

Geometric Irreversibility and Dissipation Bounds in Non-Reciprocal Neural Attractor Dynamics

Liangyu Luo

niu985@sina.com

Sechenov University

Youyang Li

Sechenov University

Research Article

Keywords:

Posted Date: July 9th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10094443/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Geometric Irreversibility and Dissipation Bounds in Non-Reciprocal Neural Attractor Dynamics

Liangyu Luo^{1,*} and Youyang Li^{1,†}

¹*International School, I.M. Sechenov First Moscow State Medical University, Moscow 119048, Russia*

(Dated: June 19, 2026)

Non-reciprocal attractor networks can possess equilibrium-like stationary densities while sustaining irreversible probability currents. This separation implies that scalar landscape descriptions can correctly predict attractor occupancies and even leading action barriers, yet fail to predict the geometry and thermodynamic cost of rare transitions. Here we show that the solenoidal component of the drift controls a quantitative triad: the asymmetry of forward and reverse transition paths, the steady-state entropy production, and the precision bound for time-integrated neural currents. In a solvable bistable model, the quasipotential remains unchanged while the minimum-action escape and relaxation paths form directed loops; the corresponding path-asymmetry index scales linearly with the non-reciprocal drift, whereas entropy production scales quadratically. Monte Carlo simulations recover the Freidlin–Wentzell barrier, and thermodynamic uncertainty products remain above the stochastic-thermodynamic bound. We then extend the calculation to high-dimensional non-reciprocal attractor networks and show that the same path–current–dissipation structure persists under full-space geometric minimum-action optimization. These results identify a falsifiable nonequilibrium signature of attractor computation: two systems with the same landscape can have different transition routes, different irreversible currents, and different thermodynamic precision costs.

I. INTRODUCTION

The brain is not an equilibrium device. Neural activity is maintained by metabolic energy consumption, ionic gradients, synaptic drive and directed recurrent interactions. At the level of collective dynamics, these mechanisms produce stochastic trajectories that can settle near metastable activity states while still carrying irreversible probability currents. This matters because many computations are not represented by static states alone. A memory must sometimes be stable and sometimes switch; a decision circuit must suppress noise while remaining responsive; a cortical or neuromorphic system must transport information through state space within finite time and finite energy budgets [1–8].

Classical attractor-network theory has supplied a powerful language for memory and decision states [9–11]. In its simplest form, that language resembles equilibrium statistical mechanics: stable neural states are minima of a scalar landscape, and noise-induced transitions are controlled by barrier heights. This approximation is useful but incomplete. Real neural circuits contain asymmetric synaptic couplings, adaptation, delays, external drive and task-dependent inputs. Such ingredients break detailed balance and can generate circulating probability currents. A scalar landscape can describe the stationary density of a system only in special cases; even when a quasipotential exists, it need not describe the direction, geometry or resolved thermodynamic cost of the coarse-grained dynamics.

The most severe failure of a purely local or mean-field description occurs for rare events. Spontaneous memory errors, perceptual switches and transitions between competing decision states are not typical fluctuations around an attractor. They are exponentially unlikely paths that cross basin boundaries. Direct simulation is inefficient in the weak-noise regime, because the mean first-passage time grows exponentially as the inverse noise strength. Freidlin–Wentzell large-deviation theory addresses this problem by replacing brute-force sampling with a variational principle over paths, within the broader large-deviation and macroscopic-fluctuation perspective on nonequilibrium systems [12–15]. The minimizer is the most probable rare transition trajectory, also called an instanton or minimum-action path. In non-gradient systems, this path is not generally the time reversal of deterministic relaxation, and it cannot be inferred from local curvature or a potential barrier alone [16, 17].

Rare-event geometry is only one side of the problem. Neural computation also has an energetic dimension. Stochastic thermodynamics quantifies irreversibility by entropy production [18]. Thermodynamic uncertainty relations (TURs) further state that the relative fluctuations of a steady current are bounded by the product of entropy production and observation time [19–22]. These inequalities formalize a speed–precision–dissipation trade-off that has been observed in biological sensing and information processing [23–28]. A neural or neuromorphic system can reduce uncertainty in a time-integrated current only by increasing dissipation, increasing integration time or changing the dynamical architecture that generates the current. Related inference methods estimate dissipation from partial trajectories or coarse observations, which is essential for both neural recordings and physical devices

* niu985@sina.com

† liyongyangchn@outlook.com

[29–31].

The present work addresses three questions. First, can two attractor systems with the same stationary density and the same quasipotential barrier nevertheless have different rare-transition geometries? Second, can the geometric asymmetry of forward and reverse transition paths be connected quantitatively to steady-state entropy production? Third, can this path-level irreversibility be propagated into a thermodynamic precision bound for neural or neuromorphic current observables?

We answer these questions in a hierarchy of models: an analytically solvable non-reciprocal normal form, finite-noise transition ensembles, and high-dimensional asymmetric attractor networks optimized in the full state space. The models isolate the central mathematical feature: non-reciprocal drift creates probability currents and irreversible switching paths even when the scalar quasipotential is unchanged. The result is not merely a computational framework, but a testable nonequilibrium signature: systems that share a landscape may differ in transition routes, irreversible currents and resolved thermodynamic precision costs.

