Hebbian Physics Networks: Self-Organized Transport Geometry from Local Conservation Laws

Gunjan Auti[1], Hirofumi Daiguji[1], Gouhei Tanaka[2,3]

[1]Department of Mechanical Engineering, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Japan, [2]Program of Computational Intelligence, Graduate School of Engineering, Nagoya Institute of Technology, Nagoya, 466-8555, Japan, [3]International Research Center for Neurointelligence, Institutes for Advanced Study, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Japan

Many physical transport processes arise from local interactions that redistribute conserved quantities while progressively reducing imbalance. In conventional numerical methods, this organization is enforced through fixed discrete operators acting on prescribed grids or meshes. We present the Hebbian Physics Network (HPN), a computational framework that replaces this rigid operator scaffold with an adaptive transport geometry that evolves directly from local conservation-law violations. In HPN, physical states are associated with graph nodes, while constitutive transport weights live on graph edges, yielding a coupled dynamical system of state and structure. The key driving signal is the residual: a local measure of unresolved imbalance, such as violation of continuity, momentum balance, or energy conservation. These residuals act as thermodynamic forces that drive the joint evolution of both the physical state and the transport operator itself.

The central mechanism of HPN is a local plasticity rule for edge weights. Transport weights adapt through a residual-modulated Hebbian update that depends only on information available at neighboring nodes and edges, without any global objective function or backpropagated error signal. This makes the framework fundamentally different from optimization-based physics-informed machine learning approaches, where physical consistency is typically imposed through a global loss. In HPN, conservation is embedded directly in the update mechanism: transport is restored by local residual relaxation through the evolving operator itself. In the near-equilibrium regime, this adaptive dynamics yields an effective transport operator that converges to a symmetric positive-definite form, recovering Onsager-type reciprocity without requiring it to be imposed in advance. Away from equilibrium, the same local rule allows the system to self-organize admissible transport pathways and geometries that redistribute imbalance while remaining thermodynamically stable.

This perspective recasts physical computation as a process of self-organization. Rather than solving a fixed discretized equation with predetermined coefficients, HPN evolves both the constitutive operator and the state through local interaction. The resulting computation is interpretable at the level of the graph: one can directly monitor how residuals reorganize transport structure, how local pathways strengthen or weaken, and how macroscopic behavior emerges from distributed adaptation. The framework therefore offers a physically grounded route toward computation in systems where the effective transport structure is unknown, history-dependent, or itself part of the dynamics.

We demonstrate this framework on canonical transport problems, including scalar diffusion and incompressible lid-driven cavity flow. Starting from random initial conditions and without prescribing a global solver structure, HPN develops physically consistent transport operators and flow organization solely through residual-driven local adaptation. These results suggest that local conservation and plastic transport geometry are sufficient to produce meaningful macroscopic computation. More broadly, HPN offers a bridge between non-equilibrium thermodynamics, self-organization, and scientific computing by treating numerical solution not as the execution of a fixed operator, but as an emergent relaxation process in which state and constitutive structure coevolve.