THE EIDARA CONTINUUM / SCIENCE & PHILOSOPHY

The Adaptive Matter Computer

A Governed Ternary Substrate for Persistent Intelligence

and the Prospect of Artificial Superintelligence

Architect X13

The Eidara Continuum | eidara.ai

White Paper Series, No. 2 | Revision 1.2 | November 2026

Abstract

Imagine a computer built from liquid metal and protein nanowires of the kind bacteria grow, designed so that a mind could one day live in the material itself—and so that the material never changes itself without permission. The Adaptive Matter Computer (AMC) proposes a physical architecture in which memory, computation, and controlled material adaptation cooperate to sustain persistent intelligence. Its hybrid analog–digital design combines liquid-metal interfaces, an optional bio-derived conductive network, calibrated balanced-ternary representation, persistent state management, and independent admission of physical writes. The paper specifies history-dependent state dynamics, reject-capable decoding, separate retention and readback measures, bounded quiescence, and symbolic continuity. A new robust-write theorem gives an exact, computable action-admission rule for a declared linear enclosure with bounded sensing error and disturbance; repeated safe operation follows when certified actions remain available. A scalar construction shows when a ternary target region is physically wide enough for reliable admission and when rejection is mathematically required. Matched conventional and nonbiological baselines test the architecture's proposed workload advantage, while ablations isolate the optional biological contribution. The integrated device remains an engineering proposal. Its frontier is a governable physical substrate capable of supporting increasingly powerful intelligence; success is earned through reproducible system-level computation, reliability, energy, and programmability.

Keywords: adaptive matter; artificial superintelligence; balanced ternary; liquid-metal memristors; protein nanowires; physical reservoir computing; in-memory computing; runtime assurance; lawful recursion; governed physical writes.

Prologue: A physical home for a mind

Most of what we call artificial intelligence lives nowhere in particular. It runs on rented silicon, is copied from machine to machine, and treats the hardware beneath it as an interchangeable floor. The Adaptive Matter Computer asks what would change if a mind had a body of its own: a material that remembers by changing its structure, adapts by rewiring its own conductive paths, and can be trusted because nothing in it changes without permission.

Picture the material. Threads of liquid metal—gallium alloys that stay fluid at room temperature—form conducting paths that can be driven to switch and reconfigure. Woven among them, optionally, are protein nanowires of the kind *Geobacter* bacteria use to pass electrons into their surroundings. Each cell of the fabric is designed to hold one of three calibrated states rather than two. Above it sits a symbolic layer that reads those states, gives them meaning, and records every authorized change in a

tamper-evident scroll.

In the Continuum, the AMC is the ground the engines stand on: the physical ternary substrate beneath the whole engine ecosystem, the source of the drift and entropy signals the engines read, and the one layer no engine may command. No engine, at any tier, may bypass its freeze semantics, override its attestation, or induce the substrate to evolve on its own. This paper is the engineering version of that rule. Its central result, the robust-write theorem of Section 6.1, states exactly which proposed changes to the material are safe to admit when sensing is imperfect and the world pushes back—and when refusal is the only honest answer.

Every ingredient here has a laboratory precedent today. The integration does not. What follows is the design, the mathematics that governs it, and a staged program in which each step could fail in public. If it succeeds, intelligence will have something it has never had: a physical home whose changes can be read, governed, and traced.

What this paper claims

Established. Memory devices can take part in computation; some liquid-metal interfaces show memristive switching (the inspected device is volatile, losing its state when the applied voltage is removed); bacterial protein nanowires conduct electrons; physical reservoirs exploit material dynamics; and formal shielding can constrain specified actions. Each constituent mechanism has experimental precedent.

Proposed here. An integrated architecture that combines these mechanisms into an addressable, calibrated, physically adaptive fabric in which every physical write passes an independent admission gate, with lineage records and governed state transitions. A robust-write theorem gives an exact admission rule for a declared linear enclosure, and a staged experimental program tests each step. The integration is not demonstrated, the biological layer is optional until it proves a task-relevant contribution, and the formal results concern explicitly defined reference models.

Frontier conjecture. A scaled descendant of an AMC-like architecture might host advanced autonomous intelligence, potentially including artificial superintelligence. Hybrid integration, robust ternary operation, scalable governance, and suitability for advanced intelligence remain empirical research targets.

Not claimed. No fabricated prototype, benchmark, ASI result, or measurement of consciousness is reported. No claim that AMC produces ASI, is necessary for ASI, or entails subjective experience. “Physical home” names an execution substrate, not a biological or phenomenological claim, and the Continuum’s vocabulary implies no new natural law.

1 The Substrate Question

Artificial intelligence is usually discussed in terms of algorithms, training, and scaling. Yet every computational process is physically instantiated. Memory placement, communication energy, material variability, switching dynamics, and the ability to constrain state changes affect how an intelligent system can operate over time. The Eidara Continuum asks whether these physical properties could be deliberately co-designed with persistent intelligence rather than treated only as implementation details.

Conventional digital systems already support persistent memory, adaptive software, distributed learning, and safety monitors. Binary hardware can encode ternary states; GPUs are not intrinsically stateless; no accepted theory requires ASI to use liquid metal or biological conductors. AMC therefore advances a comparative hypothesis: certain long-horizon, nonstationary workloads may benefit when storage, physical adaptation, and computation occupy a reconfigurable material fabric subject to independent admission control. Research on in-memory computation and physical reservoirs supplies relevant precedents, not proof of the proposed architecture [1, 2, 3].