II. NON-RECIPROCAL ATTRACTOR DYNAMICS AND PATH-CURRENT SEPARATION

A. Solvable normal form

We first define a coarse-grained state space in which metastable neural states can be analysed as attractors. Let $X_t^\epsilon \in \mathbb{R}^N$ denote a neural state, such as population firing rates, latent neural coordinates or memory-overlap variables. The stochastic dynamics is written as

$$dX_t^\epsilon = F(X_t^\epsilon) dt + \sqrt{2\epsilon} \Sigma(X_t^\epsilon) dW_t, \quad (1)$$

where F is the deterministic drift, ϵ is the dimensionless noise strength and $A(x) = \Sigma(x)\Sigma^T(x)$ is the diffusion tensor. The state dependence of $\Sigma(x)$ is not only a mathematical convenience. In coarse-grained neural models, Poisson spiking produces fluctuations whose variance scales with firing rate, synaptic release noise depends on local presynaptic activity, and channel noise depends on membrane state. Diffusion approximations to these jump processes therefore naturally lead to multiplicative noise. The additive-noise examples below are chosen for transparent validation; the Freidlin–Wentzell action and the geometric minimum-action method carry over by replacing the Euclidean norm with the $A^{-1}(x)$ diffusion metric. A recurrent rate-network example is

$$F(x) = -T^{-1}x + WS(x - \vartheta) + I, \quad (2)$$

where S is a nonlinear activation function. The coupling matrix can be decomposed as $W = W^s + W^a$, where $W^s = (W + W^T)/2$ and $W^a = (W - W^T)/2$. The antisymmetric part represents non-reciprocal interactions.

When this component is present, the drift is not generally a gradient flow and the steady state may carry a nonzero probability current.

For a diffusion process, the probability density satisfies

$$\partial_t \rho(x, t) = -\nabla \cdot J(x, t), \quad J = F\rho - \epsilon \nabla \cdot (A\rho). \quad (3)$$

Detailed balance requires $J_{ss} = 0$. Broken detailed balance is therefore not simply a statement about a non-symmetric matrix; it is a statement about a measurable circulation in state space. A system can have a clear stationary density and still be irreversible if that density is accompanied by a circulating current.

To make this distinction visible, we study the two-coordinate normal form

$$dX_t = F(X_t) dt + \sqrt{2\epsilon} dW_t, \quad X_t = (x_t, y_t), \quad (4)$$

with

$$F(x, y) = -\nabla U(x, y) + \chi R \nabla U(x, y), \quad R = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \quad (5)$$

and

$$U(x, y) = \frac{1}{4}(x^2 - 1)^2 + \frac{1}{2}y^2. \quad (6)$$

The two stable states $A = (-1, 0)$ and $B = (1, 0)$ represent coarse-grained memory or decision states, while $S = (0, 0)$ is the index-one saddle separating their basins. The parameter χ controls non-reciprocity. When $\chi = 0$, the system is an equilibrium gradient flow. When $\chi \neq 0$, the rotational term preserves the fixed points but creates state-space circulation.

At $\chi = 0.8$, deterministic trajectories spiral into the two stable memory states rather than descending directly along potential contours (Fig. 1). The eigenvalues at A and B are $-1.5 \pm 1.015i$, indicating damped rotational relaxation, whereas the saddle has eigenvalues ± 1.281 . Thus the same attractor layout supports a non-equilibrium flow field. This is the simplest possible demonstration of the main point: stability and irreversibility are different dynamical properties.

B. Same quasipotential, different currents

The benchmark model has an analytic stationary density,

$$\rho_{ss}(x, y) \propto \exp[-U(x, y)/\epsilon], \quad (7)$$

because the rotational term is orthogonal to ∇U . This property makes the model useful: the stationary density looks equilibrium-like, but the steady current is nonzero:

$$J_{ss}(x, y) = \chi R \nabla U(x, y) \rho_{ss}(x, y). \quad (8)$$

Thus density alone would miss the non-equilibrium character of the system. Figure 2 shows two density peaks




FIG. 1. **Non-equilibrium bistable neural flow.** Contours show the scalar quasipotential $U(x, y)$, while streamlines show the full drift $F = -\nabla U + \chi R \nabla U$ at $\chi = 0.8$. The fixed points A and B are stable memory states and S is an index-one saddle. The rotational component leaves the attractor locations intact but changes the geometry of relaxation and escape.

centred on the memory states. At $\chi = 0$ the current vanishes; at $\chi = 0.8$ the density is unchanged but the current circulates. The figure is the state-space analogue of an observation that a system can visit the same activity states repeatedly while still traversing them in a preferred temporal direction.

C. Path-current-dissipation relation

The normal form makes the relation between geometry, current and dissipation explicit. The drift decomposes into a gradient component and a solenoidal component,

$$F_{\text{sol}}(x) = \chi R \nabla U(x). \quad (9)$$

Since $R \nabla U$ is orthogonal to ∇U , this force changes the local direction of motion without changing the stationary density or the quasipotential barrier. It does, however, generate the steady current in Eq. (8). The entropy production rate associated with the resolved diffusion dynamics is

$$\langle \sigma \rangle = \frac{\chi^2}{\epsilon} \langle |\nabla U|^2 \rangle_{\rho_{\text{ss}}}, \quad (10)$$

so dissipation is a second-order response to the non-reciprocal drift.

The transition geometry responds at first order. Perturbing the minimum-action escape and relaxation branches around the gradient limit gives

$$\mathcal{A}_{\text{path}}(\chi) = C_{\text{geom}} \chi + O(\chi^3), \quad \chi \ll 1, \quad (11)$$

where $\mathcal{A}_{\text{path}}$ denotes the forward-backward loop asymmetry defined below and C_{geom} depends on the shape of

U and on the chosen path metric. Combining Eqs. (10) and (11) yields the central scaling relation

$$\mathcal{A}_{\text{path}}^2 = C_{\text{diss}} \epsilon \langle \sigma \rangle + O(\chi^4), \quad \chi \ll 1. \quad (12)$$

Thus, in the solvable normal form, geometric path irreversibility is the square-root signature of entropy production. The result is not claimed as a universal theorem for arbitrary neural dynamics; it is a controlled, falsifiable relation showing how an observable path-level asymmetry can track the dissipative current hidden beneath an apparently equilibrium-like landscape. A perturbative derivation is given in Appendix A.

III. RARE TRANSITION GEOMETRY

A. Minimum-action paths

The state vector in Eq. (1) is a coarse-grained coordinate. In an experiment it could correspond to latent variables inferred from population recordings, electrocorticography, fMRI time series or behaviourally defined decision coordinates. In a model it can represent firing-rate modes, memory overlaps or reduced collective variables. The framework requires that the chosen variables support an approximate Markov description on the timescale of interest and that the effective diffusion be specified. If hidden variables are important, the inferred entropy production should be interpreted as an observed-scale estimate rather than a total microscopic entropy production.

