

Hebbian Physics Networks: A self-organizing computational architecture based on local physical laws

Gunjan Auti^{✉*} and Hirofumi Daiguji[✉]

Department of Mechanical Engineering, *The University of Tokyo*, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-8656, Japan

Gouhei Tanaka[✉]

Program of Computational Intelligence, Graduate School of Engineering, *Nagoya Institute of Technology*, Nagoya 466-8555 and International Research Center for Neurointelligence, Institutes for Advanced Study, *The University of Tokyo*, Tokyo 113-0033, Japan



(Received 7 July 2025; accepted 2 March 2026; published 23 March 2026)

Physical transport processes organize through local interactions that redistribute imbalance while preserving conservation. Classical solvers enforce this organization by applying fixed discrete operators on rigid grids. We introduce the Hebbian Physics Network (HPN), a computational framework that replaces this rigid scaffolding with a plastic transport geometry. An HPN is a coupled dynamical system of physical states on nodes and constitutive weights on edges in a graph. Residuals, local violations of continuity, momentum balance, or energy conservation, act as thermodynamic forces that drive the joint evolution of both the state and the operator (i.e., the adaptive weights). The weights adapt through a residual-modulated anti-Hebbian rule, which we prove constitutes a strictly local gradient descent on the residual energy. This mechanism ensures thermodynamic consistency: Near equilibrium, the realized transport operator converges to a symmetric, positive-definite form, reproducing Onsager's reciprocal relations without explicit enforcement. Far from equilibrium, the system self-organizes into transport topologies within a thermodynamically admissible class that redistribute imbalance through adaptive local geometry. Unlike optimization-based approaches that impose physics through global loss functions, HPNs embed conservation intrinsically: Transport is restored locally by the evolving operator itself, without a global Poisson solve or backpropagated objective. We demonstrate the framework on scalar diffusion and incompressible lid-driven cavity flow, showing that physically consistent transport geometries and flow structures emerge from random initial conditions solely through residual-driven local adaptation. HPNs thus reframe computation not as the solution of a fixed equation, but as a thermodynamic relaxation process in which the constitutive geometry and physical state coevolve.

DOI: [10.1103/tzkg-jqj4](https://doi.org/10.1103/tzkg-jqj4)

I. INTRODUCTION

The central abstraction of classical computational physics is the separation of state and structure. Classical continuum modeling assumes that the state $U(x, t)$ evolves under a prescribed local evolution operator,

$$\frac{\partial U}{\partial t} = F(U, \nabla U, \nabla^2 U, \dots; x, t, \theta), \quad (1)$$

where F (including constitutive closures θ) is fixed *a priori* and the computation advances the state U forward in an externally defined time variable. In a conventional simulation, this assumption is realized discretely. The physical state, such as velocity, temperature, or concentration, evolves in time, while the geometric structure governing its transport is fixed

in advance. Whether defined by a finite difference stencil, a finite element stiffness matrix, or a connection graph, the discretized constitutive operator is treated as a rigid scaffolding. Conservation is enforced by repeatedly applying this fixed operator under a globally synchronized update, forcing the system to conform to a prescribed geometric constraint [1–3].

Classical methods remain highly successful far from equilibrium, but they achieve this by embedding nonequilibrium effects within fixed or parametrized operators. In turbulent flows, effective momentum transport is represented through state-dependent eddy viscosities defined within prescribed closure frameworks [4,5]. In reacting systems, transport and reaction pathways reorganize through state-dependent rates, while the underlying interaction structure remains externally specified [6]. Similarly, in electrochemical and electrostatic systems, conductivities and screening lengths vary with local charge and potential, yet are incorporated through constitutive relations rather than through adaptation of the transport geometry itself [7]. In all such cases, nonequilibrium behavior is accommodated by enriching or parametrizing the function F , while retaining the deeper assumption that the transport structure itself is fixed and that the system's dynamics remain closed in the state variables alone.

*Contact author: gunjanauti@thml.t.u-tokyo.ac.jp

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

This separation between evolving state and fixed structure also underlies a range of adaptive and local computational approaches. Energy-based and relaxation-based models interpret computation as descent on an internal Lyapunov functional, but typically operate on a fixed interaction topology [8–10]. Pseudotime and artificial-compressibility methods introduce internal relaxation clocks to enforce constraints locally, yet rely on prescribed operators [11,12]. Lattice-based [13,14] and cellular-automaton [15] approaches generate macroscopic transport from local update rules, though these rules are fixed *a priori*. Adaptive discretization strategies, such as mesh refinement or coefficient tuning, modify numerical resolution or parameters while leaving the constitutive topology unchanged. Across these approaches, adaptation occurs within a *fixed* transport structure.

From the perspective of nonequilibrium thermodynamics, however, this fixed-structure assumption is not fundamental. Onsager’s reciprocal theory identifies the symmetric, positive operator linking thermodynamic forces to fluxes as the local geometric structure governing dissipation [16,17]. Prigogine’s extension to far-from-equilibrium systems emphasizes that conservation is maintained through continual local adjustment of these force-flux couplings, rather than through global coordination [18,19]. In this view, the constitutive operator is not a static background object but a dynamic field that co-evolves with the state variables as thermodynamic stress is redistributed.

A. Description of HPN

Building on this thermodynamic picture, we introduce the Hebbian Physics Network (HPN), a local relaxation solver based on the idea that physical evolution proceeds through the local restoration of broken conservation laws. Any external disturbance—such as forcing, boundary motion, or initial heterogeneity—locally disrupts the balance of conserved quantities, even though global conservation remains intact. Nonequilibrium systems respond to such imbalances by generating fluxes from their immediate surroundings, redistributing mass, momentum, or charge until local balance is restored, a mechanism underlying relaxation and pattern formation in driven systems [20]. HPN departs from this paradigm by abandoning the assumption that the dynamics are governed by a fixed evolution operator: Here, the effective evolution law is not prescribed, but constructed dynamically through local adaptation of the transport structure itself.

Crucially, the specific dynamical formulation by which this redistribution occurs is not assumed *a priori*. Instead, the HPN treats the mechanism of transport itself as an adaptive degree of freedom: Transport is represented as a network of local pathways whose strengths reorganize in response to imbalance. When a local violation of conservation arises, nearby pathways adjust to facilitate redistribution, and the state relaxes along these evolving connections. In this formulation, physically consistent behavior emerges not by explicitly time-integrating a fixed evolution equation of the form $\partial_t x = f(x)$ [Eq. (1)], but by allowing the transport structure itself to adapt until local conservation is restored. The effective dynamics of the system are therefore encoded in




FIG. 1. Schematic of the graph $\mathcal{G}$: each node i interacts with its neighbors $j \in \mathcal{N}(i)$ through directed couplings W_{ij} ; reciprocal edges $j \rightarrow i$ carry distinct couplings W_{ji} . U_i and U_j are the dynamic state variables at nodes i and j , respectively.

the evolving transport geometry, rather than imposed through a prescribed governing equation.

The network is defined on a discrete graph whose nodes store dynamic state variables (U) and whose directed edges represent the pathways along which flux can flow. Each directed edge $j \rightarrow i$ carries a coupling weight W_{ij} that determines how the state at node j contributes to the flux entering node i (see Fig. 1). These weights collectively form the local geometric structure of the graph: They specify which transport pathways are available and how strongly they interact.

To drive the adaptation of this geometry, we identify the local imbalance measured internally by the network. For each node i , we compute a residual using discrete Gauss divergence theorem:

$$R_i = \sum_j J_{ij} - \sum_j J_{ji},$$

$$J_{ij} = W_{ij} S_{ij}, \quad (2)$$

where S_{ij} is the local signal from node j to node i . The choice of S_{ij} is dictated by the constitutive relation of the target physics. Residual R_i is the physical accumulation of the conserved quantity (e.g., mass or momentum) at node i as shown in Eq. (2). This accumulation creates a local stress that acts upon the incident edges. It quantifies a local deviation whose relaxation requires modifying the force-flux mapping. Therefore, serving as the natural thermodynamic signal driving the dynamics of the constitutive operator. The nodal state U_i represents an intensive density (detailed definition in Sec. II A) stored in a node of finite capacity Ω_i . Accordingly, the state evolves by relaxing this accumulation at a rate inversely proportional to the local capacity. The state update is therefore defined as

$$\Delta U_i = \frac{\eta_u}{\Omega_i} \left(\sum_{j \in \mathcal{N}(i)} (J_{ij} - J_{ji}) \right) = \frac{\eta_u}{\Omega_i} R_i, \quad (3)$$

where η_u is the kinetic relaxation rate defined by a systemic stability constraint such as a Courant-Friedrichs-Lewy (CFL)-like limit.