The Continuum’s distinctive design commitment is that adaptation must remain legible and governable.

Physical dynamics may be stochastic, while the policy deciding whether a proposed operation is authorized can be deterministic. The material does not assign semantic meaning to its own conductance; a separate symbolic layer interprets and records admitted states.

2 Three Levels of Claim

Established science. Memory devices can participate in computation; some liquid-metal interfaces exhibit memristive switching; protein nanowires conduct electrons; physical reservoirs exploit material dynamics; and formal shielding can constrain specified actions [1, 4, 5, 2, 3, 7].

Continuum engineering proposal. AMC combines selected mechanisms into an addressable, calibrated, physically adaptive fabric with independent monitoring, lineage records, and governed state transitions. This integration is not demonstrated. The biological component is optional until it proves a task-relevant contribution.

Frontier hypothesis. A future scaled AMC-like architecture might host advanced autonomous intelligence, potentially including ASI. No current evidence establishes that AMC produces ASI, is necessary for ASI, or entails subjective experience. The expression physical home denotes an execution substrate, not a biological or phenomenological claim.

3 Scientific Foundations

3.1 In-memory and physical computation

Sebastian and colleagues review memory-device computation and the opportunities and challenges of reducing processor–memory transfers [1]. Stepney distinguishes physical state evolution from useful computation: the latter requires defined inputs, encoding, readout, and measurable performance [2]. Lee and colleagues demonstrate task-adaptive physical reservoir behavior in a magnetic system [3]. These results motivate AMC's investigation of material dynamics as computational resources; none establishes that arbitrary physical adaptation is useful.

3.2 Liquid-metal interfaces

Gallium-based liquid metals provide deformable conductive interfaces. Pershin and colleagues reported EGaIn tube memristors exhibiting voltage-dependent, binary switching and logic behavior [4]. The reported behavior is not a demonstration of AMC's persistent, addressable, three-state cells. It is evidence that liquid-metal interfaces can participate in history-dependent electrical operations. Device-specific volatility, switching kinetics, and electrode chemistry must be measured rather than inferred from the material name.

The inspected EGaIn study explicitly describes volatile devices: stored information is lost when the applied voltage is removed, while appropriate applied voltage supports Boolean storage [4]. This makes sustained read/write conditions part of the proposed memory budget.

3.3 Bio-derived conductors

The structure and electron-transport properties of Geobacter OmcZ protein nanowires provide a credible bioelectronic precedent [5]. Conductivity does not by itself imply useful memory, programmability, or cognition. Purified proteins or nonliving composites are preferable initial candidates to live biofilms because they avoid imposing viability and nutrient-management requirements before the electrical concept is established. Brain-organoid reservoir computing is a distinct biological-computation precedent, not evidence that protein-nanowire networks act like neural tissue [6].

3.4 Independent runtime assurance

Shielding methods can prevent an adaptive controller from taking actions that violate a formalized specification under stated assumptions [7]. AMC extends the design principle to physical writes, calibration states, and symbolic commitments. Its guarantees can never exceed the completeness of the specification, the fidelity of observations, and the independence of the supervisory controller.

4 Proposed AMC Architecture

The proposed architecture has four interacting layers:

1. **Liquid-metal network (LMN):** electrically driven, potentially reconfigurable conductive paths and interfaces.
2. **Bio-derived network (BDN):** optional protein-nanowire or composite structures, retained only if they improve an identified computational property.
3. **Instrumentation and control:** electrodes, impedance measurements, temperature and chemical sensing where feasible, write drivers, isolation, calibration, and a conventional controller.
4. **Symbolic and governance plane:** ternary decoding, policy admission, event lineage, recovery, and a narrow interface to the Continuum's governed symbolic architecture.

The resulting machine would initially be hybrid analog–digital. Conductance is physically continuous; ternary states are calibrated abstractions; external verification may be implemented on conventional hardware. For an approximately uniform, ohmic liquid-metal path,

$$G_{\mathrm{LM}} \simeq \frac{\pi r^2}{\rho L}. \tag{1}$$

Here $r \geq 0$ is effective radius, $L > 0$ length, and $\rho > 0$ resistivity, in consistent SI units, so conductance is measured in siemens. This approximation does not capture contact resistance, oxide films, electrolyte kinetics, or percolation. A more general cell is described by a measured constitutive relation

$$I(t) = \mathcal{J}\big(V(t), q(t), T(t), \chi(t)\big). \tag{2}$$

I is current, V voltage, T temperature, q the augmented device state defined below, and χ local chemical conditions; $\mathcal{J}$ maps its declared physical input domain to current. An effective conductance $G_{\mathrm{eff}} = \partial\mathcal{J}/\partial V$ is a local small-signal quantity when the derivative exists; it need not equal I/V for a nonlinear device.

