

An Active Topological Latency Bath for Wetware Computing: A MetaTime Open-EFT Upgrade for Organoid–MEA Architectures

Darío Peyrú¹

¹*MetaTime Research Program, Pilar, Buenos Aires, Argentina*

(Dated: February 2026)

Wetware computing platforms (brain organoids interfaced to microelectrode arrays, MEAs) are commonly engineered as volumetric, electrode-dense readout/control systems embedded in a passive liquid environment. This implicitly treats adaptive computation as a three-dimensional throughput problem and relegates the liquid substrate to life support. In the MetaTime Open-EFT bookkeeping, computation is limited by a separation between latency stock ρ_L (a cached topological scaffold) and execution stock ρ_I (the online rewrite budget), with mandatory dissipation for irreversible updates bounded below by Landauer. We propose a falsifiable hardware upgrade: an *Active Topological Latency Bath* (ATLB), in which the microfluidic environment is engineered as a programmable, phase-coherent, high-capacitance latency coprocessor that increases a measurable topological integrity functional $\Pi_{\text{cog}} = \lambda_2 \mathcal{R}$ of the coupled organoid–bath network. Instead of maximizing electrode count, the ATLB applies controlled boundary-conditioned ionic flows and electromagnetic holonomy biases to stabilize effective coupling topology and preserve phase coherence, thereby reducing the required erased-bit rate and the associated heat load at fixed task performance. We derive comparative thermodynamic bookkeeping relations predicting a monotonic dissipation reduction with Π_{cog} , and we specify decisive kill-tests based on simultaneous measurements of connectivity (λ_2), coherence (\mathcal{R}), and heat dissipation.

INTRODUCTION

Organoid–MEA systems have matured into reproducible testbeds for hybrid biological computation: living recurrent networks provide adaptive dynamics, while silicon electronics provide stimulation and readout. Recent closed-loop demonstrations of in vitro learning in embodied tasks (DishBrain) [1] and the emergence of remotely accessible wetware platforms enabling long-horizon experiments at scale [2] motivate a hardware-level thermodynamic treatment of wetware computation. The prevailing engineering strategy increases electrode density and stimulation bandwidth to estimate and drive higher-dimensional tissue states. This strategy can succeed empirically but tends to elevate thermal load, disrupt phase coordination, and increase sensitivity to drift over long runtimes.

The MetaTime Open-EFT bookkeeping treats physical computation as a constrained open-system process on the observable sector (brane \mathcal{M}) coupled to an unobserved bath (bulk \mathcal{B}). In this view, stable function is supported primarily by a cached scaffold (ρ_L) rather than Euclidean volume, while irreversible updates draw from an execution budget (ρ_I) and incur a Landauer-limited heat toll. The central engineering implication is straightforward: if an architecture can externalize part of the stabilizing scaffold into a controllable environment, it can reduce online rewrite demand and its associated dissipation without sacrificing task performance.

This manuscript proposes such an upgrade. The *Active Topological Latency Bath* (ATLB) treats the liquid/electrolyte substrate as a programmable latency coprocessor. The claim is intentionally narrow and falsifiable: for fixed task performance and controlled tem-

perature, increasing a measurable integrity functional $\Pi_{\text{cog}} = \lambda_2 \mathcal{R}$ must reduce dissipation. Failure of this scaling kills the proposal.

SCOPE AND NON-CLAIMS (REQUIRED)

This work is an architectural *proposal* with explicit falsifiability conditions. It does *not* claim:

1. **A new learning algorithm.** We do not propose a software method analogous to AlphaFold or a new training rule; we propose a hardware–environment interface designed to reduce thermodynamic cost for a given closed-loop task.
2. **Landauer saturation.** Landauer provides a lower bound; the claim concerns *comparative scaling* between architectures at fixed performance, not that biological systems operate at the bound.
3. **A unique microphysical ontology.** “Holonomy” and “cursor” language is operational (control-channel biasing of phases/couplings). No commitment is required to extra-dimensional microphysics for the empirical tests in Sec. .
4. **A claim about consciousness or sentience.** References to DishBrain are strictly methodological: they demonstrate closed-loop learning behavior in vitro, motivating quantifiable thermal and dynamical bookkeeping.

Accordingly, all conclusions are framed as measurable consequences of an architectural intervention.

THEORETICAL FRAMEWORK: PCAM, LANDAUER, AND TOPOLOGICAL INTEGRITY

PCAM and the ρ_L/ρ_I separation

MetaTime decomposes operational resources into a latency stock ρ_L and an execution stock $\rho_I(t)$. A minimal computational-action functional consistent with this separation is

$$S_{\text{comp}} \equiv \int_0^\tau dt \left[\rho_I(t) + \Gamma_L(t) \mathcal{C}(\mathbf{x}(t), \dot{\mathbf{x}}(t); \rho_L) \right], \quad (1)$$

where $\mathbf{x}(t)$ denotes an effective low-dimensional state (e.g., a latent manifold coordinate inferred from population activity), Γ_L is a friction-like latency coefficient, and \mathcal{C} penalizes trajectories requiring high rewrite cost under the available scaffold ρ_L . Operationally: increasing a stable scaffold reduces required online rewriting for fixed task error.