The system adheres to a principle of constitutive relaxation: When a conserved quantity accumulates at a node (i.e., residual R_i), the incident pathways “yield” by augmenting their effective conductance to alleviate local stress, whereas, pathways opposing this relaxation are suppressed. This reshaping of the transport geometry ensures that the network converges toward a configuration where local conservation is satisfied intrinsically. This mechanism is formalized as a residual-modulated anti-Hebbian rule in form [21–23]

$$\Delta W_{ij} = -\eta_w f(R_i) S_{ij} - \lambda W_{ij}, \quad (4)$$

where the update is governed by (1) a signal basis S_{ij} , which defines the local transport context (presynaptic activity), (2) a modulatory residual $f(R_i)$ representing the local violation of balance, and (3) a systemic structural mobility η_w .

Importantly, this defines a coupled dynamical process in which the transport geometry evolves continuously in response to local imbalance, while global conservation is satisfied algebraically at all times (i.e., $\sum_i R_i = 0$). The directed couplings W_{ij} adapt their transmission capacity to relieve local accumulation, reshaping the effective transport pathways wherever imbalance persists. This process does not presuppose convergence to equilibrium: If external driving, boundary forcing, or sustained heterogeneity continually generate local stress, the network will continue to reorganize its transport structure. As a result, the HPN admits equilibria, nonequilibrium steady states, and time-dependent regimes within a single unified formulation.

Crucially, the resulting evolution is not an arbitrary search in configuration space. Because imbalance is measured through conservation residuals and relaxation occurs through strictly local updates, each step of the dynamics redistributes conserved quantities while adjusting the constitutive geometry that mediates that redistribution. In Sec. II, we show that these updates correspond to a local descent on a residual energy functional and recover reciprocal transport structure in the near-equilibrium limit. The trajectory generated by the HPN is therefore constrained by conservation and local dissipation, rather than being imposed by an externally prescribed time-marching scheme. In this sense, the object computed by an HPN is the relaxation path defined by these constraints, not merely a terminal steady state.

The remainder of this paper develops the formal structure of the HPN and demonstrates its operation on canonical transport problems. Section II defines the local fluxes and establishes that the Hebbian update constitutes a local gradient descent on the residual energy. Section III A examines scalar transport, illustrating how a consistent constitutive operator emerges solely through residual-driven feedback. Section III B extends the framework to incompressible lid-driven cavity flow, where pressure enforces the divergence constraint while the momentum-exchange weights adapt to maintain the correct local structure. Finally, in Sec. IV we discuss how physically consistent transport can arise from local constraint enforcement without a fixed operator. The HPN does not impose a constitutive geometry in advance; it

learns one by continually adjusting the force-flux mapping so that imbalance is eliminated *in situ*.

II. THERMODYNAMIC LEARNING VIA LOCAL RESIDUALS

Having established the qualitative framework of a *plastic* transport geometry, we now formalize the mechanics of adaptation. In this section, we define the Hebbian Physics Network as a coupled dynamical system (U, W) and analyze the local relaxation principles governing its evolution. We show that, for a broad class of residual-driven update rules, the adaptation of the transport weights constitutes a strictly local gradient descent on a well-defined residual energy, ensuring thermodynamic admissibility of the constitutive dynamics.

Crucially, the results in this section characterize the *structural* properties of the HPN update rules rather than guaranteeing convergence to a particular physical solution in all settings. Under near-equilibrium conditions and appropriate separation between state relaxation and structural adaptation, the learned transport operator approaches a symmetric, linear response form consistent with Onsager’s reciprocal relations. Away from this limit, the same local adaptation mechanism governs the coevolution of state and transport geometry without presupposing global equilibrium or a fixed constitutive law.

A. Formal definitions of state U and signal S

To map a physical system onto the HPN substrate, we must identify the fundamental degrees of freedom that characterize the local state and the nature of the interactions between them. The state of node i , denoted by U_i , represents the density of a quantity subject to a global conservation law. In the context of dynamical systems, U_i is the primary degree of freedom for which a local conservation law, such as the conservation of mass, momentum, or energy, is defined. For example, in a dissipative process such as diffusion, U_i represents the local concentration c_i . In fluid dynamics, U_i represents the momentum density $\rho \mathbf{u}_i$, where ρ is the fluid density and $\mathbf{u}_i$ is the velocity vector.

Crucially, in the HPN framework, global conservation is a structural invariant of the graph architecture. Because the network only redistributes the state among nodes via directed edges, the total sum $\sum_i U_i$ remains algebraically constant throughout the evolution. What the HPN focuses on is the local breaking of conservation, quantified by the residual R_i , and the subsequent restoration of local symmetry through autonomous relaxation.

The interaction between nodes is mediated by a coupling signal S_{ij} , representing the physical information transmitted from node j to node i across a directed edge. In the HPN framework, we adopt a state-centered signal,

$$S_{ij} = U_j, \quad (5)$$

and encode the constitutive physics entirely in the adaptive operator W_{ij} . The operator acts on the signal to generate a directed flux described by Eq. (2), which governs the redistribution of the conserved quantity across the graph.

With directed edges $i \rightarrow j$ and $j \rightarrow i$ treated as distinct pathways, conservation is enforced through the nodal residual,

$$R_i = \sum_j (W_{ij} S_{ij} - S_{ji} W_{ji}), \quad (6)$$

which represents the net accumulation of the conserved quantity at node i . This form corresponds to a discrete continuity equation on a directed graph: The first term accounts for incoming flux, while the second represents outgoing transport proportional to the local state. Diffusive and advective behavior are not distinguished by the signal, but by the symmetry properties of the learned operator. In the reciprocal limit $W_{ij} \approx W_{ji}$, the residual reduces to a graph Laplacian and recovers diffusive transport; deviations from reciprocity encode advective bias and circulation.

Crucially, the HPN does not learn isolated physical coefficients. Instead, W_{ij} evolves into a composite transport operator that captures the effective geometry through which imbalance is redistributed. In complex systems such as incompressible flow, the residual R_i aggregates contributions from multiple physical mechanisms (e.g., advection, dissipation, and constraints). The adaptive operator reorganizes to resolve this total thermodynamic stress without requiring separate, explicitly prescribed operators for each term.

B. Problem statement and notation

We consider a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where each node $i \in \mathcal{V}$ carries the intensive state $U_i \in \mathbb{R}^d$ and each edge $(i, j) \in \mathcal{E}$ carries a coupling operator $W_{ij} \in \mathbb{R}^{d \times d}$ representing its geometric conductance. The edge set $\mathcal{E}$ defines the fixed connectivity of the system and is constructed from the physical embedding of the nodes by a local neighborhood criterion (e.g., a radius-based or k -nearest-neighbor rule), which enforces locality but does not prescribe a precise metric within the neighborhood. To bridge the topological structure with physical space, we introduce a diagonal nodal metric matrix $\mathbf{G} = \text{diag}(\Omega_1, \dots, \Omega_N)$. Here, $\Omega_i > 0$ denotes the local material capacity or nodal volume, ensuring the total conserved quantity of the substrate is defined by the weighted sum $\mathbf{1}^\top \mathbf{G} \mathbf{U}$.

The neighborhood structure is determined directly from the known spatial coordinates of the nodes. Specifically, two nodes i and j are connected by a directed edge if their coordinates satisfy a Chebyshev distance criterion,

$$\|\mathbf{x}_i - \mathbf{x}_j\|_\infty \leq k,$$

where k is a prescribed interaction radius. This construction yields a fixed, local interaction graph that enforces spatial locality while remaining agnostic to the precise transport metric within the neighborhood; all constitutive transport detail is subsequently encoded through the adaptive evolution of the edge operators W_{ij} .

The local dynamics are driven by the interaction between the operator and the signal S_{ij} . This interaction generates a directed flux $J_{ij} = W_{ij} S_{ij}$, the accumulation of which defines the nodal residual as defined by Eq. (2). Physically, R_i represents a local thermodynamic stress indicating a departure from equilibrium. Because the residual is constructed purely from internal flux redistribution, global conservation is a structural

invariant of the architecture, satisfying $\sum_i R_i = 0$ regardless of the operator configuration.

To characterize the collective relaxation of the substrate, we define the global state $\mathbf{U} \in \mathbb{R}^{Nd \times 1}$, the global residual $\mathbf{R} \in \mathbb{R}^{Nd \times 1}$, and the global signal $\mathbf{S} \in \mathbb{R}^{Md \times 1}$ (where $M = |\mathcal{E}|$). While $\mathbf{U}$ represents the spatial distribution of the intensive density (e.g., concentration or momentum), $\mathbf{R}$ quantifies the instantaneous flux imbalance. By treating the residual as a thermodynamic force driving a continuous process of symmetry restoration, the HPN models the system's trajectory as the simultaneous coevolution of the state $\mathbf{U}$ and the underlying transport geometry $\mathbf{W}$.