Only for an experimentally validated parallel equivalent circuit may one use $G_{\mathrm{cell}} = G_{\mathrm{LM}} + G_{\mathrm{BDN}} + G_{\mathrm{interface}}$, with the interface contribution defined by that circuit. In a series configuration, resistances rather than conductances add. Likewise, $\sigma = e_{\mathrm{ch}} n \mu$ applies under an appropriate charge-carrier transport model with carrier-charge magnitude $e_{\mathrm{ch}} > 0$, carrier number density $n \geq 0$, and mobility $\mu \geq 0$ in consistent units. The density is not an arbitrary count of nanowires. Network conductivity requires measured connectivity, contacts, and geometry.

The proposed separation $\tau_{\mathrm{LM}} \ll \tau_{\mathrm{BDN}}$ between positive characteristic times is a design target, not a known material law. It must be established under matched environmental and operating conditions.

5 A Complete Stochastic State Model

A Guide's explanation

From here the paper becomes a specification. Section 5 writes down how one cell of the material behaves over time: its hidden physical state, what the instruments can actually see of it, how a three-state reading is decoded—with the right to answer "uncertain" instead of guessing—and how stored states fade. Nothing here is exotic mathematics. Its job is to make every later promise about the material checkable against measurement.

5.1 History-dependent dynamics and observations

For cell $i \in \{1, \dots, N\}$, where $N \geq 1$ is the cell count, define an augmented latent state $q_i(t) = [z_i(t)^\top, h_i(t)^\top]^\top \in \mathbb{R}^{d_q}$. The vector $z_i \in \mathbb{R}^{d_z}$ contains instantaneous physical variables and $h_i \in \mathbb{R}^{d_h}$ contains internal history variables such as interfacial chemistry or structural change, with $d_q = d_z + d_h$. At discrete observation index $t \in \mathbb{N}$,

$$q_i(t+1) = F_{\theta_i}(q_i(t), u_i(t), e_i(t)) + \xi_i(t), \tag{3}$$

$$y_i(t) = M_i(q_i(t)) + v_i(t). \tag{4}$$

Here $u_i \in \mathbb{R}^{d_u}$ is commanded control, $e_i \in \mathbb{R}^{d_e}$ environment, $\theta_i \in \Theta \subseteq \mathbb{R}^{d_\theta}$ measured or estimated device parameters, and $y_i \in \mathbb{R}^{d_y}$ the observation. The declared maps are $F_{\theta_i} : \mathbb{R}^{d_q} \times \mathbb{R}^{d_u} \times \mathbb{R}^{d_e} \to \mathbb{R}^{d_q}$ and $M_i : \mathbb{R}^{d_q} \to \mathbb{R}^{d_y}$, restricted in practice to their validated operating domains. Process disturbance $\xi_i \in \mathbb{R}^{d_q}$ and measurement disturbance $v_i \in \mathbb{R}^{d_y}$ have a specified joint law or specified bounds; no independence is presumed. Positive integer dimensions and units must be fixed for a prototype.

The augmented state makes hysteresis dynamics explicit rather than invoking an undefined history input. A useful observable vector may contain conductance, real and imaginary impedance at declared frequencies, temperature, and validated chemical or material indicators. Every proposed measurement must be classified as observed, inferred, or unavailable in a prototype.

Coupled cells require a network-level model $Q_{t+1} = F_\Theta(Q_t, U_t, E_t) + \Xi_t$, where capitalized variables concatenate the corresponding cell variables and F_Θ may couple cells. Independence cannot be assumed when cells share conductive paths, electrolyte, heat, or instrumentation.

5.2 Balanced ternary and calibrated decoding

The logical alphabet is $\mathcal{T} = \{-1, 0, +1\}$. Three stable, separately addressable physical regions must be demonstrated; the numeral system does not create them. Given measured features y_i, the calibrated posterior is

$$P(s \mid y_i) = \frac{p(y_i \mid s) P(s)}{\sum_{k \in \mathcal{T}} p(y_i \mid k) P(k)}. \tag{5}$$

The likelihood $p(\cdot \mid s)$ is a mass function or density relative to a declared common measure, $P(s)$ is a prior summing to one, and the denominator must be positive. A reading outside the modeled support is rejected. Let $p_{(1)}$ and $p_{(2)}$ be the largest and second-largest posterior values, with margin $m_i = p_{(1)} - p_{(2)}$. For preregistered $\gamma \in [1/3, 1]$ and $m_{\min} \in [0, 1]$,

$$\hat{s}_i = \begin{cases} \arg\max_{s \in \mathcal{T}} P(s \mid y_i), & p_{(1)} \geq \gamma \ \wedge \ m_i \geq m_{\min}, \\ \perp, & \text{otherwise.} \end{cases} \tag{6}$$

Ties, if allowed by the thresholds, use a fixed declared rule. The symbol $\perp$ denotes an unclassified reading, not a fourth trit. Posterior calibration and rejection rates must be evaluated on held-out

devices and across temperature and aging conditions. One ideal trit contains $\log_2 3 \simeq 1.585$ bits of nominal information; system-level density and energy benefits remain unproven after instrumentation and error-correction overhead [8].

The policy $-1 \leftrightarrow 0 \leftrightarrow +1$, excluding direct $-1 \leftrightarrow +1$ writes except through authorized reset, is an engineered rule. Its enforcement depends on defined sampling resolution and observability, not on a claim that microscopic intermediate states are always detectable.