The two-dimensional benchmark was chosen to isolate the effect of non-reciprocity. It should not be read as a direct model of the human brain. The point of the benchmark is that it separates four quantities that are often conflated: attractor locations, quasipotential barriers, transition-path geometry and entropy production. This separation makes it possible to validate the computational pipeline in a system where analytic checks are available.

Freidlin-Wentzell action from short-time transition probabilities.

For additive noise,

$$dX_t^\epsilon = F(X_t^\epsilon)dt + \sqrt{2\epsilon} B dW_t, \quad A = BB^T, \quad (13)$$

the short-time increment is approximately Gaussian:

$$X_{t+\Delta t} - X_t \approx F(X_t)\Delta t + \sqrt{2\epsilon\Delta t} B \xi, \quad \xi \sim \mathcal{N}(0, I). \quad (14)$$

The logarithmic cost of observing an increment $\Delta x \approx \dot{x}\Delta t$ is

$$\frac{1}{4\epsilon} (\dot{x} - F)^T A^{-1} (\dot{x} - F) \Delta t. \quad (15)$$

Multiplying such costs along a path and taking the continuum limit yields Eq. (37). The inverse diffusion tensor




FIG. 2. **Same landscape, different dynamics.** (A) At $\chi = 0$, the model has the same bistable density landscape but no stationary current. (B) At $\chi = 0.8$, the density is unchanged while the probability current circulates around the attractor basins. (C) The integrated current norm grows linearly with χ , whereas entropy production grows quadratically.

acts as the metric: deviations along highly noisy directions are cheaper than deviations along weakly noisy directions. The factor of $1/4$ follows from the convention $\sqrt{2\epsilon}\Sigma dW_t$.

Hamiltonian formulation and quasipotential.

The Lagrangian associated with the path action is

$$L(x, \dot{x}) = \frac{1}{4}[\dot{x} - F(x)]^T A^{-1}(x)[\dot{x} - F(x)]. \quad (16)$$

The conjugate momentum is

$$p = \frac{1}{2}A^{-1}(x)[\dot{x} - F(x)], \quad (17)$$

and the Hamiltonian is

$$H(x, p) = p^T F(x) + p^T A(x)p. \quad (18)$$

The quasipotential satisfies the stationary Hamilton–Jacobi equation

$$H(x, \nabla V(x)) = 0. \quad (19)$$

For the benchmark model, $A = I$ and $F = -\nabla U + \chi R \nabla U$. Substituting $V = U$ into Eq. (19) gives

$$\nabla U^T(-\nabla U + \chi R \nabla U) + |\nabla U|^2 = \chi \nabla U^T R \nabla U = 0, \quad (20)$$

because R is antisymmetric. Thus the quasipotential equals the scalar potential even though the dynamics violates detailed balance. This special property lets us verify the fitted action barrier against an analytic value.

Geometric action.

Because the optimal transition time is unknown and can diverge in the weak-noise limit, the time-parametrized action is often replaced by a geometric action. Let $\phi(t) = \gamma(s(t))$ with $s \in [0, 1]$. Minimizing the time action over positive reparametrizations gives

$$S_g[\gamma] = \frac{1}{2} \int_0^1 [\|F(\gamma)\|_A \|\gamma'(s)\|_A - \langle F(\gamma), \gamma'(s) \rangle_A] ds. \quad (21)$$

The minimizer of S_g is the same geometric curve as the minimizer of the time-dependent action. A general gMAM implementation discretizes the path, descends on the discrete action, removes tangential updates that only reparametrize the curve and redistributes nodes to maintain approximately uniform arclength.

For a general model, the discrete path is represented by nodes $\gamma_0, \dots, \gamma_M$, with γ_0 fixed at the source attractor and γ_M fixed at a candidate separatrix point. A midpoint discretization is

$$S_g^{(h)} = \frac{1}{2} \sum_{m=0}^{M-1} [\|F(\bar{\gamma}_m)\|_{A_m} \|\Delta\gamma_m\|_{A_m} - \langle F(\bar{\gamma}_m), \Delta\gamma_m \rangle_{A_m}], \quad (22)$$

where

$$\Delta\gamma_m = \gamma_{m+1} - \gamma_m, \quad \bar{\gamma}_m = \frac{\gamma_{m+1} + \gamma_m}{2}, \quad A_m = A(\bar{\gamma}_m). \quad (23)$$

The interior nodes can be updated by projected descent,

$$\tilde{\gamma}_m^{(n+1)} = \gamma_m^{(n)} - \eta_n \mathcal{P}_m^{(n)} \frac{\partial S_g^{(h)}}{\partial \gamma_m}. \quad (24)$$

The projection $\mathcal{P}_m$ removes the tangential component of the gradient, because tangential motion changes only the parametrization of the curve. After each update, nodes are redistributed to equal arclength in the diffusion metric. Convergence should be checked by monitoring the action decrease, the projected gradient norm and the Hamiltonian residual along the reconstructed path. In high dimensions, automatic differentiation is often the most reliable way to differentiate the discrete action.

Instanton computation in the benchmark.

In the benchmark, the uphill instanton can be computed from the Hamiltonian structure. Since $V = U$, the escape dynamics is Eq. (40). The escape branch from an attractor to the saddle is obtained by integrating the stable manifold of the uphill dynamics backward from the saddle. The relaxation branch from the saddle to the target attractor is obtained by integrating the deterministic

drift forward. The discrete geometric action reported in the data files is evaluated on the resulting path using a midpoint rule.

Fixed points and stability.

For the normal form, the fixed points are analytic: $A = (-1, 0)$, $S = (0, 0)$ and $B = (1, 0)$. In a general rate network, fixed points are found by solving $F(x) = 0$ from multiple initial guesses and clustering duplicate solutions. Stability is determined by the Jacobian eigenvalues. A fixed point is locally stable if all eigenvalues have negative real part. An index-one saddle has exactly one eigenvalue with positive real part and typically organizes the dominant transition channel between neighbouring basins.