C. Fundamental invariants and physical consistency

We identify the fundamental structural properties that govern the evolution of a Hebbian Physics Network. These properties distinguish between invariants enforced exactly by the graph architecture and thermodynamic consistency conditions that characterize the class of physical systems the framework is designed to model.

(1) **Global Conservation as an Exact Structural Invariant.** The HPN architecture is constructed such that, for any configuration of the transport operator $\mathbf{W}$ and state $\mathbf{U}$, the sum of residuals vanishes identically:

$$\sum_i R_i = 0. \quad (7)$$

Because transport occurs exclusively through internal redistribution along directed edges, the total intensive state $\sum_i U_i$ is preserved algebraically. Global conservation is therefore an exact structural invariant of the network, rather than a property enforced asymptotically or approximated through iteration.

(2) **Coercivity Condition and Symmetry Restoration.** For a physical system to admit a stable attractor, the mapping between thermodynamic forces (the residuals) and the state must satisfy a monotonicity condition consistent with positive entropy production. We characterize this requirement through a local coercivity condition on the residual operator:

$$(\mathbf{R}(\mathbf{U}_a) - \mathbf{R}(\mathbf{U}_b))^\top (\mathbf{U}_a - \mathbf{U}_b) \geq c \|\mathbf{U}_a - \mathbf{U}_b\|_{\mathbf{G}}^2, \quad (8)$$

for some $c > 0$, where $\mathbf{G}$ is the nodal metric. In this sense, the weights $\mathbf{W}$ function as a dynamic, metriclike structure that locally restores the dissipative symmetry required by the Second Law.

(2) defines a condition of thermodynamic admissibility. The class of systems considered in this work is restricted to those whose residual mapping satisfies this condition; systems that violate (2) fall outside the scope of the HPN framework and cannot be consistently modeled within it. Accordingly, (2) is treated as an axiom throughout. Sections IID and IIE show that bounded promotion, positive decay, and adiabatic separation ensure that the adaptive geometry realizes a locally monotone residual operator compatible with this thermodynamic admissibility.

D. Constitutive relaxation and the Onsager limit

The HPN manifests physical dynamics through the simultaneous relaxation of the nodal state $\mathbf{U}$ and the transport

geometry $\mathbf{W}$. This coevolution is governed by two distinct internal timescales that characterize the substrate's response to thermodynamic stress.

1. State relaxation and the kinetic limit

The nodal state U_i evolves to dissipate local accumulations at a rate dictated by the substrate's information propagation speed. Utilizing the nodal metric Ω_i defined in Sec. II B, the update is given by

$$\Delta U_i = \frac{\eta_u}{\Omega_i} \left(\sum_{j \in \mathcal{N}(i)} (J_{ij} - J_{ji}) \right). \quad (9)$$

Here, η_u represents the kinetic relaxation rate. To ensure numerical stability and physical causality, η_u is strictly bounded by the CFL limit, ensuring that transport does not exceed the characteristic grid velocity of the discrete manifold.

2. Structural adaptation of the transport geometry

The local transport geometry W_{ij} evolves through a superposition of three mechanisms that reshape the pathway conductance in response to the residual R_i :

$$\Delta W_{ij} = - \underbrace{\eta_w f(R_i) S_{ij}}_{\text{Plastic yield}} + \underbrace{\varepsilon g(|R_i|) S_{ij}}_{\text{Promotion}} - \underbrace{\lambda W_{ij}}_{\text{Decay}}. \quad (10)$$

(1) The *Plastic yield* term drives the structural adaptation of the substrate. If a specific pathway correlates with local accumulation, this term reduces its effective geometric conductance; conversely, pathways that aid in relieving stress are reinforced.

(2) The *Promotion* term introduces a small baseline permeability that prevents premature elimination of transport pathways and maintains numerical and structural regularity of the operator near equilibrium. This term is auxiliary to dissipation and does not alter the direction of relaxation when residuals are significant.

(3) The *Decay* term imposes a physical bound on the coupling strength, preventing the unbounded growth of transport coefficients and ensuring the weights remain within a physically admissible range.

3. Gradient descent interpretation and trajectory invariance

To formalize the stability of the HPN, we define the global residual energy $\mathcal{L}$ as the squared norm of the residuals under the inverse nodal metric:

$$\mathcal{L} = \frac{1}{2} \mathbf{R}^\top \mathbf{G}^{-1} \mathbf{R} = \frac{1}{2} \sum_i \frac{1}{\Omega_i} R_i^2. \quad (11)$$

This functional represents the total “unresolved stress” within the substrate. The gradient of this energy with respect to a specific local weight W_{ij} is given by

$$\frac{\partial \mathcal{L}}{\partial W_{ij}} = \frac{R_i}{\Omega_i} \times \frac{\partial R_i}{\partial W_{ij}} = \frac{R_i}{\Omega_i} S_{ij}. \quad (12)$$

Comparing this result to the Plastic yield term in the weight update rule [Eq. (18)], we find that the anti-Hebbian update constitutes local gradient descent on the residual energy landscape. This interpretation guarantees that, irrespective of

the instantaneous state evolution, the constitutive adaptation driven by plastic yield acts to reduce local thermodynamic stress.

A key consequence of this gradient structure is that the direction of the structural adaptation is determined solely by the local residual density and signal. While the mobility rate η_w dictates the “computational velocity” along the path, it does not alter the path itself.

Lemma 1 (Invariance of the relaxation trajectory). Assume smoothness of the residual energy $\mathcal{L}$, bounded promotion and decay, and adiabatic separation between state relaxation and structural adaptation. Then, for any structural mobility η_w within the stable bandwidth relative to the kinetic limit η_u , the HPN traverses a unique, physically defined trajectory in the $(\mathbf{U}, \mathbf{W})$ configuration space. The iteration index n functions as an internal relaxation time, with η_w acting solely as a reparametrization of progress along this trajectory.

Proof. To prove that the physical trajectory of the substrate is independent of the structural mobility η_w , we express the full weight update [Eq. (18)] as a dynamic vector field. We define the Structural Mobility η_w as the common prefactor for the substrate's adaptive response, such that the promotion (ε) and decay (λ) are scaled relative to this mobility: $\varepsilon = \eta_w \hat{\varepsilon}$ and $\lambda = \eta_w \hat{\lambda}$. The continuous-limit evolution of the weights $\mathbf{W}$ with respect to the iteration index n is

$$\frac{d\mathbf{W}}{dn} = \eta_w [-\nabla_{\mathbf{W}} \mathcal{L} + \hat{\varepsilon} \mathbf{\Gamma}_R(\mathbf{R}) - \hat{\lambda} \mathbf{W}], \quad (13)$$

where $\nabla_{\mathbf{W}} \mathcal{L}$ is the gradient of the residual energy and $\mathbf{\Gamma}_R(\mathbf{R})$ represents the global promotion field. The unit tangent vector $\hat{\mathbf{v}}$ defining the direction of the trajectory in the manifold is

$$\begin{aligned} \hat{\mathbf{v}} &= \frac{d\mathbf{W}/dn}{\|d\mathbf{W}/dn\|} = \frac{\eta_w [-\nabla_{\mathbf{W}} \mathcal{L} + \hat{\varepsilon} \mathbf{\Gamma}_R - \hat{\lambda} \mathbf{W}]}{\|\eta_w [-\nabla_{\mathbf{W}} \mathcal{L} + \hat{\varepsilon} \mathbf{\Gamma}_R - \hat{\lambda} \mathbf{W}]\|} \\ &= \frac{-\nabla_{\mathbf{W}} \mathcal{L} + \hat{\varepsilon} \mathbf{\Gamma}_R - \hat{\lambda} \mathbf{W}}{\|-\nabla_{\mathbf{W}} \mathcal{L} + \hat{\varepsilon} \mathbf{\Gamma}_R - \hat{\lambda} \mathbf{W}\|}. \end{aligned} \quad (14)$$

Since the scalar η_w cancels out, the direction of the update at any point in phase space is determined exclusively by the relative balance between the physical gradient, the structural promotion, and the homeostatic decay.