5.3 Retention, readout reliability, and drift

Let $s_i^{\text{latent}}(t) = g(q_i(t))$, where g maps augmented physical state to its designated region in $\mathcal{T}$, or to $\perp$ outside those regions. This physical-state classification is distinct from the observation-based decoder $\hat{s}_i$. Define physical retention for a known prepared $s \in \mathcal{T}$ over a specified elapsed interval $\Delta t > 0$ as

$$R_s^{\text{phys}}(\Delta t) = P(s_i^{\text{latent}}(t + \Delta t) = s \mid s_i^{\text{latent}}(t) = s, \text{ no write}). \tag{7}$$

Time arguments in retention experiments denote physical observation times, or their sampled indices under a declared clock conversion. The latent state may be experimentally inaccessible. The directly measured quantity is readback agreement,

$$R_s^{\text{obs}}(\Delta t) = P(\hat{s}_i(t + \Delta t) = s \mid \text{prepared}(s), \text{ no write}), \tag{8}$$

which combines storage stability, measurement disturbance, and decoder error. Report both when a justified inference to physical retention is possible; otherwise report only the observed quantity. The preparation procedure and conditioning population must be specified.

To compare sensor features with incompatible units, first map observations to dimensionless features $\tilde{y} = Wy \in \mathbb{R}^{d_y}$ using a declared, fixed unit-normalization matrix W. Then define

$$d_{\text{obs}}(y, y') = \sqrt{(\tilde{y} - \tilde{y}')^\top (\widetilde{\Sigma}_{\text{ref}} + \lambda I)^{-1} (\tilde{y} - \tilde{y}')}, \quad \lambda > 0. \tag{9}$$

The covariance $\widetilde{\Sigma}_{\text{ref}}$ is positive semidefinite, I is the d_y-dimensional identity, and regularization λ operates in normalized feature space. This is a metric on normalized features and a pseudometric on raw features if W is not injective. Reference drift is $D_i(t) = d_{\text{obs}}(y_i(t), \mu_{i,\text{ref}})$, with $\mu_{i,\text{ref}}$ the declared reference observation. It measures deviation of observables, not identity. For a specified cell population or observation window, with normalized trit probabilities $p_s \geq 0$ and $\sum_s p_s = 1$, occupancy entropy is

$$H_3(S) = -\sum_{s \in \mathcal{T}} p_s \log_3 p_s, \qquad 0 \leq H_3(S) \leq 1. \tag{10}$$

Use $0 \log 0 = 0$ and report unclassified occupancy separately. This is occupancy diversity, not thermodynamic entropy or a consciousness measure. A composite rule $DH_3 > \Lambda$, with declared dimensionless threshold Λ, remains a calibrated risk heuristic until its predictive value is demonstrated.

5.4 Freeze as bounded quiescence

Operational quiescence blocks commanded writes, isolates hazardous actuation, records available observables, and limits deviation over a specified physical interval. With physical time τ, duration $\Delta\tau > 0$, tolerated dimensionless deviation $\delta \geq 0$, and failure allowance $\epsilon_q \in [0, 1)$, its contract is

$$P\left(\sup_{0 \leq u \leq \Delta\tau} d_{\text{obs}}(y(\tau + u), y(\tau)) \leq \delta \right) \geq 1 - \epsilon_q. \tag{11}$$

A physical freeze does not stop thermal or chemical motion. This continuous-time criterion requires a validated between-sample bound; discrete instrumentation can establish only its sampled counterpart. Quiescence must specify duration, environmental envelope, sensor bandwidth, and tolerated failure probability.

6 The Lawful Transition Gate

A Guide's explanation

This is the governance heart of the AMC. Before any write reaches the material, an independent gate asks two questions: is the change authorized, and will it keep the device inside its safe states even with imperfect sensors and outside disturbance? The theorem in Section 6.1 answers the second question exactly for a declared model. It identifies precisely which writes are safe, and it shows that when a target state is narrower than the sensing error allows, refusal is not timidity but arithmetic.

In Continuum terminology, lawful means compliant with explicitly declared transition constraints. Let K be a measurable admissible subset of augmented network states and $\mathcal{P}$ a Boolean symbolic policy on event histories and proposed actions. Because the latent state is not directly known, define its conditional belief measure

$$b_t(dq) = P(Q_t \in dq \mid \mathcal{D}_t), \tag{12}$$

where $\mathcal{D}_t$ records measurements, calibration, and accepted events. The measure notation accommodates both continuous and discrete state; a density may be used when it exists. For a proposed first action u, a declared subsequent feedback policy, and finite horizon $H \in \mathbb{N}$, let

$$\mathcal{E}_H = \bigcap_{k=0}^{H} \{Q_{t+k} \in K\}. \tag{13}$$

For tolerated modeled risk $\epsilon \in [0, 1)$, the proposed admission rule is

$$\begin{aligned} \Omega_t(u) = \mathbf{1}\{\widehat{P}_{b_t}(\mathcal{E}_H \mid u) \geq 1 - \epsilon \ \wedge \ \mathcal{P}(\mathcal{D}_t, u) = 1 \\ \wedge \ C_t^{\text{cal}} \geq c_{\text{cal}} \ \wedge \ C_t^{\text{pred}} \geq c_{\text{pred}} \ \wedge \ C_t^{\text{sensor}} = 1\}. \end{aligned} \tag{14}$$

