mandl, D. J.; Ong, L.; Frye, S. W.; Campbell, P. E.; Landis, D. R.; Young, J. P.; Pollack, N. H. The Earth Observing One (EO-1) Satellite Mission: Over a Decade in Space. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **2013**, *6* (2), 243-256. DOI: 10.1109/jstars.2013.2249496.
99. Estlin, T. A.; Bornstein, B. J.; Gaines, D. M.; Anderson, R. C.; Thompson, D. R.; Burl, M.; Castaño, R.; Judd, M. AEGIS Automated Science Targeting for the MER Opportunity Rover. *ACM Transactions on Intelligent Systems and Technology* **2012**, *3* (3), 1-19. DOI: 10.1145/2168752.2168764.
100. Francis, R.; Estlin, T.; Doran, G.; Johnstone, S.; Gaines, D.; Verma, V.; Burl, M.; Frydenvang, J.; Montano, S.; Wiens, R. C.; et al. AEGIS autonomous targeting for ChemCam on Mars Science Laboratory: Deployment and results of initial science team use. *Sci Robot* **2017**, *2* (7). DOI: 10.1126/scirobotics.aan4582.
101. Selva, D.; Golkar, A.; Korobova, O.; Cruz, I. L. i.; Collopy, P.; de Weck, O. L. Distributed Earth Satellite Systems: What Is Needed to Move Forward? *Journal of Aerospace Information Systems* **2017**, *14* (8), 412-438. DOI: 10.2514/1.1010497.
102. Kuindersma, S.; Deits, R.; Fallon, M.; Valenzuela, A.; Dai, H.; Permenter, F.; Koolen, T.; Marion, P.; Tedrake, R. Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot. *Autonomous Robots* **2015**, *40* (3), 429-455. DOI: 10.1007/s10514-015-9479-3.

103. Radosavovic, I.; Xiao, T.; Zhang, B.; Darrell, T.; Malik, J.; Sreenath, K. Real-world humanoid locomotion with reinforcement learning. *Sci Robot* **2024**, *9* (89), eadi9579. DOI: 10.1126/scirobotics.adi9579.
104. Ji, S.; Chen, Z.; Zhang, C.; Gao, J.; Qiao, H. A Unified Reinforcement Learning Framework for Payload-Aware Humanoid Control With Integrated Inertial and State Estimation. *IEEE/ASME Transactions on Mechatronics* **2026**, *31* (3), 3654-3664. DOI: 10.1109/tmech.2025.3648874.
105. Bongard, J.; Zykov, V.; Lipson, H. Resilient machines through continuous self-modeling. *Science* **2006**, *314* (5802), 1118-1121. DOI: 10.1126/science.1133687.
106. Cully, A.; Clune, J.; Tarapore, D.; Mouret, J. B. Robots that can adapt like animals. *Nature* **2015**, *521* (7553), 503-507. DOI: 10.1038/nature14422.
107. Markvicka, E. J.; Bartlett, M. D.; Huang, X.; Majidi, C. An autonomously electrically self-healing liquid metal-elastomer composite for robust soft-matter robotics and electronics. *Nat Mater* **2018**, *17* (7), 618-624. DOI: 10.1038/s41563-018-0084-7.
108. Pozarycki, T. A.; Hwang, D.; Barron, E. J., 3rd; Wilcox, B. T.; Tutika, R.; Bartlett, M. D. Tough Bonding of Liquid Metal-Elastomer Composites for Multifunctional Adhesives. *Small* **2022**, *18* (41), e2203700. DOI: 10.1002/sml.202203700.
109. Root, S. E.; Wu, C.; Choi, H.; Sun, E.; Ngaruka, G.; Park, H.; Ramos Figueroa, A. L.; Berman, A.; Patino, D. U.; Shi, Y.; et al. A damage-perceptive, self-healing electronic skin with millimeter resolution. *Device* **2025**, *3* (8). DOI: 10.1016/j.device.2025.100802.
110. Mathews, N.; Christensen, A. L.; O'Grady, R.; Mondada, F.; Dorigo, M. Mergeable nervous systems for robots. *Nat Commun* **2017**, *8* (1), 439. DOI: 10.1038/s41467-017-00109-2.
111. Del Dottore, E.; Adhami, R.; Shahabi, E.; Solfiti, E.; Martini, M.; Mariani, S.; Parmiggiani, A.; Mondini, A.; Sinibaldi, E.; Mazzolai, B. Peripheral control enabled by distributed sensing in an octopus-inspired soft robotic arm for autonomous underwater grasping. *Nature Machine Intelligence* **2026**, *8* (5), 708-721. DOI: 10.1038/s42256-026-01230-y.
112. Rosenberg, L. B. Virtual fixtures: Perceptual tools for telerobotic manipulation. In Proceedings of IEEE Virtual Reality Annual International Symposium, 1993.
113. Lopez, E.; Kwok, K. W.; Payne, C. J.; Giataganas, P.; Yang, G. Z. Implicit Active Constraints for Robot-Assisted Arthroscopy. *IEEE Int Conf Robot Autom* **2013**, *2013*, 5390-5395. DOI: 10.1109/ICRA.2013.6631350.
114. Shademan, A.; Decker, R. S.; Opfermann, J. D.; Leonard, S.; Krieger, A.; Kim, P. C. Supervised autonomous robotic soft tissue surgery. *Sci Transl Med* **2016**, *8* (337), 337ra364. DOI: 10.1126/scitranslmed.aad9398.
115. Sefati, S.; Gao, C.; Iordachita, I.; Taylor, R. H.; Armand, M. Data-Driven Shape Sensing of a Surgical Continuum Manipulator Using an Uncalibrated Fiber Bragg Grating Sensor. *IEEE Sens J* **2021**, *21* (3), 3066-3076. DOI: 10.1109/jsen.2020.3028208.
116. Kim, D. H.; Lu, N.; Ma, R.; Kim, Y. S.; Kim, R. H.; Wang, S.; Wu, J.; Won, S. M.; Tao, H.; Islam, A.; et al. Epidermal electronics. *Science* **2011**, *333* (6044), 838-843. DOI: 10.1126/science.1206157.
117. Moreddu, R.; Mahmoodi, N.; Kassanos, P.; Vigolo, D.; Mendes, P. M.; Yetisen, A. K. Stretchable Nanostructures as Optomechanical Strain Sensors for Ophthalmic Applications. *ACS Applied Polymer Materials* **2021**, *3* (11), 5416-5424. DOI: 10.1021/acsapm.1c00703.