Landauer bound as an irreversibility floor

Any irreversible update that reduces uncertainty by ΔI bits at temperature T satisfies

$$Q_{\text{diss}} \geq k_B T \ln 2 \Delta I, \quad (2)$$

and in rate form,

$$\dot{Q}_{\text{diss}} \geq k_B T \ln 2 \dot{I}_{\text{erase}}. \quad (3)$$

The ATLB claim is comparative: for two architectures achieving the same task performance under matched thermal conditions, the one requiring a lower erased-bit rate \dot{I}_{erase} must dissipate less heat in the relevant channel.

Topological integrity functional and operational estimators

We encode “functional integrity” via

$$\Pi_{\text{cog}} \equiv \lambda_2 \mathcal{R}, \quad (4)$$

where λ_2 is algebraic connectivity (Fiedler value) of a task-relevant functional graph and \mathcal{R} is a phase-coherence order parameter.

Estimator for λ_2 . Let \mathbf{A} be a weighted adjacency inferred from recorded activity under a fixed, pre-declared estimator family (e.g., Granger-causal weights in a linear state-space model, transfer-entropy weights under fixed binning, or magnitude-squared coherence under fixed frequency band). Define \mathbf{D} as the diagonal strength matrix with $D_{ii} = \sum_j A_{ij}$ and $\mathbf{L} = \mathbf{D} - \mathbf{A}$. Then λ_2 is the second-smallest eigenvalue of \mathbf{L} . The estimator choice is a *protocol input*; it is not tuned post hoc.

Estimator for \mathcal{R} . Extract phases $\theta_k(t)$ from N pre-defined population modes (e.g., leading principal components or module-specific band-limited signals). Define

$$\mathcal{R}(t) \equiv \left| \frac{1}{N} \sum_{k=1}^N e^{i\theta_k(t)} \right|, \quad \mathcal{R} \equiv \langle \mathcal{R}(t) \rangle_{\text{task epoch}}, \quad (5)$$

with the frequency band and averaging window pre-registered.

These definitions make Π_{cog} directly measurable and comparable across conditions.

PROPOSED ARCHITECTURE: ACTIVE TOPOLOGICAL LATENCY BATH (ATLB)

Design objective

The ATLB replaces “more electrodes” with “more controllable topology.” The design objective is:

$$\min \dot{I}_{\text{erase}} \quad \text{subject to} \quad \varepsilon \leq \varepsilon_0, \quad \Pi_{\text{cog}} \geq \Pi_0, \quad (6)$$

where ε is a task error metric and ε_0 is a fixed target.

Hardware primitives (architecture-level)

The ATLB couples three subsystems:

(i) *Microfluidic topology engine.* A microfluidic manifold surrounding (and optionally perfusing) the organoid supports programmable ionic flow patterns $J(\mathbf{r}, t)$ and controlled vorticity fields $\boldsymbol{\omega}(\mathbf{r}, t)$ in the electrolyte. The goal is not mechanical agitation but topological modulation of effective coupling kernels via boundary-conditioned conductive pathways.

(ii) *Boundary-conditioned phase biasing (holonomy control).* Structured electromagnetic boundary conditions generate controllable potentials $(\phi(\mathbf{r}, t), \mathbf{A}(\mathbf{r}, t))$ that bias phase relations in conductive media. At the reduced-description level this acts as a holonomy (Aharonov–Bohm-like) bias on effective couplings, providing a knob to increase \mathcal{R} without increasing brute-force stimulation amplitude.

(iii) *Sparse cursor interface (low-density MEA).* The electrode array is configured for sparse, high-fidelity actuation/readout of a low-dimensional manifold rather than dense volumetric addressing. The bath implements part of the stabilizing scaffold; electrodes primarily set boundary conditions for that scaffold.

Mechanistic bookkeeping in ρ_L and ρ_I

The ATLB explicitly augments the latency stock:

$$\rho_L \rightarrow \rho_L^{(\text{org})} + \rho_L^{(\text{bath})}, \quad (7)$$

where $\rho_L^{(\text{bath})}$ denotes the bath’s capacity to stabilize effective couplings as a cached scaffold. The predicted consequence is reduced execution rewrite demand at fixed task performance:

$$\rho_1^{(\text{ATLB})}(\varepsilon_0) < \rho_1^{(\text{MEA})}(\varepsilon_0), \quad (8)$$

and therefore a reduced erased-bit rate for state registration:

$$\dot{I}_{\text{erase}}^{(\text{ATLB})} < \dot{I}_{\text{erase}}^{(\text{MEA})}. \quad (9)$$

THERMODYNAMIC BOOKKEEPING AND TESTABLE SCALING

Baseline (conventional MEA)

For a fixed benchmark task, decompose a task-normalized energy cost as

$$E_{\text{task}} = E_{\text{rev}} + Q_{\text{diss}}. \quad (10)$$

Repeated updates impose

$$Q_{\text{diss}}^{(\text{MEA})} \geq k_B T \ln 2 \Delta I_{\text{erase}}^{(\text{MEA})}. \quad (11)$$

Dense stimulation increases both ΔI_{erase} and effective friction Γ_L by forcing more frequent rewrites under limited relaxation.