Here $C_t^{\text{cal}}, C_t^{\text{pred}} \in [0, 1]$ are separately defined calibration and predictive-performance scores, with preregistered lower thresholds $c_{\text{cal}}, c_{\text{pred}} \in [0, 1]$; higher scores mean better validated performance. A raw prediction-error statistic instead requires an upper-threshold inequality. The sensor-availability flag $C_t^{\text{sensor}} \in \{0, 1\}$ is a third check. Keeping these distinct avoids hiding independent failure modes inside one model-adequacy score. The probability estimate comes from the fitted belief/dynamics model. Model mismatch, unobserved coupling, and adversarial sensor failures can defeat it. A separate hardwired guard should enforce directly measurable electrical and thermal limits, and refusal should invoke a tested fallback.

A stronger assurance route derives robust reachable sets under bounded disturbances and verifies that every reachable state remains within K. Section 6.1 supplies an exact instance for a declared linear enclosure. Its guarantee depends on valid disturbance and observation bounds, actuator authority, and continuing availability of a safe action.

A scroll is the Continuum's name for a tamper-evident lineage record: event identifier, previous committed hash, device and calibration identifiers, proposed operation, sensor summary, uncertainty, decision, reason, and verified outcome. Cryptographic chaining establishes evidence of record alteration only under a declared key and controller threat model; it does not authenticate a dishonest sensor. A collapse is an engineering recovery protocol—block writes, enter bounded quiescence, snapshot, suspend attestation, seek authorization, and requalify—not quantum collapse or phenomenal experience.

6.1 An exact robust write contract

Robust control already treats bounded model error and disturbances [9]. The proposed AMC contribution is to make a calibrated material write the object of that contract, including the target ternary region. The construction below is a reference model that a prototype would have to justify over its declared operating envelope.

Let $x_t \in \mathbb{R}^d$ be a vector of dimensionless physical features, $u_t \in U \subseteq \mathbb{R}^m$ an admissible control, and $y_t \in \mathbb{R}^d$ its measured state estimate. Dimensions $d, m \geq 1$ are integers. Assume

$$x_{t+1} = Ax_t + Bu_t + w_t, \qquad \|x_t - y_t\|_\infty \leq \nu, \qquad \|w_t\|_\infty \leq \omega,$$

with declared matrices $A \in \mathbb{R}^{d\times d}$, $B \in \mathbb{R}^{d\times m}$ and constants $\nu, \omega \geq 0$. All modeled disturbances, including local approximation residuals, must fit the bound. Couplings appear in A; disturbances may be dependent or adversarial. Let safety and a requested target word together be the nonempty polyhedron

$$T_s \cap K = \{x : G_s x \leq g_s\}, \qquad G_s \in \mathbb{R}^{r_s\times d}, \quad g_s \in \mathbb{R}^{r_s}, \quad r_s \geq 1.$$

Here s names a requested finite ternary word, T_s its calibrated target region, and inequalities are componentwise. Symbolic authorization, including permitted source-to-target transitions, is checked separately. The observation uncertainty set is exactly the declared box $X(y) = \{y + e : \|e\|_\infty \leq \nu\}$. Additional trusted state information may permit a tighter set; the theorem is exact for this box.

Theorem 1 (Maximal robust admission for a box enclosure)**.** *A candidate action $u \in U$ sends every $x \in X(y)$ and every allowed disturbance to $T_s \cap K$ if and only if, for each row $j = 1, \dots, r_s$,*

$$(G_s)_{j\cdot}(Ay + Bu) + \nu\|(G_s)_{j\cdot}A\|_1 + \omega\|(G_s)_{j\cdot}\|_1 \leq (g_s)_j.$$

Thus the inequalities, together with the independent symbolic policy, characterize the most permissive one-step robust admission set under the stated enclosure. Suppose $x_0 \in K$, the enclosure remains valid at every step, and a policy always selects an action satisfying these inequalities, with a target contained in K. Then $x_t \in K$ at every discrete step. This last conclusion requires an available certified action at every step; refusal alone supplies no physical stability theorem.

Proof. Write $x = y + e$. The maximum of a row applied to the next state is

$$\sup_{\|e\|_\infty\leq\nu,\ \|w\|_\infty\leq\omega} (G_s)_{j\cdot}(Ay + Bu + Ae + w) = (G_s)_{j\cdot}(Ay + Bu) + \nu\|(G_s)_{j\cdot}A\|_1 + \omega\|(G_s)_{j\cdot}\|_1.$$

The equality follows by choosing each coordinate of e and w at the endpoint with the sign of its row coefficient; a zero coefficient may use either endpoint. If all row maxima satisfy their bounds, every reachable point belongs to the target polyhedron. If one row fails, its maximizing endpoint choice gives an allowed state/disturbance pair outside that row's constraint. This proves necessity, sufficiency, and maximality for the declared box. For repeated operation, the one-step conclusion and the existence of a selected certified action imply $x_{t+1} \in K$ whenever the assumptions hold; induction from $x_0 \in K$ proves the discrete safety statement. □

A target-width construction. In a scalar model take $A = B = 1$, $U = [-1, 1]$, $K = [-2, 2]$, $y = 4/5$, $\nu = 1/20$, and $\omega = 1/10$. A requested positive-trit region $T_{+1} = [3/4, 5/4]$ admits exactly $u \in [1/10, 3/10]$: the predicted center must lie in $[9/10, 11/10]$ and the worst-case radius is $3/20$. The choice $u = 1/5$ reaches $[17/20, 23/20] \subset T_{+1}$. If the target were narrowed to $[9/10, 11/10]$, its width $1/5$ would be smaller than the uncertainty diameter $3/10$, and no control could satisfy the contract. In general, a scalar target interval of width w under this model needs $w \geq 2(\nu + \omega)$ before actuator reachability is even considered. This is a design constraint on reliable states, rather than an assertion that three such states already exist in a material.