Generalization to high-dimensional recurrent networks.

The normal form can be replaced by a high-dimensional recurrent network without changing the pipeline. This extension connects the present benchmark to recent work on non-reciprocal spin-glass networks and task-dependent entropy flow in neuronal population models [32, 33]. A common rate model is

$$dX_t = [-X_t + W \tanh(\beta X_t) + I]dt + \sqrt{2\epsilon\sigma_0}dW_t, \quad X_t \in \mathbb{R}^N \quad (25)$$

The coupling matrix can be decomposed as

$$W = W^s + \chi W^a. \quad (26)$$

The symmetric component can store patterns through a Hebbian rule,

$$W_{ij}^s = \frac{J}{N} \sum_{\mu=1}^P \xi_i^\mu \xi_j^\mu, \quad W_{ii}^s = 0, \quad (27)$$

while the antisymmetric component can be generated from a random matrix G as

$$W^a = \frac{G - G^T}{2\sqrt{N}}. \quad (28)$$

Memory state μ can be monitored by the overlap

$$q_\mu(t) = \frac{1}{N} \sum_i \xi_i^\mu X_i(t). \quad (29)$$

Transitions are detected when a different overlap becomes dominant and remains dominant for a commitment time. In such networks, the minimum-action path generally lives in the full high-dimensional state space even if it is visualized by overlaps or principal components.

High-dimensional systems introduce multiple transition channels. The dominant escape route need not pass through the Euclidean nearest saddle, and it need not be visible in a simple projection. A practical strategy is to identify candidate saddles or basin-boundary points by deterministic basin analysis, initialize gMAM paths to

each candidate and select the smallest action. Continuation in non-reciprocity, input strength or noise anisotropy can improve convergence. The final object is a directed graph of metastable states, where each edge carries an action barrier, a transition path, a fitted rate and a path-asymmetry measure.

For the scalability test reported below we used controlled N -dimensional embeddings of the same non-equilibrium normal form, with $N \in \{50, 100, 200, 500\}$ and four random seeds per size. The potential was

$$U_N(x) = \frac{1}{4}(x_1^2 - 1)^2 + \frac{1}{2} \sum_{i=2}^N \lambda_i x_i^2, \quad (30)$$

and the drift was

$$F_N(x) = -\nabla U_N(x) + \chi K \nabla U_N(x), \quad (31)$$

where $K = -K^T$ is a sparse random antisymmetric matrix. The $N = 100$ realization contains 148 nonzero coupling pairs, and the same sparse construction is used for the other system sizes. This choice was made to keep the high-dimensional benchmark in a locally stable, interpretable attractor regime: a dense antisymmetric matrix would generate strong rotational mixing in many transverse directions and could obscure whether failure came from the variational algorithm or from a deliberately turbulent landscape. The coupling density is therefore not claimed as a cortical connectivity estimate; it is a controlled stress test that preserves local attractor stability while introducing enough non-reciprocal structure to bend the transition channel. The fixed points remain $A = (-1, 0, \dots, 0)$, $B = (1, 0, \dots, 0)$ and $S = (0, \dots, 0)$, but the escape path is optimized in the full ambient state space. The geometric action was discretized with 48 path nodes for the multi-size scan and minimized by projected descent with arclength redistribution. This example is intentionally more controlled than a fully random recurrent network: it tests whether the variational pipeline converges in high dimension while preserving analytic checks on the barrier and entropy production.

Monte Carlo first-passage simulations.

Stochastic trajectories are integrated by Euler–Maruyama:

$$X_{n+1} = X_n + F(X_n)\Delta t + \sqrt{2\epsilon\Delta t}\xi_n, \quad \xi_n \sim \mathcal{N}(0, I). \quad (32)$$

Trajectories are initialized near A . A first passage from A toward B is recorded when the trajectory first crosses the separatrix $x = 0$. The empirical rate is $k = 1/\langle\tau\rangle$. The smallest noise value has a small censored fraction; this is reported in `transition_rates.csv`. At lower noise strengths, direct Monte Carlo would become inefficient and should be replaced by importance sampling, adaptive multilevel splitting or transition path sampling.

Enhanced sampling and uncertainty quantification.

In lower-noise regimes, a direct first-passage estimate can require an exponentially large number of trajectories. Enhanced sampling methods are then needed. Importance sampling can bias trajectories toward the minimum-action path and reweight them to recover unbiased estimates. Adaptive multilevel splitting repeatedly clones trajectories that make progress along a reaction coordinate. Transition path sampling samples the ensemble of reactive trajectories directly. The minimum-action path computed by gMAM is useful in all three settings because it supplies an approximate reaction coordinate and an initial transition tube.

Uncertainty quantification is equally important. In empirical applications, the drift and diffusion are inferred from finite data, and the rare-event path often travels through low-density regions where inference is least certain. One should therefore report uncertainty in the model parameters, in the quasipotential barrier and in the transition path. Possible approaches include bootstrap resampling of trajectories, Bayesian latent SDE inference, ensembles of fitted drift models and sensitivity analysis with respect to diffusion anisotropy. A strong claim that an observed transition follows an instanton should show that the predicted path is stable under these uncertainties.

Entropy production and TUR estimation.

For the benchmark, entropy production is evaluated using the analytic steady density in Eq. (7). The expectation in Eq. (46) is computed by numerical quadrature. Current statistics are estimated from independent steady-state trajectory segments of duration $T = 60$. The Stratonovich current in Eq. (47) is discretized by the midpoint rule:

$$\hat{J}_T = \sum_{n=0}^{K-1} \chi R \nabla U \left(\frac{X_{n+1} + X_n}{2} \right)^T (X_{n+1} - X_n). \quad (33)$$

The mean and variance across trajectory segments are inserted into Eq. (49). A finite-sample estimate below one would indicate insufficient sampling, non-stationarity or an inappropriate current observable.