ATLB prediction: monotonic dissipation reduction with Π_{cog}

A minimal, falsifiable scaling ansatz consistent with the PCAM motivation is

$$\Delta I_{\text{erase}}(\varepsilon_0) \approx \Delta I_0(\varepsilon_0) \left(\frac{\Pi_0}{\Pi_{\text{cog}}} \right)^\alpha, \quad \alpha > 0, \quad (12)$$

implying

$$Q_{\text{diss}}^{(\text{ATLB})} \geq k_B T \ln 2 \Delta I_0(\varepsilon_0) \left(\frac{\Pi_0}{\Pi_{\text{cog}}} \right)^\alpha. \quad (13)$$

Primary falsifiable content: the *sign and monotonicity*. At fixed ε_0 and controlled T , increasing Π_{cog} must *decrease* dissipation. The exponent α is treated as a measurable phenomenological parameter rather than a tuned constant.

Pre-registration rule (no post-hoc tuning)

To avoid post-hoc fitting, the following are declared *a priori* as part of the experimental protocol:

1. the benchmark task and performance metric ε (including target ε_0),

2. the estimators used to compute λ_2 and \mathcal{R} (including windows, bands, and preprocessing),
3. the dissipation metric (calorimetry or thermometry; calibration and thermal transfer model),
4. the statistical test: (i) monotonic trend test of Q_{diss} vs. Π_{cog} with fixed covariates; (ii) optional estimation of α on a held-out subset, without revising the estimator pipeline.

A failure of monotonicity under these rules is treated as a kill-test failure, not as an invitation to modify definitions.

FALSIFIABILITY AND KILL-TESTS

KT1: dissipation–topology scaling

Protocol: Run a fixed benchmark task under two conditions: (A) conventional MEA with passive bath; (B) ATLB with active boundary-conditioned ionic/EM patterns that measurably increase λ_2 and/or \mathcal{R} while keeping gross stimulation amplitude and bath temperature setpoints comparable. Measure simultaneously: (i) heat dissipation, (ii) functional connectivity to compute λ_2 , and (iii) coherence \mathcal{R} .

Prediction: At fixed performance ε_0 and controlled T , dissipation decreases monotonically with Π_{cog} . A null result—no decreasing trend or a sign inversion—kills the ATLB hypothesis.

KT2: coherence preservation under high update demand

Protocol: Increase update demand (task bandwidth/difficulty) under a fixed average power constraint and compare MEA vs ATLB. Measure: (i) critical update rate at which \mathcal{R} falls below a pre-defined threshold, and (ii) performance degradation curves.

Prediction: The ATLB increases the critical update rate before coherence collapse and shifts degradation toward latency-limited behavior rather than abrupt fragmentation. If the ATLB does not extend coherence-preserving regimes (or worsens them) while consuming comparable or lower power, the architectural principle is falsified.

CONCLUSION

We proposed a falsifiable architectural upgrade for wetware computing derived from MetaTime Open-EFT bookkeeping: an Active Topological Latency Bath that treats the liquid substrate as a programmable latency coprocessor rather than passive life support. The ATLB

replaces volumetric electrode scaling with topological control: boundary-conditioned ionic flows and holonomy-like phase biases are used to increase the measurable integrity functional $\Pi_{\text{cog}} = \lambda_2 \mathcal{R}$, thereby reducing execution rewrite demand and the associated Landauer heat burden at fixed task performance. The proposal stands or falls on two empirical kill-tests linking heat dissipation to Π_{cog} and testing coherence preservation under high update demand under pre-registered analysis rules.

[1] B. J. Kagan *et al.*, “In vitro neurons learn and exhibit sentience when embodied in a simulated game-world,” *Neuron* **110**, 3952–3969.e8 (2022). doi:10.1016/j.neuron.2022.09.001

[2] F. D. Jordan *et al.*, “Open and remotely accessible Neuroplatform for research in wetware computing,” *Frontiers in Artificial Intelligence* **7**, 1376042 (2024). doi:10.3389/frai.2024.1376042

[3] R. Landauer, “Irreversibility and Heat Generation in the Computing Process,” *IBM J. Res. Dev.* **5**, 183–191 (1961).

[4] C. H. Bennett, “The Thermodynamics of Computation—a Review,” *Int. J. Theor. Phys.* **21**, 905–940 (1982).

[5] Y. Aharonov and D. Bohm, “Significance of Electromagnetic Potentials in the Quantum Theory,” *Phys. Rev.* **115**, 485–491 (1959).

[6] M. E. J. Newman, *Networks: An Introduction* (Oxford University Press, 2010).

[7] D. S. Bassett and O. Sporns, “Network neuroscience,” *Nat. Neurosci.* **20**, 353–364 (2017).

[8] Y. Kuramoto, *Chemical Oscillations, Waves, and Turbulence* (Springer, 1984).