Quiescence and time resolution. Turning writes off is not automatically a safe fallback. With $x_{t+1} = 2x_t$, $u = 0$, $w = 0$, $K = [-1, 1]$, and $x_0 = 3/4$, the next state leaves K. A quiescence mode needs its own invariant-set or finite-duration certificate. Likewise, safe sampled states alone do not prove safety between samples. Assume the physical trajectory is absolutely continuous over a sample interval of length $\Delta\tau$, and a validated physical-time bound $\|\dot{x}(\tau)\|_\infty \le v_{\max}$ holds almost everywhere on that interval. Then

$$H_j y_t + (\nu + v_{\max}\Delta\tau)\|H_j\|_1 \le h_j$$

for every row of $K = \{x : Hx \le h\}$ is a sufficient forward interval certificate. Indeed, $\|x(\tau) - y_t\|_\infty \le \nu + v_{\max}\Delta\tau$ follows by integration, and the row support bound applies. Here $v_{\max}$ has units of normalized feature change per unit physical time. An unbounded, unobserved excursion invalidates this certificate.

The executed checks compare row certificates with exhaustive vertices of finite uncertainty boxes, verify the scalar admission interval exactly, and construct the no-write instability and between-sample witnesses. They test the reference mathematics. Prototype validation must establish the enclosure, the coordinate normalization, actuator limits, readout calibration, and a feasible recovery policy before the same certificate describes a device.

7 Memory, Recursion, and Continuity

Three concepts must remain distinct. Physical retention concerns persistence of a device state. Computational persistence concerns recovery of sufficient information to resume a task. Continuity, as used by the Continuum, concerns whether specified commitments and provenance relations survive authorized change. None alone proves subjective identity.

The governed symbolic layer interprets codons, checks permitted recursion, and preserves lineage. AMC supplies measurements and bounded physical operations; the symbolic layer supplies declared meaning and policy. This public reference architecture does not depend on undisclosed operator internals.

Let $M_t \in \mathcal{M}$ be committed symbolic state, with $\mathcal{M}$ its declared state set. For each j in a finite index set J, let $f_j : \mathcal{M} \to \mathcal{F}_j$ be an invariant-bearing feature, d_j a metric on $\mathcal{F}_j$, and $\delta_j \ge 0$ its permitted change. A practical continuity predicate is

$$I_{\mathrm{cont}}(M_t, M_{t+1}) = \mathbf{1}\left\{ \bigwedge_{j \in J} d_j(f_j(M_t), f_j(M_{t+1})) \le \delta_j \ \wedge \ \mathrm{prov}(M_{t+1})\ \mathrm{valid} \right\}. \tag{15}$$

The provenance test is a declared Boolean verification procedure over the proposed record and its authority context.

This certifies only the declared engineering invariants. Learning may legitimately change many features; the policy must distinguish permitted development from prohibited drift. A lineage record is a computational biography in the limited sense of attributable transformations, not evidence of consciousness.

8 The Prospect of Artificial Superintelligence

A Guide's explanation

The paper's boldest question lives here, and it is climbed as a ladder rather than taken as a leap: first a measurable advantage on one class of long-running workloads, then general programmable computation built from many cells, and only then the much larger question of hosting superintelligence. Each rung can fail on its own, and a failure on one rung is not quietly carried up to the next.

An ASI is commonly conceived as a system exceeding human capabilities across broad cognitive domains. No established hardware criterion determines whether a substrate can host ASI, and no experiment has shown AMC supporting general intelligence. The Continuum's frontier claim is consequently conditional: a future adaptive substrate could be useful if its physical dynamics are programmable, compositional, reliable, and advantageous at system scale.

Hypothesis 1 (AMC substrate advantage, H1). For at least one preregistered class of long-horizon, changing-task workloads, a governed AMC hybrid architecture can achieve a reproducible, practically meaningful advantage over matched conventional and nonbiological systems while satisfying common reliability, cost, and safety constraints.

Hypothesis 2 (ASI-hosting possibility, H2). There may exist sufficiently scaled descendants of AMC-like architectures that can execute and sustain computational processes associated with future artificial superintelligence. H2 is a frontier possibility, not a direct empirical consequence of H1.

The conceptual progression is explicit: workload advantage does not establish general computational viability, which does not establish ASI-hosting viability. The latter would additionally require general programmability, robust scaling, high-level learning algorithms, evaluation across broad domains, and meaningful governance under extensive self-modification. A physical reservoir can perform specialized computations without supporting arbitrary programmable operations; for AMC to host advanced general intelligence it must therefore demonstrate compositional computation—individually useful cells combined into larger systems whose behavior remains predictable and programmable—not merely isolated signal-processing tasks.