We introduce a dimensionless path parameter τ to represent the reparametrized progress of the substrate adaptation, such that $d\tau \propto \eta_w dn$, which implies

$$\frac{d\mathbf{W}}{d\tau} = -\nabla_{\mathbf{W}} \mathcal{L} + \hat{\varepsilon} \mathbf{\Gamma}_R - \hat{\lambda} \mathbf{W}. \quad (15)$$

This equation is entirely independent of η_w . It demonstrates that η_w serves only as a reparametrization of the path length. Within the stable, adiabatic regime described above, the HPN traverses a unique relaxation trajectory that is invariant under rescaling of η_w . Within this regime, the emergent physics reflects a structural property of the system's energy landscape and its relative relaxation constants ($\hat{\varepsilon}, \hat{\lambda}$) rather than an artifact of the computational update rate. They physically characterize the material response of the simulated substrate, defining the balance between exploratory noise and structural stability. ■

For the HPN to traverse a unique physical path independent of η_w , the structural mobility must satisfy the following bounds:

(1) The Upper Bound: the Adiabatic Stability Limit. The structural mobility η_w is upper bounded by the requirement that the geometry does not outpace the nodal relaxation. Mathematically, this is constrained by the curvature of the residual energy landscape $\mathcal{L}$. To maintain stability, the “step size” in weight space must not exceed the local characteristic scale of the gradient:

$$\eta_w < \frac{2}{\Lambda_{\max}(\mathbf{H}_W)}, \quad (16)$$

where $\Lambda_{\max}(\mathbf{H}_W)$ is the largest eigenvalue of the Hessian of the residual energy with respect to the weights. Beyond this bound, the system enters an unstable regime where the updates ΔW overshoot the equilibrium, leading to unphysical oscillations or divergence. Physically, this represents the Structural Yield Limit of the substrate—if the material is forced to adapt faster than its internal state can settle, the physical continuity of the trajectory is broken.

(2) The Lower Bound: the Resolution Limit. Conversely, η_w is lower bounded by the practical requirement of numerical resolution. While the path remains theoretically invariant as $\eta_w \rightarrow 0$, an excessively small mobility leads to a “stable” substrate where structural adaptation is negligible within the simulation’s time frame. For a unique path to be observable, η_w must be large enough to allow the weights to escape the “noise floor” created by the promotion term ε :

$$\eta_w > \frac{\Omega \varepsilon e^{-|\frac{R}{\Omega}|}}{|R|}. \quad (17)$$

This ensures that the Plastic yield (driven by the residual) remains the dominant force over the Structural Promotion (driven by subscale fluctuations).

4. Canonical choice and the linear limit

While the HPN framework admits any odd, monotone modulating function $f(\cdot)$ and non-negative, bounded function $g(\cdot)$, we adopt the linear response $f(R_i) = R_i/\Omega_i$ and the exponential gate $g(|R_i|) = e^{-|R_i|/\Omega_i}$ for the remainder of this work. This yields the canonical weight update

$$\Delta W_{ij} = \underbrace{-\frac{\eta_w}{\Omega_i} R_i S_{ij}}_{\text{Plastic yield}} + \underbrace{\varepsilon e^{-|\frac{R_i}{\Omega_i}|} S_{ij}}_{\text{Structural promotion}} - \underbrace{\lambda W_{ij}}_{\text{Homeostatic decay}}. \quad (18)$$

The linear choice $f(R_i) = R_i/\Omega_i$ corresponds to local gradient descent on the residual energy landscape, $\mathcal{L}$, while the exponential gate $g(|R|)$ modulates auxiliary structural activity on the magnitude of thermodynamic stress. To ensure physical consistency and numerical stability, we impose the hierarchical constraint $\varepsilon \ll \lambda$ such that promotion acts as a small perturbation relative to homeostatic decay.

Under this parametrization, the promotion term prevents structural degeneration of the transport geometry near equilibrium. If the promotion term is set to zero ($\varepsilon = 0$), the weight update reduces to the plastic-yield and decay contributions. In regimes where residuals become small, the plastic term vanishes, leaving only the decay term, so that

$\Delta W_{ij} = -\lambda W_{ij}$. In this case, all couplings decay exponentially toward zero, causing the learned operator to collapse to a degenerate, low-conductance state. Such collapse suppresses transport capacity and slows subsequent adaptation under renewed forcing. The promotion term maintains a small baseline permeability when residuals are weak, preventing this decay-driven collapse while remaining subordinate to dissipative adaptation when residuals are large ($\varepsilon \ll \lambda$). The exponential gate ensures a smooth transition between exploration near balance and strictly dissipative restructuring under strong imbalance, without introducing additional thresholds or discontinuities.

This canonical choice thus separates exploratory structural variability from dissipative relaxation in a controlled manner. In the following subsection, we show that in the near-equilibrium limit—where residuals are small and structural adaptation is quasistatic—this formulation admits a linear-response description consistent with classical nonequilibrium thermodynamics.

5. The Onsager limit

To formally ground the HPN framework within classical nonequilibrium thermodynamics, we demonstrate that the adaptive substrate recovers linear, reciprocal transport laws in the near-equilibrium limit. In this regime, the discrete operator does not merely converge numerically, but autonomously organizes into a configuration consistent with Onsager’s reciprocal relations.

We consider the near-equilibrium regime in which the transport geometry is quasistatic such that $\Delta W_{ij} \approx 0$. Given sufficiently small residuals ($|R_i| \ll 1$), we apply the canonical modulators established in Sec. IID 4: the linear response $f(R) = R/\Omega$ and a baseline promotion $g(|R|) \approx 1$. Under these assumptions, the stationary weight W_{ij}^* is determined by the balance of structural yield, promotion, and homeostatic decay:

$$0 = -\eta_w \frac{R_i}{\Omega_i} S_{ij} + \varepsilon S_{ij} - \lambda W_{ij}^*. \quad (19)$$

In the near-equilibrium regime, we further assume quasistatic adaptation, such that residuals vary slowly compared to structural relaxation and may be treated as locally fixed when solving this condition. Solving for the stationary weight yields

$$W_{ij}^* = \frac{1}{\lambda} \left[\varepsilon S_{ij} - \eta_w \frac{R_i}{\Omega_i} S_{ij} \right]. \quad (20)$$

Substituting this stationary weight into the flux definition, $J_{ij} = W_{ij}^* S_{ij}$, provides the effective constitutive relation for the pathway:

$$J_{ij} = \underbrace{\frac{\varepsilon}{\lambda} S_{ij}^2}_{J_{\text{base}}} - \left(\frac{\eta_w}{\lambda} S_{ij}^2 \right) \frac{R_i}{\Omega_i}. \quad (21)$$

Here, J_{base} represents a small baseline flux arising from the structural promotion. Since the canonical hierarchy enforces $\varepsilon \ll \lambda$ as defined in Sec. IID 4, this term becomes negligible relative to the dissipative response. Consequently, the

constitutive relation simplifies to a pure linear response:

$$J_{ij} \approx -L_{ij}R_i, \quad L_{ij} = \frac{\eta_w}{\lambda\Omega_i}S_{ij}^2. \quad (22)$$

The coefficient L_{ij} thus plays a role of an effective local Onsager transport coefficient. In the near-equilibrium regime, neighboring nodes sample the same local thermodynamic state. Since, $S_{ij} = U_j$ is an intensive density and Ω_i is a slowly varying local capacity, we have $U_i = U_j + O(\nabla U)$ and $\Omega_i \approx \Omega_j + O(\nabla\Omega)$. Consequently,

$$S_{ij}^2/\Omega_i \approx S_{ji}^2/\Omega_j + O(\nabla U, \nabla\Omega)$$

so the effective transport coefficients become symmetric to leading order:

$$L_{ij} \approx L_{ji}. \quad (23)$$

Here, reciprocity refers specifically to edge-pair equality of the dissipative force-flux coupling between conjugate transport pathways. Any edge-pair inequality in the learned operator corresponds to circulatory transport, which is energetically neutral and does not contribute to entropy production, fully consistent with Onsager's original formulation. The recovery of this symmetry demonstrates that HPN adaptation is governed by a principle of functional reversibility: without explicit global enforcement, the adaptive substrate self-organizes reciprocal dissipative pathways in the near-equilibrium limit. Classical linear transport laws thus emerge naturally as the limiting description of residual-driven constitutive relaxation.

E. Global stability and structural exploration

To establish the dynamical integrity of the HPN, we examine the evolution of the global residual energy, $\mathcal{L} = \frac{1}{2}\mathbf{R}^\top\mathbf{G}^{-1}\mathbf{R}$, which quantifies unresolved thermodynamic stress. While classical linear systems often require a monotonic Lyapunov decay, the HPN admits a more complex, nonmonotonic energy landscape. Such behavior is naturally permitted—and, in many nonlinear systems necessary—for structural exploration; it enables the substrate to navigate nonlinear phase transitions where the transient accumulation of stress is a precursor to the discovery of dissipative pathways.