When the stationary density and current are not analytic, they can be estimated from trajectories. A grid estimator partitions the state space into cells, estimates residence time in each cell and estimates net probability flux through cell faces. Kernel estimators replace hard cells by smooth weights and are often less noisy in low dimensions. In high dimensions, direct density-current estimation becomes difficult; one may instead estimate entropy production from time-reversal asymmetry, projected currents or variational lower bounds. All such estimators are sensitive to sampling rate. If the sampling interval is too coarse, fast irreversible loops are missed and entropy production is underestimated.

The current observable should be chosen before testing the TUR. In a decision model, it might be accumulated evidence along a decision coordinate. In a sequential memory model, it might be circulation around a se-

quence of stored patterns. In a neuromorphic device, it might be charge transport or progress through a state-machine cycle. The TUR applies to time-antisymmetric currents in a steady state. Applying it to a state occupation, a non-stationary transient or a current with nearly zero mean can produce misleading numerical values.

Synthetic population readout.

The synthetic population readout is generated from one stochastic transition. Each of 80 units is assigned a preferred memory coordinate and a random context gain. Its activity is a noisy sigmoidal function of the latent coordinates. Principal component analysis is then applied to the activity matrix. This procedure is used only to illustrate how a curved state-space transition could appear in population recordings; it is not used to validate the theory.

Empirical workflow for neural recordings.

A complete empirical workflow would proceed as follows. First, infer latent coordinates from neural recordings and identify metastable states in that latent space. Second, fit a stochastic dynamical model to estimate drift and diffusion. Third, compute candidate minimum-action paths between metastable states using the fitted model. Fourth, divide observed transitions into training and held-out sets. The training set may be used to fit the model, but the held-out transitions should be used to test whether the predicted transition tube explains the data better than an equilibrium surrogate or a straight-line interpolation. Fifth, estimate probability currents and time-reversal asymmetry in the same latent coordinates. Finally, define a task-relevant current and test whether its fluctuations are consistent with the TUR bound within sampling uncertainty.

This workflow makes the empirical claim falsifiable. If observed transition paths do not concentrate around the predicted tube, the inferred drift or noise model is incomplete. If the stationary density is well fit but the transition directionality is not, an equilibrium landscape model may explain occupancy but not dynamics. If the inferred entropy production changes across task conditions, one can ask whether those changes track behavioural speed, precision or error rate. The theory therefore suggests concrete experiments rather than only a metaphor of landscapes and flows.

Physical interpretation of the entropy-production estimator.

For a diffusion process, the steady probability current can be written as

$$J_{ss}(x) = \rho_{ss}(x) v_{ss}(x), \quad (34)$$

where v_{ss} is the local mean velocity of probability mass in state space. The entropy-production rate in Eq. (45) is the steady-state average of the squared local velocity measured in the inverse diffusion metric:

$$\langle \sigma \rangle = \frac{1}{\epsilon} \int \rho_{ss}(x) v_{ss}^T(x) A^{-1}(x) v_{ss}(x) dx. \quad (35)$$

This expression makes the physical meaning transparent. A current is costly when it moves probability mass persistently through directions in which diffusion is weak. Conversely, a current aligned with strongly noisy directions is less costly in the effective stochastic description. For coarse-grained neural or neuromorphic state variables, this means that the resolved thermodynamic cost depends not only on the magnitude of directed flow but also on the noise geometry of the observed variables.

This interpretation also explains why a scalar potential is insufficient for thermodynamic analysis. The stationary density determines how often the system visits each region, but the entropy production depends on how probability mass circulates through those regions. A system with two density peaks and no current is equilibrium-like. A system with the same two density peaks and a circulating current is dissipative. Distinguishing these cases requires temporal information. In empirical recordings, this means that state-occupancy histograms, by themselves, cannot establish whether a neural system is operating near or far from detailed balance.

The estimator used in the benchmark is analytic because the density and current are known. In data applications, the same expression should be viewed as a target rather than a literal recipe. Directly estimating J_{ss}/ρ_{ss} can be unstable in regions of low density. Regularized estimators, kernel smoothing or variational lower bounds are often needed. Moreover, the estimate depends on the chosen coordinates. If the coordinates compress away a hidden cyclic variable, the observed current and entropy production decrease. Thus, an empirical value of entropy production should always be reported together with the state representation, sampling interval, noise model and estimator.

Testing non-equilibrium models against equilibrium surrogates.

A rigorous empirical study should compare the non-equilibrium model with an equilibrium surrogate. The surrogate should not be a strawman. It should be constructed to reproduce the same stationary density or the same attractor structure as well as possible. In the benchmark, the equilibrium surrogate is obtained by setting $\chi = 0$, which preserves the potential landscape but removes the circulating current. In a data-driven setting, one could fit a gradient drift to the inferred stationary density or project the drift onto its detailed-balance component.

The comparison should then use held-out dynamical observables. Stationary density alone is not enough because the surrogate is designed to match it. Useful tests include transition-path likelihood, forward-backward path asymmetry, cyclic flux around closed loops, prediction of intermediate states during switching and time-reversal classification accuracy. If the non-equilibrium model predicts held-out transition trajectories better than the equilibrium surrogate, then the current is not merely an interpretive addition; it is necessary for predictive accuracy.

TABLE I. Core numerical parameters.

Parameter	Value	Role
Memory states	$(-1, 0), (1, 0)$	Attractors
Saddle	$(0, 0)$	Separatrix
Non-reciprocity	$\chi = 0.8$	Phase paths
Rate noises	$0.055\text{--}0.110$	LDT scaling
TUR noise	$\epsilon = 0.08$	Current statistics
TUR duration	$T = 60$	Integration time
Rate trials	360 per ϵ	First passage
TUR segments	2200 per χ	Mean/variance

This model-comparison logic is especially important for high-dimensional neural data, where overfitting is a real danger. A flexible non-equilibrium model will almost always fit training trajectories better than a constrained equilibrium model. The meaningful question is whether it predicts held-out rare transitions and held-out temporal asymmetries. Cross-validation across trials, behavioural conditions or recording sessions should therefore be part of any empirical application.