118. Moreddu, R.; Nasrollahi, V.; Kassanos, P.; Dimov, S.; Vigolo, D.; Yetisen, A. K. Lab-on-a-Contact Lens Platforms Fabricated by Multi-Axis Femtosecond Laser Ablation. *Small* **2021**, *17* (38), e2102008. DOI: 10.1002/sml.202102008.
119. Gao, W.; Emaminejad, S.; Nyein, H. Y. Y.; Challa, S.; Chen, K.; Peck, A.; Fahad, H. M.; Ota, H.; Shiraki, H.; Kiriya, D.; et al. Fully integrated wearable sensor arrays for multiplexed in situ perspiration analysis. *Nature* **2016**, *529* (7587), 509-514. DOI: 10.1038/nature16521.
120. Nyein, H. Y. Y.; Bariya, M.; Kivimaki, L.; Uusitalo, S.; Liaw, T. S.; Jansson, E.; Ahn, C. H.; Hangasky, J. A.; Zhao, J.; Lin, Y.; et al. Regional and correlative sweat analysis using high-throughput microfluidic sensing patches toward decoding sweat. *Sci Adv* **2019**, *5* (8), eaaw9906. DOI: 10.1126/sciadv.aaw9906.
121. Xu, C.; Song, Y.; Sempionatto, J. R.; Solomon, S. A.; Yu, Y.; Nyein, H. Y. Y.; Tay, R. Y.; Li, J.; Heng, W.; Min, J.; et al. A physicochemical-sensing electronic skin for stress response monitoring. *Nat Electron* **2024**, *7* (2), 168-179. DOI: 10.1038/s41928-023-01116-6.
122. Tehrani, F.; Teymourian, H.; Wuerstle, B.; Kavner, J.; Patel, R.; Furnidge, A.; Aghavali, R.; Hosseini-Toudeshki, H.; Brown, C.; Zhang, F.; et al. An integrated wearable microneedle array for the continuous monitoring of multiple biomarkers in interstitial fluid. *Nat Biomed Eng* **2022**, *6* (11), 1214-1224. DOI: 10.1038/s41551-022-00887-1.
123. Emberland, K. E.; Rortveit, G. [Norwegian health data - an inaccessible treasure]. *Tidsskr Nor Laegeforen* **2016**, *136* (18), 1506. DOI: 10.4045/tidsskr.16.0613.
124. Wongpiromsarn, T.; Ghasemi, M.; Cubuktepe, M.; Bakirtzis, G.; Carr, S.; Karabag, M. O.; Neary, C.; Gohari, P.; Topcu, U. Formal Methods for Autonomous Systems. *Foundations and Trends® in Systems and Control* **2023**, *10* (3-4), 180-407. DOI: 10.1561/26000000029.
125. Sifakis, J.; Harel, D. Trustworthy Autonomous System Development. *ACM Transactions on Embedded Computing Systems* **2023**, *22* (3), 1-24. DOI: 10.1145/3545178.
126. Luckcuck, M.; Farrell, M.; Dennis, L. A.; Dixon, C.; Fisher, M. Formal Specification and Verification of Autonomous Robotic Systems. *ACM Computing Surveys* **2019**, *52* (5), 1-41. DOI: 10.1145/3342355.
127. Harder, A.; Kulkarni, A.; Behl, M. HALO: Fault-Tolerant Safety Architecture for High-Speed Autonomous Racing. *ACM Transactions on Cyber-Physical Systems* **2026**, *10* (4), 1-29. DOI: 10.1145/3807774.
128. Adadi, A.; Berrada, M. Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI). *IEEE Access* **2018**, *6*, 52138-52160. DOI: 10.1109/access.2018.2870052.
129. Venkatesh, K. P.; Raza, M. M.; Kvedar, J. C. Health digital twins as tools for precision medicine: Considerations for computation, implementation, and regulation. *NPJ Digit Med* **2022**, *5* (1), 150. DOI: 10.1038/s41746-022-00694-7.
130. De Domenico, M.; Allegri, L.; Caldarelli, G.; d'Andrea, V.; Di Camillo, B.; Rocha, L. M.; Rozum, J.; Sbarbati, R.; Zambelli, F. Challenges and opportunities for digital twins in precision medicine from a complex systems perspective. *NPJ Digit Med* **2025**, *8* (1), 37. DOI: 10.1038/s41746-024-01402-3.
131. Bruckner, S.; Kirsten, T.; Schwarz, P.; Cotte, F.; Tsisis, M.; Gilbert, S. The Social Contract for Health and Wellness Data Sharing Needs a Trusted Standardized Consent. *Mayo Clin Proc Digit Health* **2023**, *1* (4), 527-533. DOI: 10.1016/j.mcpdig.2023.07.008.

132. Charles, S.; Mishra, P. Reconfigurable Network-on-Chip Security Architecture. *ACM Transactions on Design Automation of Electronic Systems* **2020**, 25 (6), 1-25. DOI: 10.1145/3406661.