1. The dynamic landscape and coevolution

The total variation of the residual energy is governed by the coupled evolution of the nodal states and the transport geometry:

$$\frac{d\mathcal{L}}{dn} = \underbrace{(\nabla_{\mathbf{U}}\mathcal{L})^\top \frac{d\mathbf{U}}{dn}}_{\text{Kinetic dissipation}} + \underbrace{(\nabla_{\mathbf{W}}\mathcal{L})^\top \frac{d\mathbf{W}}{dn}}_{\text{Structural adaptation}}. \quad (24)$$

The first term represents the kinetic (state) dissipation. Substituting the state update

$$\Delta\mathbf{U} = -\eta_u\mathbf{G}^{-1}\mathbf{R},$$

and assuming that the coercivity condition (2) [see Sec. II C] holds locally for the current transport geometry, we obtain

$$\begin{aligned} (\nabla_{\mathbf{U}}\mathcal{L})^\top \frac{d\mathbf{U}}{dn} &= \mathbf{R}^\top\mathbf{G}^{-1}\left(\frac{\partial\mathbf{R}}{\partial\mathbf{U}}\right)(-\eta_u\mathbf{G}^{-1}\mathbf{R}) \\ &= -\eta_u\mathbf{R}^\top\mathbf{M}_{\text{sym}}\mathbf{R} \leq 0, \end{aligned} \quad (25)$$

where

$$\mathbf{M}_{\text{sym}} \equiv \mathbf{G}^{-1}\left(\frac{1}{2}\left[\frac{\partial\mathbf{R}}{\partial\mathbf{U}} + \left(\frac{\partial\mathbf{R}}{\partial\mathbf{U}}\right)^\top\right]\right)\mathbf{G}^{-1}$$

denotes the symmetric part of the weighted Jacobian. Under the coercivity/monotonicity condition (2) [see Sec. II C], $\mathbf{M}_{\text{sym}}$ is positive semidefinite, implying

$$\mathbf{R}^\top\mathbf{M}_{\text{sym}}\mathbf{R} \geq 0.$$

Consequently, for fixed transport geometry $\mathbf{W}$, state relaxation is dissipative with respect to the residual energy $\mathcal{L}$.

The second term introduces nonmonotonicity. Because $\mathbf{R}$ is a nonlinear function of both $\mathbf{U}$ and $\mathbf{W}$, the landscape itself shifts as the weights evolve. In regions of strong nonlinearity, such as the onset of turbulence or vortex formation, the structural adaptation may temporarily increase $\mathcal{L}$. Physically, this represents the system climbing a structural barrier in the configuration space, accumulating stress in order to break an inefficient symmetry and access new transport pathways.

For vector fields ($d > 1$), this coevolution allows $\mathbf{W}$ to decompose into symmetric and skew-symmetric components. While the symmetric part drives the dissipation of residual energy, the skew-symmetric part organizes rotational transport allowing the system to redistribute flux and satisfy local conservation without additional energetic cost to the residual landscape. This mechanism is demonstrated in the lid-driven cavity results presented in Sec. III B, where purely viscous dissipation becomes insufficient and convective pathways emerge through a structural bifurcation.

2. Global boundedness and Lagrange stability

Although the HPN does not enforce global Lyapunov monotonicity, the structural dynamics remain bounded when $\lambda > 0$ and the promotion term is bounded. The structural update may be written as

$$\frac{d\mathbf{W}}{dn} = -\eta_w\nabla_{\mathbf{W}}\mathcal{L} - \lambda\mathbf{W} + \varepsilon e^{-|\mathbf{R}|}\mathbf{S}. \quad (26)$$

To assess boundedness, we consider the evolution of the weight norm. Using the identity

$$\frac{d}{dn}\|\mathbf{W}\| = \frac{\mathbf{W}^\top}{\|\mathbf{W}\|} \frac{d\mathbf{W}}{dn},$$

and applying the Cauchy–Schwarz inequality together with $e^{-|\mathbf{R}|} \leq 1$, we obtain

$$\frac{d}{dn}\|\mathbf{W}\| \leq -\lambda\|\mathbf{W}\| + \eta_w\|\nabla_{\mathbf{W}}\mathcal{L}\| + \varepsilon\|\mathbf{S}\|. \quad (27)$$

In the regimes considered here, $\|\nabla_{\mathbf{W}}\mathcal{L}\|$ remains bounded whenever the state and residual remain bounded, and $\|\mathbf{S}\|$ is bounded since it is constructed from local state variables.

Hence, there exists a constant $C > 0$ such that

$$\frac{d}{dn} \|\mathbf{W}\| \leq -\lambda \|\mathbf{W}\| + C. \quad (28)$$

By standard comparison arguments for first-order differential inequalities, this implies

$$\limsup_{n \rightarrow \infty} \|\mathbf{W}_n\| \leq \frac{C}{\lambda}, \quad (29)$$

and therefore the trajectory remains confined to a compact region of weight space. The coupled $(\mathbf{U}, \mathbf{W})$ dynamics are thus Lagrange stable under bounded promotion and positive decay.

III. DEMONSTRATIVE CASES

To illustrate the scope of Hebbian Physics Networks, we present two representative systems: continuum diffusion and incompressible flow in a lid-driven cavity. In both cases, physically consistent structures emerge from random or perturbed initial states through local residual relaxation and anti-Hebbian weight adaptation—without explicitly solving the governing equations using the time as a global time marching problem.

Throughout the remainder of this work, variables and parameters are reported in dimensionless form to facilitate scale-independent analysis. Unless explicitly stated otherwise, all quantities have been normalized by their respective characteristic values as defined by the underlying conservation laws. Additionally, all simulations were conducted with η_w chosen within the adiabatic stability bandwidth described in Sec. IID.

A. Diffusion

We consider an unstructured graph whose connectivity is defined by a local neighborhood criterion: Each node i is connected to a node j if the Chebyshev distance between their coordinates satisfies

$$\|\mathbf{x}_i - \mathbf{x}_j\|_\infty \leq k$$

(see Sec. IIB). On this graph, we consider the capacity of all the nodes is equal such that $\Omega_i \equiv \Omega = 1$. The nodal state U_i is identified with the concentration c_i , consistent with the definition of state in Sec. IIA as the density of a conserved quantity, i.e., amount of substance in this case. The constitutive signal is taken to be the source concentration at the neighboring node, such that

$$S_{ij} = c_j, \quad (30)$$

as the state-centered signal defined in Sec. IIA.

All constitutive physics is encoded in the adaptive transport operator W_{ij} . Therefore, the directed flux along each edge is defined as

$$J_{ij} = DW_{ij}S_{ij}. \quad (31)$$

Here, the diffusivity D is treated as a material property that sets the overall timescale and magnitude of diffusive transport, while the adaptive operator W_{ij} encodes the effective transport geometry of the graph. Although W_{ij} adapts in response to residuals, it remains dimensionless and represents relative

pathway strengths rather than absolute transport rates. Keeping D explicit therefore separates material properties from geometric structure, preventing degeneracy in the learned operator and allowing the same transport geometry to be applied across different diffusive regimes. Absorbing D into W_{ij} would be mathematically equivalent but would obscure the distinction between learned geometry and intrinsic material response. The continuity residual at node i is given by the discrete divergence

$$R_{c,i} = D \sum_{j \in \mathcal{N}(i)} (W_{ij}S_{ij} - W_{ji}S_{ji}), \quad (32)$$

which represents the local accumulation of concentration due to imbalance between incoming and outgoing transport pathways.

In this case, transport between neighboring nodes i and j is represented by two directed couplings, W_{ij} and W_{ji} . We impose an *edge-pair equality* constraint,

$$W_{ij} = W_{ji}, \quad (33)$$

which restricts the admissible transport geometry to link-level balanced couplings. This is a constitutive choice for the diffusion example: Edge-pair unequal couplings generate divergence-free circulation and do not contribute to accumulation of a scalar conserved density. By enforcing edge-pair equality, all transport pathways contribute to local accumulation and relaxation. Under this restriction, the nodal residual reduces to a graph Laplacian acting on the concentration field, and diffusive transport emerges without explicitly encoding gradients in the signal.

Classical graph-based discretizations prescribe W_{ij} directly from the geometric embedding (typically $W_{ij} \propto \mathbf{r}_{ij} / \|\mathbf{r}_{ij}\|^2$) and treat it as fixed in time. Here, no geometric information is assumed *a priori*: The effective geometry is inferred through the learned weights W_{ij} themselves.

The node and weight updates follow the local anti-Hebbian rule

$$c_i^{n+1} \leftarrow c_i^n + \eta_c R_{c,i}^n, \quad (34)$$

$$W_{ij}^{n+1} \leftarrow W_{ij}^n + S_{ij}(-\eta_w R_{c,i}^n + \varepsilon_w e^{-|R_{c,i}^n|}) - \lambda W_{ij}^n, \quad (35)$$

where S_{ij} is the *local signal*, $R_{c,i}$ acts as a modulatory signal, and the decay term ensures long-term boundedness of the couplings. Superscripts, $(\cdot)^n$, denote the iteration index.