Numerical quality control.

The reproducible benchmark uses fixed random seeds, explicit CSV outputs and figure-generation scripts. This is not only a convenience; it is part of the scientific claim. Rare-event calculations are sensitive to discretization, sampling windows and path initialization. Every reported figure should be traceable to a data file, and every data file should be regenerated from a script that states the model parameters.

For the action calculation, the main checks are path resolution and endpoint stability. Increasing the number of path points should not change the geometric channel qualitatively. Moving the endpoint slightly along the separatrix should not produce a lower-action channel unless a genuinely different transition route exists. For the rate calculation, the main checks are time-step refinement, censoring fraction and sensitivity of the fitted barrier to the noise window. For the TUR calculation, the main checks are burn-in length, segment duration, current mean stability and bootstrap uncertainty of the current variance.

The diagnostic plot in Fig. 11 illustrates one such check. The fitted barrier shifts when the smallest-noise points are removed, which is expected because finite-noise corrections become more important. This does not invalidate the large-deviation scaling; it shows why the asymptotic regime must be identified rather than assumed. In more complex models, the same diagnostic should be combined with enhanced sampling or independent first-passage estimates.

Reproducible data products.




FIG. 3. **Forward and reverse transition channels do not coincide.** Solid curves are fluctuation-driven instanton branches and dashed curves are deterministic relaxation branches from the saddle. In an equilibrium gradient system, the escape and return routes would collapse onto the same curve after time reversal. Here non-reciprocity bends the two switching channels into distinct phase-space routes.

B. Minimum-action paths in the normal form

Large-deviation theory predicts the most probable transition path.

Rare transitions are described by path probabilities rather than by the stationary density alone. For a smooth path $\phi : [0, T] \rightarrow \mathbb{R}^N$, Freidlin–Wentzell theory gives

$$\mathbb{P}\{X_t^e \approx \phi_t, 0 \leq t \leq T\} \asymp \exp[-I_T[\phi]/\epsilon], \quad (36)$$

where the action is

$$I_T[\phi] = \frac{1}{4} \int_0^T [\dot{\phi}_t - F(\phi_t)]^T A^{-1}(\phi_t) [\dot{\phi}_t - F(\phi_t)] dt. \quad (37)$$

The quasipotential relative to attractor m_α is

$$V_\alpha(x) = \inf_{T>0} \inf_{\phi(0)=m_\alpha, \phi(T)=x} I_T[\phi]. \quad (38)$$

The leading transition rate from m_α to m_β obeys

$$k_{\alpha \rightarrow \beta} = C_{\alpha\beta} \exp[-\Delta V_{\alpha\beta}/\epsilon] [1 + o(1)]. \quad (39)$$

This barrier is a path-space quantity, not simply a local curvature or a Kramers prefactor [34].

For the benchmark model, the quasipotential is $V = U$ because $\nabla U^T R \nabla U = 0$. However, the transition *path* is not the gradient path. The optimal uphill escape branch satisfies

$$\dot{x} = F(x) + 2\nabla V(x) = \nabla U(x) + \chi R \nabla U(x). \quad (40)$$

The downhill branch from the saddle follows the deterministic drift F . The forward transition $A \rightarrow S \rightarrow B$ and the reverse transition $B \rightarrow S \rightarrow A$ therefore form different loops in phase space (Fig. 3). This is the geometric signature of irreversibility in rare neural switching.

C. Forward–backward loop asymmetry

The separation between the forward and backward transition loops can be quantified by resampling both loops to equal arclength and computing a root-mean-square distance between the forward loop and the time reverse of the backward loop. We denote this index by $\mathcal{A}_{\text{path}}$. The key point is not only that $\mathcal{A}_{\text{path}}$ grows with non-reciprocity. In the small- χ regime, the normal form exhibits the quantitative relation anticipated in Eq. (12):

$$\mathcal{A}_{\text{path}}^2 \propto \epsilon \langle \sigma \rangle. \quad (41)$$

Thus path geometry is a first-order response to the solenoidal drift, whereas entropy production is the corresponding second-order thermodynamic cost (Fig. 4). This square-root relation is the central physical result of the solvable benchmark.

Finite-noise transition ensembles organize around the instanton.

Large-deviation theory is an asymptotic theory. At finite noise, individual reactive trajectories fluctuate around the most probable path rather than lying exactly on it. To visualize this relationship, we simulated an ensemble of stochastic transitions at $\epsilon = 0.12$ and extracted the reactive segments leaving the A basin. The resulting trajectories form a broad transition tube, but the minimum-action channel passes through the centre of the ensemble (Fig. 5). This is the practical meaning of an instanton in a neural context: it predicts the organizing centre of a noisy cloud of successful state switches.

D. Monte Carlo verification of action barriers

The benchmark model permits an independent rate test. The quasipotential barrier between A and the saddle is

$$\Delta V = U(0, 0) - U(-1, 0) = 0.25. \quad (42)$$

We simulated 360 first-passage trajectories for each noise strength

$$\epsilon \in \{0.055, 0.065, 0.080, 0.095, 0.110\} \quad (43)$$

and estimated $k_{A \rightarrow B} = 1/\langle \tau \rangle$. A linear fit of $\log k_{A \rightarrow B}$ against $1/\epsilon$ gives

$$\widehat{\Delta V} = 0.2516, \quad R^2 = 0.9962. \quad (44)$$

The fitted barrier is within one percent of the analytic value (Fig. 6). This agreement verifies that the observed transition statistics are governed by the large-deviation action rather than by local Gaussian fluctuations around the attractor.




FIG. 4. **Path irreversibility is the geometric square-root signature of dissipation.** (A) The path-asymmetry index $\mathcal{A}_{\text{path}}$ grows approximately linearly with non-reciprocity in the small- χ regime. (B) The entropy production rate grows quadratically, appearing as a linear relation in χ^2 . (C) Combining the two scalings gives $\mathcal{A}_{\text{path}}^2 \propto \epsilon(\sigma)$ for small χ . Dashed lines show small- χ fits.