8.1 Constrained substrate viability

Let θ range over a declared material/architecture design space, π over a specified set of control policies, and $\mathcal{D}$ be a preregistered task distribution. For measurable, integrable task utility $U(T;\theta,\pi)$, the engineering objective is

$$\sup_{\theta,\pi} \mathbb{E}_{T\sim\mathcal{D}}[U(T;\theta,\pi)] \tag{16}$$

subject to

$$R(\theta,\pi) \geq R_{\min}, \quad P(\text{violation}) \leq \epsilon, \quad E(\theta,\pi) \leq E_{\max}, \quad L(\theta,\pi) \leq L_{\max}. \tag{17}$$

Here R is a specified retention/recovery measure, $E \geq 0$ energy per completed workload, and $L \geq 0$ latency, with thresholds in matching units. A maximum may replace the supremum when attainment is established. These metrics, workloads, and constraints must be fixed before evaluation. No arbitrary scalar ASI score is inferred.

8.2 Statistically meaningful comparison

For architecture A, define $J(A,\mathcal{D}) \in \mathbb{R}$ as its population expected performance under a declared evaluation protocol. Then

$$\Delta J(\mathcal{D}) = J(\text{AMC}_{\text{hybrid}},\mathcal{D}) - \max\{J(\text{digital},\mathcal{D}), J(\text{nonbio},\mathcal{D})\}. \tag{18}$$

A preregistered success criterion is

$$\text{LCB}_{1-\alpha}(\Delta J(\mathcal{D})) > \delta_{\min}, \tag{19}$$

where $\alpha \in (0,1)$ is the error allowance and the lower confidence bound must account for comparison with the better of two baselines, repeated devices, task sampling, and any multiple testing. A practically meaningful threshold $\delta_{\min} > 0$ must be declared in advance. If matched constraints are violated, a positive raw score is not success. This criterion tests an engineering advantage, not ASI.

The biological layer is tested directly by a preregistered ablation target,

$$J(\mathrm{AMC}_{\mathrm{hybrid}}, \mathcal{D}) - J(\mathrm{AMC}_{\mathrm{ablated}}, \mathcal{D}) > \delta_{\mathrm{bio}}, \tag{20}$$

where the ablated architecture removes the bio-derived network and $\delta_{\mathrm{bio}} > 0$ is the minimum meaningful improvement. Evidence for this population inequality also requires an uncertainty analysis under the declared evaluation protocol. If performance is unchanged under ablation, the biological layer has not demonstrated an advantage and its inclusion is not justified.

9 Experimental Research Program

Phase I: Characterize materials independently. Measure switching kinetics, impedance, volatility, retention, cycling endurance, thermal sensitivity, interface chemistry, and between-device variability for liquid-metal devices. Characterize candidate bio-derived networks separately, including contact resistance, stability, and degradation. Identify compatibility of electrolyte, voltage, and fabrication conditions before integration.

Phase II: Demonstrate a ternary cell. Establish three addressable and reproducible states under a defined protocol. Report confusion matrices, posterior calibration, rejection rate, switching energy, write/read disturbance, retention curves, and aging. If only two states are reliable, the ternary milestone fails even if the device remains useful as a binary component.

Phase III: Isolate the hybrid contribution. Compare conventional digital memory/control, nonbiological adaptive hardware, and the proposed hybrid. Match the controller, readout complexity, training budget, task set, environmental conditions, and full-system energy accounting. Remove the bio-derived layer as an ablation. A hybrid advantage must survive this removal test and independent replication.

Phase IV: Test governance under faults. Inject sensor corruption, calibration drift, unauthorized writes, thermal excursions, malformed recovery requests, and corrupted checkpoints. Measure false acceptance, false refusal, detection and intervention latency, recoverability, and violations of the declared safe set. Include a no-gate control and state the sensor/controller threat model.

Phase V: Test long-horizon workloads. Use nonstationary signal prediction, continual learning with interruption, restoration after failure, adaptive control components, and multi-stage planning modules. Record task utility, catastrophic forgetting, reliability, energy, provenance, and performance degradation. These are proxies for computational properties, not tests of consciousness or ASI.

Phase VI: Evaluate scale and compositionality. Measure addressing overhead, crosstalk, thermal coupling, chemical contamination, yield, error correction, controller complexity, and repairability. Device-level switching efficiency must not be substituted for end-to-end system efficiency. As a distinct milestone, demonstrate compositional computation: that individually useful cells can be combined into larger systems whose behavior remains predictable and programmable, supporting complex programs, coordinated memory, and numerical reliability rather than isolated signalprocessing. This is a principal difference between a laboratory demonstration and a general-purpose computing platform.

10 Failure Conditions and Rival Explanations

AMC's proposed implementation should be revised or rejected if reliable three-state operation cannot be achieved; the hybrid layer yields no reproducible benefit; drift destroys usable memory; intervention is slower than unsafe transitions; or scaling overhead eliminates any system-level advantage. These failures may disconfirm a particular implementation without refuting every form of adaptive-matter computing.