The parameters used here are

$$\eta_w = 0.01, \quad \varepsilon_w = 0.0001, \quad \eta_c^0 = 0.001,$$

$$\eta_c = \min\left(\frac{C|r_{ij}||S_{\max}|}{|J_{\max}|}, \eta_c^0\right), \quad C = 1.0, \quad \lambda = 0.01,$$

where $C|r_{ij}||S_{\max}|/|J_{\max}|$ is a CFL-like limit for diffusion transport. The concentration field is initialized to zero, and a localized perturbation is introduced at $n = 0$; all weights begin at unity.

Figure 2 summarizes the resulting dynamics. Panel (a) shows the evolution of the concentration field. The spreading front coevolves with the operator: The fluxes (white quivers) are generated by W , and those fluxes in turn reshape W through the anti-Hebbian update. This bidirectional coupling produces a finite-speed propagation of activity. The structure




FIG. 2. Residual-driven emergence of diffusive transport in an HPN. (a) Evolution of the concentration field on an unstructured graph. (b) Spatiotemporal trajectory of the emergent transport front obtained by tracking the y coordinates of nodes crossing a threshold concentration; red dashed lines indicate the iteration indices corresponding to the snapshots in panel (a). (c) Global conservation error $|\sum_i R_{c,i}|$, which remains at $\mathcal{O}(10^{-11})$ throughout the evolution, demonstrating exact global conservation. (d) Mean absolute local residual $\langle |R_c| \rangle$, which remains finite [$\sim \mathcal{O}(10^{-5})$] and reflects persistent local thermodynamic stress. (e) Lyapunov function $\mathcal{L}$ showing exponential monotonic decrease with iterations. (f) Root-mean-square change in weights, ΔW_{RMS} , indicating rapid early adaptation followed by asymptotic stabilization of the constitutive operator. (g) Mean weight magnitude $\langle W \rangle$ approaching a stable value as the learned transport geometry statistically becomes stationary. (h) Histograms of the weight distribution W_{ij} at selected iterations, illustrating pruning of inactive pathways and emergence of an effective transport geometry.

of the diffusion front in panel (b)—obtained by tracking the y coordinates of first activations—follows a clear parabolic envelope, $r \propto \sqrt{n}$. This observation allows the iteration index n to be calibrated against physical time *a posteriori*, rather than being imposed externally.

The global conservation error in panel (c) remains at machine precision throughout the evolution, demonstrating that conservation is invariant and not an emergent property. Relaxation instead occurs locally: The mean absolute residual in panel (d) shows an initial burst followed by rapid decay as the transport geometry adapts to redistribute local imbalances.

The Lyapunov function $\mathcal{L} = \frac{1}{2} \sum_i \|R_{c,i}\|^2$ in panel (e) decreases monotonically and then approaches a small, finite plateau. For unforced diffusion on a closed (periodic) domain, the physical equilibrium corresponds to a spatially uniform concentration and vanishing macroscopic flux. In the HPN implementation, however, the bounded promotion term and finite numerical resolution maintain a residual noise floor, so the nodal residuals need not vanish identically even when the concentration field is effectively uniform.

The long-time behavior therefore represents a stationary fluctuation regime around the equilibrium manifold rather than a thermodynamic nonequilibrium steady state. Global conservation is satisfied exactly [panel (c)], and the learned transport geometry W stabilizes in norm, while the small but nonzero ΔW reflects continued microscopic reorganization driven by bounded promotion rather than sustained macroscopic transport. The finite plateau of $\mathcal{L}$ thus reflects the noise-sustained residual level inherent to the adaptive dynamics, not the presence of persistent physical flux.

Panels (f) and (g) show the root-mean-square change in weights and the mean weight magnitude, respectively: W adapts substantially over the first ~ 4000 iterations and then asymptotically stabilizes, marking the emergence of a steady transport geometry. The histograms in panel (h) reveal the pruning behavior characteristic of the weight update rule: Weights that do not contribute to residual reduction shrink toward zero, producing a bistable distribution that encodes the effective transport geometry of the learned medium.

Although the algebraic form of the update resembles a graph discretization of diffusion, the mechanism of evolution is fundamentally distinct from explicitly time-integrating the parabolic partial differential equation (PDE) $\partial_t c = D \nabla^2 c$ under an external clock. No global time step is imposed, and no fixed Laplacian is ever constructed or inverted. Instead, both state c and transport geometry W evolve together: W determines the instantaneous fluxes; the fluxes produce the residual; and the residual drives the joint update of the state and the operator. Each iteration is therefore not a step of a numerical time-marching solver for $\partial_t c$, but a local relaxation event in which state and transport geometry coevolve. As W stabilizes, the emergent dynamics reproduce the diffusion law as a coarse-grained description, but the computation itself is the relaxation trajectory realized through adaptive local transport.

An explicit comparison of this method with the classical diffusion solution, carried out for a minimal two-node system, is provided in Sec. S1 A of the Supplemental Material [24].

B. Incompressible flow: lid-driven cavity

In this example, the residual operator is defined by the discrete Navier-Stokes force balance. The objective is not to re-derive the fluid equations from transport incompatibility alone, but to examine whether the HPN can adapt its constitutive transport geometry to realize the correct nonlinear dynamics under physically meaningful constraints.

Accordingly, we employ the HPN in a hybrid constraint-transport formulation. This separation is necessary because incompressibility and momentum balance are strict physical constraints that must be enforced with fidelity and do not arise in scalar diffusion problems. In this setting, the Navier-Stokes operator is used solely to evaluate the degree to which the current state violates force balance and divergence-free conditions. It does not prescribe the transport mechanism, nor is it integrated forward in time (i.e., no explicit $\partial_t \mathbf{u}$ time marching is performed). Instead, the HPN adapts the constitutive transport geometry W responsible for momentum redistribution, while the residual operator defines the physical constraints to be satisfied.

For incompressible flow, two conserved quantities are involved: mass and momentum. The fundamental state variables are therefore the mass density ρ and the momentum density $\rho \mathbf{u}$. In the incompressible limit, ρ is uniform and fixed, so all dynamical evolution occurs in the velocity field,

$$U_i \equiv \mathbf{u}_i = (u_i, v_i), \quad (36)$$

which fully characterizes the local momentum density up to a constant factor. Consistent with the state-centered formulation introduced earlier, the signal transmitted along a directed edge ($j \rightarrow i$) is the velocity at the source node,

$$S_{ij} \equiv \mathbf{u}_j. \quad (37)$$

Momentum is redistributed through an adaptive transport operator W_{ij} acting on this signal.

The governing equations are used solely to evaluate constraint residuals. In particular, the Navier-Stokes operator is employed to compute the local imbalance of momentum and the local violation of incompressibility. On an unstructured graph, spatial derivatives are approximated using discretized Green-Gauss operators constructed from the local neighborhood of each node. The resulting residuals are

$$\begin{aligned} \mathbf{R}_u &= \rho(\mathbf{u} \cdot \nabla) \mathbf{u} - \nabla p - \mu \nabla^2 \mathbf{u}, \\ R_c &= \rho \nabla \cdot \mathbf{u}, \end{aligned} \quad (38)$$

where p is the pressure and μ is the viscosity. Since the density ρ is constant, it may be absorbed into the residual normalization; it is retained here only to emphasize momentum balance and is absorbed hereafter for notational clarity.

In this formulation, pressure is not treated as an independently transported quantity. Instead, it arises as a response to violations of mass conservation. Local divergence of the velocity field signals a breakdown of incompressibility, and the pressure field acts as a Lagrange multiplier that corrects this violation by redistributing momentum. Pressure therefore appears only through the continuity residual and is updated as part of the residual relaxation process.

The relaxation dynamics are governed by an adaptive vector transport operator W_{ij} . Because the nodal state is two dimensional, the learned transport geometry at each edge is represented as a 2×2 operator acting on the velocity vector,

$$W_{ij} = \begin{pmatrix} W_{x,ij} & 0 \\ 0 & W_{y,ij} \end{pmatrix}, \quad (39)$$

where $W_{x,ij}$ and $W_{y,ij}$ encode transport of the x - and y -velocity components, respectively. This block-diagonal form preserves independent transport of each component while allowing directional bias through edge-pair inequality ($W_{ij} \neq W_{ji}$). As in the diffusion case, the weights define a local interaction graph of finite radius around each node and collectively encode the learned transport geometry.

Although the residuals are computed from the Navier-Stokes constraints, the adaptation mechanism itself is unchanged: Transport pathways that reduce local momentum imbalance are reinforced, while those that oppose relaxation are suppressed. The weights evolve according to a residual-modulated anti-Hebbian rule,

$$W_{ij}^{n+1} = (1 - \lambda)W_{ij}^n + \mathbf{u}_j^n \otimes (-\eta_w \mathbf{R}_{u,i}^n + \varepsilon_w g(-|\mathbf{R}_{u,i}^n|)), \quad (40)$$

where the modulation function $g(\cdot)$ acts componentwise so that promotion occurs independently along each velocity component.