FIG. 5. **Finite-noise transition ensemble.** Gray curves are reactive stochastic transition segments and the red curve is the large-deviation transition channel. The ensemble is broad because the noise is finite, but it is organized around the predicted minimum-action route.




FIG. 6. **Large-deviation scaling of first-passage rates.** The logarithm of the empirical transition rate is approximately linear in inverse noise strength. The fitted slope corresponds to $\widehat{\Delta V} = 0.252$, close to the analytic barrier 0.25. Error bars denote propagated standard errors from the mean first-passage time.

IV. HIGH-DIMENSIONAL NON-RECIPROCAL ATTRACTOR NETWORKS

A. Network construction

The high-dimensional normal form in Eqs. (30)–(31) preserves the same two attracting memory states and saddle while embedding the transition problem in a sparse non-reciprocal ambient network. The construction is intentionally controlled: it tests scaling of the variational calculation and of the path-current-dissipation relation without replacing the benchmark by an uninterpretable random dynamical system.

B. Full-space gMAM calculation

The two-dimensional benchmark is useful for intuition, but a theory intended to identify testable state-space signatures must also survive in an ambient space where the transition path is not visually obvious. We therefore repeated the large-deviation calculation in the high-dimensional antisymmetric networks defined in Eqs. (30)–(31), using $N = 50, 100, 200, 500$ and four random realizations of K for each size. The geometric minimum-action descent was initialized from a straight path between the source memory state and the saddle and optimized in the full state space, not in a preselected two-dimensional projection.

For each optimized escape branch we computed a path-asymmetry statistic $\mathcal{A}_{\text{path}}$, defined as the root-mean-square transverse excursion of the minimum-action path away from the one-dimensional equilibrium escape axis.

This statistic is zero at $\chi = 0$ and grows approximately linearly with non-reciprocity for all tested N (Fig. 7A). The corresponding entropy production rate grows linearly with χ^2 , with a slope that increases with system size because larger sparse networks contain more dissipative transverse channels (Fig. 7B). A representative $N = 200$ run shows simultaneous decrease of the geometric action and the projected gradient residual during gMAM descent (Fig. 7C).

C. Scaling with dimension and disorder

The high-dimensional thermodynamic constraint is also stable. For the housekeeping current generated by the antisymmetric force, current-segment ensembles were used to compute bootstrap confidence intervals for the uncertainty–cost product. Across $N = 50$ –500 and $\chi = 0.4, 0.8, 1.2$, the estimates remain above the TUR threshold $\mathcal{Q} = 1$ within the displayed uncertainty (Fig. 7D). Thus the high-dimensional result is no longer a single illustrative PCA projection; it is a scaling and robustness test of rare-event geometry, dissipation and current precision.

V. THERMODYNAMIC PRECISION BOUNDS

A. Entropy production

The same non-reciprocal current that bends transition paths also has a thermodynamic cost in the resolved coarse-grained Markovian dynamics. For additive diffusion, the steady-state entropy production rate is

$$\langle \sigma \rangle = \frac{1}{\epsilon} \int \frac{J_{ss}^T(x) J_{ss}(x)}{\rho_{ss}(x)} dx. \quad (45)$$

For the benchmark model this is the same quadratic law derived in Eq. (10),

$$\langle \sigma \rangle = \frac{\chi^2}{\epsilon} \langle |\nabla U|^2 \rangle_{\rho_{ss}}. \quad (46)$$

We then define a time-integrated housekeeping current,

$$\mathcal{J}_T = \int_0^T \chi R \nabla U(X_t) \circ dX_t, \quad (47)$$

where $\circ$ denotes the Stratonovich convention.

B. TUR for neural current observables

The thermodynamic uncertainty relation gives

$$\frac{\text{Var}(\mathcal{J}_T)}{\langle \mathcal{J}_T \rangle^2} \langle \sigma \rangle T \geq 2, \quad (48)$$

or equivalently

$$\mathcal{Q} = \frac{\text{Var}(\mathcal{J}_T)}{\langle \mathcal{J}_T \rangle^2} \frac{\langle \sigma \rangle T}{2} \geq 1. \quad (49)$$

At fixed $\epsilon = 0.08$ and $T = 60$, the empirical relative variance decreases as entropy production increases, but all measurements remain above the TUR bound. The uncertainty–cost product ranges from $\mathcal{Q} = 1.02$ at $\chi = 0.2$ to $\mathcal{Q} = 2.44$ at $\chi = 1.2$ (Fig. 8). Thus improved current precision is not free: it is paid for by dissipation, time or architectural changes.

C. Speed–precision–dissipation design map

Equation (48) can be read as a design inequality:

$$\frac{\text{Var}(\mathcal{J}_T)}{\langle \mathcal{J}_T \rangle^2} \geq \frac{2}{\langle \sigma \rangle T}. \quad (50)$$

For a neural or neuromorphic device, this inequality states that a short integration time demands higher dissipation for the same relative precision. Figure 9 plots the lower bound on relative variance as a function of non-reciprocity and integration time. The map gives a compact visual form of the speed–precision–energy trade-off: fast and accurate currents require either high entropy production or an architecture closer to the bound.

Population readouts can recover the curved transition geometry.

To connect the state-space theory to neural population data, we generated a synthetic readout of one rare transition. Eighty units were assigned preferred memory coordinates and random context gains. Their activities were generated from the simulated latent trajectory and then projected by principal component analysis. The population activity changes gradually across units, and the PCA trajectory recovers a curved low-dimensional path (Fig. 10). This analysis is illustrative rather than empirical, but it shows how the theoretical instanton could be compared with latent trajectories extracted from neural recordings.

Estimator diagnostics identify finite-noise and finite-sample effects.

Two diagnostics are important for interpreting the numerical results. First, the rate-fit barrier depends mildly on the noise window. Including the smallest noise strength gives the estimate closest to the analytic barrier; removing it shifts the fitted barrier upward because finite-noise corrections and prefactors become more visible. Second, the finite-sample estimates of $\mathcal{Q}$ remain above the TUR threshold for every χ (Fig. 11). These diagnostics are essential in applications where the true barrier is unknown.