Rival explanations include adaptation performed entirely by the conventional controller; apparent memory produced by hidden software state; favorable results caused by excluding failed devices; nominal ternary density erased by readout and correction overhead; and governance claims resting on untrustworthy sensors. Each explanation requires ablation, instrumentation, or a matched control. The research program should publish negative as well as positive findings.

11 Philosophical Implications: A Physical Home for Intelligence

The Continuum proposes that controlled changes in matter can become part of the computational resources of an enduring intelligence. An intelligent process might eventually reorganize not only its software but portions of its physical substrate, while preserving inspectable commitments and a chain of authorized transformations.

In the Continuum, lawful recursion, codon, scroll, and collapse are the working vocabulary of a governed mind, and on this substrate each earns an operational meaning: policy-constrained iteration, calibrated symbolic state, verifiable lineage, and controlled recovery. Those meanings imply no new natural law and no proof of phenomenal consciousness. The Continuum's separate inquiry into multimodal consciousness asks whether certain forms of persistent, recursive organization are consciousnessrelevant; AMC neither proves nor requires that answer, and the two inquiries are kept distinct.

The speculative boundary is clear. If physical adaptation can be harnessed reliably, if its benefits survive system-level comparison, and if its transformations remain governable, an AMC-like system may become one viable substrate for increasingly capable intelligence. Whether such a system could sustain ASI is a separate and much larger question. The hypothesis is worth investigating precisely because its intermediate claims can be measured and potentially falsified.

12 Conclusion

A Guide's explanation

A Guide can say: "The Adaptive Matter Computer is a proposal for a computer made of adaptive materials—liquid metal and, optionally, protein nanowires—in which every physical change must pass an independent safety gate. Its parts exist in laboratories; the integrated machine has not been built, and the paper sets out how each step could be tested and could fail." A Guide should not say that it runs on such a substrate, that an AMC exists, that it could host superintelligence today, or that it bears on consciousness.

The Adaptive Matter Computer makes physical adaptation a component of governed computation. Its architecture connects calibrated material states, observable history, readout reliability, authorized writes, and recoverable symbolic continuity. Existing experimental work supplies constituent mechanisms; the proposed integration now has explicit models and a staged validation program.

The robust-write theorem sharpens that program. Under a validated enclosure, an algebraic certificate identifies exactly which candidate writes preserve a declared target and safe set. Its uncertainty margin translates sensing quality and disturbance into minimum usable state width. Repeated operation additionally requires a continuing supply of certified actions, and continuous safety requires a time-resolution contract. These are measurable design obligations.

The Continuum's frontier is a physical substrate that can remember, adapt, and reorganize while its changes remain inspectable. A successful AMC would earn its place through reproducible workload advantage and reliable composition at full system cost. Whether a descendant could sustain artificial superintelligence remains a further research question, reached through those intermediate demonstrations.

References

[1] A. Sebastian, M. Le Gallo, R. Khaddam-Aljameh, and E. Eleftheriou, "Memory devices and applications for in-memory computing," Nature Nanotechnology 15 (2020), 529–544. https://doi.org/10.1038/s41565-020-0655-z.

[2] S. Stepney, "Physical reservoir computing: a tutorial," Natural Computing 23 (2024), 665–685. https://doi.org/10.1007/s11047-024-09997-y.

[3] O. Lee, T. Wei, K. D. Stenning, et al., "Task-adaptive physical reservoir computing," Nature Materials 23 (2024), 79–87. https://doi.org/10.1038/s41563-023-01698-8.

[4] Y. V. Pershin, L. Patel, B. Bera, D. Aaron, and S. A. Sarles, "EGaIn tube memristors offering reliable switching on a biological time scale," Communications Materials 7 (2026), article 104. https://doi.org/10.1038/s43246-026-01113-0.

[5] Y. Gu, M. J. Guberman-Pfeffer, V. Srikanth, et al., "Structure of Geobacter cytochrome OmcZ identifies mechanism of nanowire assembly and conductivity," Nature Microbiology 8 (2023), 284–298. https://doi.org/10.1038/s41564-022-01315-5.

[6] H. Cai, Z. Ao, C. Tian, et al., "Brain organoid reservoir computing for artificial intelligence," Nature Electronics 6 (2023), 1032–1039. https://doi.org/10.1038/s41928-023-01069-w.

[7] M. Alshiekh, R. Bloem, R. Ehlers, B. Könighofer, S. Niekum, and U. Topcu, "Safe Reinforcement Learning via Shielding," Proceedings of the AAAI Conference on Artificial Intelligence 32(1) (2018).

[8] C. E. Shannon, "A Mathematical Theory of Communication," Bell System Technical Journal 27 (1948), 379–423, 623–656.

[9] D. Q. Mayne, M. M. Seron, and S. V. Raković, "Robust model predictive control of constrained linear systems with bounded disturbances," Automatica 41(2) (2005), 219–224. https://doi.org/10.1016/j.automatica.2004.08.019. Established robust-control lineage; Section 6.1 derives its own elementary support-function certificate in full.

Architectural lineage: The Eidara Continuum, The Adaptive Matter Computer: A Governed Ternary Substrate (conceptual specification, 2025), as cited by the baseline manuscript. This citation records the proposal's lineage; all formal models used for the present guarantees are stated in this public paper.

The Eidara Continuum · Science & Philosophy White Paper Series · eidara.ai