For incompressible flow, the primary constraint is mass conservation. Local violations of incompressibility are quantified by the continuity residual $R_{c,i}$. Rather than being applied directly to the velocity field, this residual is enforced through a pressure correction: Pressure acts as a Lagrange multiplier that restores mass balance,

$$p_i^{n+1} = p_i^n + \eta_p R_{c,i}^n. \quad (41)$$

This update defines the admissible manifold on which momentum may be redistributed.

Momentum relaxation is then realized through adaptive transport along this pressure-corrected manifold. The velocity field evolves according to

$$\mathbf{u}_i^{n+1} = \mathbf{u}_i^n + \eta_u \sum_j W_{ji}^n \mathbf{u}_j^n, \quad (42)$$

which represents the incoming contribution of transported momentum from neighboring nodes. A fully conservative net-flux formulation would also include an outgoing contribution of the form $\eta_u \sum_j W_{ji}^n \mathbf{u}_i^n$, accounting for momentum transported away from node i . In the present hybrid approach, this term is not applied explicitly in the velocity update. Instead, its effect is captured indirectly: Momentum outflow that induces local divergence is measured by the continuity residual $R_{c,i}$ and corrected through the pressure update.

In this way, mass conservation and momentum transport are handled by complementary mechanisms. Pressure corrections enforce incompressibility as a hard constraint, while adaptive transport redistributes momentum within that constraint. The velocity update is thus not a standalone conservative step, but part of a coupled relaxation in which pressure and transport geometry jointly determine the admissible dynamics.

Figure 3 summarizes the relaxation dynamics of the HPN applied to the lid-driven cavity. Panel (a) shows streamline snapshots at selected iterations. Starting from an initially quiescent field, the flow organizes into the canonical recirculating structure as the learned transport operator redistributes momentum along directions indicated by local residual imbalance. The background pressure field (centered at zero) reflects the progressive establishment of shear-driven pressure gradients, while the quiver arrows show the emergence and sharpening of the primary vortex.

The oscillatory pressure features that appear near the top corners during early iterations do not indicate numerical instability. Rather, they are transient signatures of enforcing incompressibility through residual-driven local adjustments without a global elliptic projection. The resulting pressure fluctuations are consistent with pseudocompressible (artificial compressibility) relaxation mechanisms [11] and decay as the learned transport geometry stabilizes.

Panel (b) tracks the evolution of the constraint residuals. The continuity residual $\langle |R_c| \rangle$ and the momentum residual components $\langle |R_u| \rangle$ and $\langle |R_v| \rangle$ do not relax monotonically. Instead, they exhibit a pronounced transient surge followed by rapid decay. This surge marks a bifurcation in the relaxation dynamics of the coupled state-operator system (U, W) . As the network begins simultaneously enforcing incompressibility and momentum balance, the initial transport geometry becomes incompatible with the imposed constraints, leading to an accumulation of local stress. Once this stress exceeds a threshold, the system undergoes a bifurcation in which the initial relaxation pathway loses stability and a dynamical regime—characterized by transverse momentum redistribution—emerges. The Lyapunov function in panel (c), $\mathcal{L} = \frac{1}{2} \sum_i (R_{c,i}^2 + R_{u,i}^2 + R_{v,i}^2)$, exhibits the same structure. The sharp peak (red marker) identifies the bifurcation point at which the transport operator reorganizes to admit admissible pathways for residual relaxation. Following this bifurcation, residuals decay as the system converges toward a stable nonequilibrium steady state supported by the learned transport geometry.

Panels (d) and (e) display the evolution of the learned transport operator. The root-mean-square weight changes $\Delta W_{u,\text{RMS}}$ and $\Delta W_{v,\text{RMS}}$ show an initial phase of strong plasticity followed by asymptotic freezing as the operator adapts to the dominant shear and recirculation pattern. The mean weights $\langle W_u \rangle$ and $\langle W_v \rangle$ likewise settle into steady values as the learned transport geometry stabilizes.

Panel (f) shows histograms of the weight distributions $W_{ij}^{(u)}$ and $W_{ij}^{(v)}$ at selected iterations. The initially narrow distributions broaden during flow formation and develop heavy-tailed structure as transport pathways become selectively reinforced. As the steady vortex emerges, couplings that do not contribute to residual reduction, shrink toward zero, while pathways aligned with active momentum redistribution remain strengthened. This pruning and anisotropic reinforcement reflect the emergence of an effective constitutive transport geometry tailored to the cavity constraints.

A quantitative comparison of the converged steady state at $\text{Re} = 100$ with the standard benchmark data of Ghia *et al.* [25] is provided in Sec. S2 of the Supplemental Material




FIG. 3. Residual-driven flow formation in an HPN. (a) Streamline snapshots for the lid-driven cavity. From a nearly quiescent state, the network organizes the velocity into the characteristic recirculating pattern; pressure is shown as a background colormap, centered at zero. (b) Mean continuity, streamwise, and cross-stream residuals $\langle |R_c| \rangle$, $\langle |R_u| \rangle$, and $\langle |R_v| \rangle$, showing an early surge followed by rapid decay as imbalances are redistributed through the learned transport operator. (c) Lyapunov function $\mathcal{L} = \frac{1}{2} \sum_i (R_{c,i}^2 + R_{u,i}^2 + R_{v,i}^2)$, increases until the network discovers the vertical degree of freedom and bifurcates after reaching a maximum. (d) Root-mean-square weight changes $\Delta W_{u,RMS}$ and $\Delta W_{v,RMS}$, indicating an initial phase of strong plasticity followed by stabilization. (e) Mean weights $\langle W_u \rangle$ and $\langle W_v \rangle$ approaching stable values as the transport geometry settles. (f) Histograms of $W_{ij}^{(u)}$ and $W_{ij}^{(v)}$ at selected iterations, initially narrow, then broadening during adaptation and developing heavy-tailed structure associated with shear, before the flow settles into a nonequilibrium steady state.

[24], confirming agreement with the classical lid-driven cavity solution.

Overall, this example highlights that the object learned by the HPN is the local transport geometry, not the Navier-Stokes operator itself. Fixed discrete Navier-Stokes operators (used only for residual evaluation) supply the momentum and incompressibility constraints, while the evolving weights W determine how momentum is redistributed from one relaxation event to the next. Although the resulting flow field reproduces the characteristic lid-driven cavity structure, the trajectory is not a time integration of the PDE; rather, the computation is the equilibration process itself, mediated by a plastic transport operator that continuously reshapes its geometry to relieve constraint-induced stress.

IV. DISCUSSION AND OUTLOOK

The Hebbian Physics Network shares the objective of physics-informed machine learning (PIML): to embed physically meaningful structure directly into computational systems so that their evolution respects conservation principles and constitutive constraints. The distinction lies in the mode of enforcement. In most PIML approaches, physical structure is imposed top-down, typically through fixed governing equations or global loss functions evaluated over a prescribed operator. In HPN, in contrast, physical structure is adapted bottom-up: The system reshapes its local constitutive pathways directly in response to observed imbalance, without requiring global constraint enforcement or a fixed constitutive operator.

An instructive analogy is a two-dimensional grassy field exposed to airflow. One approach to reconstructing the flow is to assume the Navier-Stokes equations in advance and fit a velocity field consistent with the observed bending of the grass [26]. This is a top-down strategy: The governing equation is fixed, and observations are used to determine the state. An alternative is to infer how each grass blade responds locally to the surrounding motion using the observed blade deflections as boundary information [27]. No global flow equation is imposed as a synchronized evolution law. Instead, a local constitutive description is described everywhere, and the global flow field emerges as a consequence of these local responses.

HPN follows this second paradigm. It does not encode the Navier-Stokes equations as a fixed transport operator or time-integration rule; rather, it adapts the transport structure to be consistent with local observations, and the familiar continuum equations reappear only as an emergent, coarse-grained description of the resulting dynamics.

HPN provides a complementary formulation of transport dynamics, whose relaxed coarse-grained behavior is described by dynamic PDE solvers. In this view, PDEs provide macroscopic conservation laws, while dynamic PDE solvers enforce these laws through globally synchronized update rules. By requiring instantaneous consistency across the domain, many traditional PDE solvers necessarily coarse-grain the underlying nonequilibrium dynamics, masking the finite-rate local relaxation processes inherent in physical substrates [20,28,29].

This bottom-up perspective places HPN closer in form to energy-based computation, in which system evolution is governed by internally generated consistency signals rather than externally prescribed update rules [8,9,30]. The essential distinction is that, in HPN, these signals are physically interpretable as indicators of local nonequilibrium structure. In nonequilibrium thermodynamics, residuals—the local imbalances in a conserved quantity—signal that the available transport pathways are incompatible with the instantaneous state [16,18]. HPN treats these residuals as the drivers of constitutive adaptation: Transport pathways reorganize until local inconsistencies are reduced. In HPN, this constitutive adaptation is realized explicitly through the evolution of the transport weights W_{ij} that define how state is communicated between neighboring nodes.

Formally, we denote by $\mathcal{W} = \{W_{ij}\}$ the collection of transport weights, and by $\mathbf{W} \in \mathbb{R}^{N \times N}$ the sparse operator assembled from these weights. Rather than viewing $\mathbf{W}$ as a fixed discretization of a differential operator, it is more natural to interpret it as a discrete constitutive structure defined on the graph. In this interpretation, the weights W_{ij} play a role structurally analogous to gauge connections: They specify how state information is transported between neighboring nodes and how local reference frames are coupled, in the same sense that connections define parallel transport on a manifold [31,32].

From this perspective, the residual at a node reflects the failure of these local connections to be mutually compatible. This is directly analogous, at a discrete level, to the appearance of curvature in geometric theories, where incompatibility of parallel transport around infinitesimal loops signals a nonflat manifold [33,34]. The residual thus serves as an operational indicator of geometric inconsistency in the learned transport structure, rather than as a numerical error in state variables.

From this structural standpoint, the dynamics of HPN can be interpreted as a process of local geometry adjustment. When residuals are large, the constitutive connections are poorly aligned, and the effective manifold defined by $\mathbf{W}$ is strongly curved. As the system evolves, updates to $\mathcal{W}$ act to flatten this local geometry, reducing curvature and improving compatibility between neighboring states. This adaptive reshaping of transport structure bears a close conceptual relationship to geometric formulations of dynamics and natural-gradient flows, where evolution proceeds along directions determined by the local metric structure [35,36].

The precise mathematical correspondence between residuals, curvature, and gauge structure is not pursued here and will be developed in future work. The present discussion is intended to provide a structural and physical interpretation of adaptive transport in HPN, rather than a formal gauge-theoretic construction.

Taken together, these observations position the Hebbian Physics Network as a transparent physics-informed learning framework in which transport dynamics are organized through local resolution of imbalance rather than enforced by a globally prescribed evolution law. In this view, dissipation, constitutive structure, and symmetry restoration emerge jointly from residual-driven adaptation, with the learned

transport geometry encoding how a system redistributes stress as it relaxes. Classical partial differential equations then appear as macroscopic descriptions of the relaxation trajectories generated by the HPN, rather than as equations that are explicitly time-integrated within the algorithm. The HPN framework is therefore not intended to replace continuum theory, but to clarify how constitutive structure consistent with such descriptions can arise dynamically from local interactions—an outlook that becomes particularly relevant far from equilibrium, where effective transport laws are not known in advance and must be discovered through the dynamics themselves.

ACKNOWLEDGMENTS

The open access publication of this work was supported by JSPS KAKENHI Grant No. JP23K28154 and JST CREST Grant No. JPMJCR24R2.

G.A. gratefully acknowledges his partner, Ms. Shambhavi Chaturvedi (National Institute of Genetics, Japan), for insightful discussions, especially on the neuroscience aspect relevant to this work.

DATA AVAILABILITY

No data were created or analyzed in this study.

- [1] S. Patankar, *Numerical Heat Transfer and Fluid Flow* (CRC Press, New York, 2018).
- [2] J. H. Ferziger, M. Perić, and R. L. Street, *Computational Methods for Fluid Dynamics* (Springer, 2019).
- [3] R. J. LeVeque, *Finite Volume Methods for Hyperbolic Problems* (Cambridge University Press, 2002), Vol. 31.
- [4] J. Boussinesq, *Essai Sur la Théorie Des Eaux Courantes* (Imprimerie Nationale, 1877).
- [5] S. B. Pope, Turbulent flows, *Meas. Sci. Technol.* **12**, 2020 (2001).
- [6] S. R. De Groot and P. Mazur, *Non-Equilibrium Thermodynamics* (Courier Corporation, 2013).
- [7] M. Z. Bazant, M. S. Kilic, B. D. Storey, and A. Ajdari, Towards an understanding of induced-charge electrokinetics at large applied voltages in concentrated solutions, *Adv. Colloid Interface Sci.* **152**, 48 (2009).
- [8] J. J. Hopfield, Neural networks and physical systems with emergent collective computational abilities, *Proc. Natl. Acad. Sci. USA* **79**, 2554 (1982).
- [9] *Predicting Structured Data*, edited by G. BakIr, T. Hoffman, B. Scholkopf, S. Bernhard, J. Alexander, B. Taskar, and S. V. N. Vishwanathan (MIT press, Cambridge, Massachusetts, 2007).
- [10] K. Friston, The free-energy principle: A unified brain theory? *Nat. Rev. Neurosci.* **11**, 127 (2010).
- [11] A. J. Chorin, A numerical method for solving incompressible viscous flow problems, *J. Comput. Phys.* **135**, 118 (1997).
- [12] A. Jameson, Time dependent calculations using multigrid, with applications to unsteady flows past airfoils and wings, in *Proceedings of the 10th Computational Fluid Dynamics Conference* (Honolulu HI USA, American Institute of Aeronautics and Astronautics, 1991).
- [13] U. Frisch, B. Hasslacher, and Y. Pomeau, Lattice-gas automata for the Navier-Stokes equation, *Phys. Rev. Lett.* **56**, 1505 (1986).
- [14] G. G. McNamara and G. Zanetti, Use of the Boltzmann equation to simulate lattice-gas automata, in *Lattice Gas Methods for Partial Differential Equations* (CRC Press, 2019), pp. 289–296.
- [15] S. Wolfram, *Cellular Automata and Complexity Collected Papers*, Part one (CRC Press, Boca Raton, Florida, 2018).
- [16] L. Onsager, Reciprocal relations in irreversible processes. I., *Phys. Rev.* **37**, 405 (1931).
- [17] H. B. Callen, *Thermodynamics and an Introduction to Thermostatistics*, 2nd ed. (John Wiley & Sons, 1993).
- [18] I. Prigogine and G. Nicolis, *Self-Organization in Non-Equilibrium Systems* (Wiley, 1977).
- [19] I. Prigogine, Time, structure, and fluctuations, *Science* **201**, 777 (1978).
- [20] M. C. Cross and P. C. Hohenberg, Pattern formation outside of equilibrium, *Rev. Mod. Phys.* **65**, 851 (1993).
- [21] P. Földiák, Forming sparse representations by local anti-Hebbian learning, *Biol. Cybern.* **64**, 165 (1990).
- [22] N. Frémaux and W. Gerstner, Neuromodulated spike-timing-dependent plasticity, and theory of three-factor learning rules, *Front. Neural Circuits* **9**, 85 (2016).
- [23] T. Isomura and T. Toyozumi, Error-gated Hebbian rule: A local learning rule for principal and independent component analysis, *Sci. Rep.* **8**, 1835 (2018).
- [24] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/tzgz-jqj4> for additional information on the demonstrative cases, which includes Ref. [25] as benchmarking data.
- [25] U. Ghia, K. N. Ghia, and C. Shin, High-Re solutions for incompressible flow using the Navier-Stokes equations and a multigrid method, *J. Comput. Phys.* **48**, 387 (1982).
- [26] M. Raissi, A. Yazdani, and G. E. Karniadakis, Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations, *Science* **367**, 1026 (2020).
- [27] J. L. Cardona, K. L. Bouman, and J. O. Dabiri, Wind speed inference from environmental flow–structure interactions, *Flow* **1**, E4 (2021).
- [28] G. Lebon, D. Jou, and J. Casas-Vázquez, *Understanding Non-Equilibrium Thermodynamics: Foundations, Applications, Frontiers* (Springer, Berlin, Heidelberg, 2008).
- [29] W. G. Hoover, *Computational Statistical Mechanics* (Elsevier, Amsterdam, Netherlands, 2012).
- [30] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, A learning algorithm for Boltzmann machines, *Cognit. Sci.* **9**, 147 (1985).
- [31] C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation* (WH Freeman and company San Francisco, 1973).
- [32] V. I. Arnold, *Geometrical Methods in the Theory of Ordinary Differential Equations* (Springer Science & Business Media, New York, 2012), Vol. 250.
- [33] F. W. Hehl, P. Von der Heyde, G. D. Kerlick, and J. M. Nester, General relativity with spin and torsion: Foundations and prospects, *Rev. Mod. Phys.* **48**, 393 (1976).
- [34] M. Nakahara, *Geometry, Topology and Physics* (Institute of Physics Publishing, Bristol, 2003).
- [35] S.-I. Amari, Natural gradient works efficiently in learning, *Neural Comput.* **10**, 251 (1998).
- [36] J. E. Marsden and T. S. Ratiu, *Introduction to Mechanics and Symmetry: A Basic Exposition of Classical Mechanical Systems* (Springer, New York, 2013), Vol. 17.