The full workflow joins rare-event and thermodynamic observables.




FIG. 7. **High-dimensional scaling and robustness.** (A) Path-asymmetry statistic $\mathcal{A}_{\text{path}}$ versus non-reciprocity for $N = 50, 100, 200, 500$. Points and bars show mean and 95% confidence intervals across four random antisymmetric matrices. (B) Entropy production rate versus χ^2 , showing the expected quadratic dependence on non-reciprocity and the increase of total dissipation with system size. (C) Representative $N = 200$ gMAM diagnostic: the geometric action decreases while the projected gradient residual contracts. (D) TUR product Q with bootstrap 95% confidence intervals for high-dimensional housekeeping-current segments. The dashed line marks the thermodynamic bound $Q = 1$.

The complete computational pipeline is shown in Fig. 12. A specified noisy neural dynamical system defines attractors, saddles and diffusion geometry. Large-deviation theory computes transition paths and action barriers. Monte Carlo or enhanced-sampling simulations validate first-passage scaling. Steady-state trajectories estimate currents, entropy production and TUR products. The important point is that all observables are computed within one dynamical model. Transition geometry and thermodynamic cost are not separate stories; they are two projections of the same non-equilibrium drift.

An equilibrium surrogate fails to describe transition directionality.

A useful way to understand the role of non-equilibrium currents is to compare the benchmark with an equilibrium surrogate that has the same stationary density. Setting $\chi = 0$ produces a gradient system with the same potential U and the same attractor locations. Because the quasipotential is also U in the non-equilibrium benchmark, the equilibrium surrogate predicts the same leading exponential barrier. If one measured only the station-

ary density or only the asymptotic slope of the transition rate, the two systems would appear closely related. They differ, however, in the object that matters for neural state switching: the transition path.

In the equilibrium surrogate, the escape path from A to the saddle is the time reversal of deterministic relaxation from the saddle to A . The return transition uses the same geometric channel. Therefore a single landscape ridge organizes both directions of switching. In the non-equilibrium system, the escape branch and the relaxation branch are deflected differently by the rotational drift. The result is a loop, not a line. This loop is not a decorative feature of the phase portrait; it implies that different intermediate states are visited during forward and reverse switching. In a neural system, this difference could correspond to direction-specific population activity during a memory error or perceptual transition.

This comparison clarifies why the benchmark was chosen. If the rotational drift changed both the barrier and the path, it would be difficult to decide which effect was responsible for changes in transition statistics. Here the




FIG. 8. **Thermodynamic uncertainty bound and Pareto structure.** Left: relative current variance decreases with entropy production and stays above the TUR lower bound. Right: the uncertainty-cost product $\mathcal{Q}$ remains greater than one for all tested values of non-reciprocity.




FIG. 9. **TUR design map.** The colour scale shows the lower bound on relative current variance as a function of non-reciprocity and integration time. White contours mark constant lower-bound precision levels.

barrier is analytically fixed while the path and entropy production vary with χ . The model therefore isolates the effect of irreversibility. It shows that non-equilibrium structure can be hidden from scalar landscape analyses even when the landscape itself is exactly known.

Quantitative synthesis of the model-generated evidence.

The model-generated results support four quantitative claims. First, two systems can have the same quasipotential barrier while exhibiting different transition routes. Second, the path-asymmetry index grows as $\mathcal{A}_{\text{path}} \sim \chi$, while entropy production grows as $\langle \sigma \rangle \sim \chi^2/\epsilon$, producing the small- χ relation $\mathcal{A}_{\text{path}}^2 \propto \epsilon \langle \sigma \rangle$. Third, first-passage rates obey the large-deviation scaling predicted by the quasipotential, with a fitted barrier of 0.2516 compared with the analytic value 0.25. Fourth, the time-integrated

housekeeping current satisfies the TUR in both the solvable benchmark and the high-dimensional scaling tests.

These results are deliberately not presented as a claim about measured human brain data. They are a controlled demonstration of a nonequilibrium signature that can be tested in neural or neuromorphic state-space data. The strength of the demonstration is that every quantity is computed from the same model: the attractors, transition paths, action barrier, rate scaling, probability current, entropy production and TUR product. Here the path-level and thermodynamic claims refer to the same state variables and the same stochastic dynamics.

Predictions for future experiments.

The results identify several experimentally testable predictions. First, if a latent neural state space has broken detailed balance, then forward and reverse transitions between the same two metastable states should not occupy the same tube after time reversal. This can be tested by aligning spontaneous state switches in population recordings and comparing the two reactive ensembles. Second, perturbations that increase effective non-reciprocity should increase transition-path asymmetry even if the occupancy landscape changes only weakly. Third, task conditions that demand faster and more precise state-space currents should show increased resolved entropy production or increased time-reversal asymmetry in the relevant latent coordinates.

Fourth, equilibrium surrogate models fitted to the same stationary density should fail to predict direction-specific transition intermediates when the real dynamics is irreversible. This is a strong test because it does not rely on absolute entropy-production estimates. One can compare predictive accuracy for held-out transition trajectories: does the non-equilibrium fitted model predict the observed transition tube better than the equilibrium model with the same occupancy? Fifth, in neuromorphic




FIG. 10. **Synthetic population readout of a rare transition.** Left: normalized activity of 80 synthetic units during one stochastic transition. Right: PCA recovers a curved low-dimensional trajectory. This illustrates how model-based transition geometry could be related to population neural recordings.

devices, current precision should improve with either increased integration time or increased dissipation, and the observed trade-off should lie above the TUR bound for the chosen current.

VI. DISCUSSION

The results support a simple conclusion: in non-reciprocal attractor dynamics, state stability, transition geometry and resolved thermodynamic cost are distinct but coupled properties. The benchmark has the same attractor positions for all χ , and its quasipotential barrier is unchanged by the rotational drift. Nevertheless, the transition paths, probability currents and entropy production change strongly with χ . A scalar landscape therefore captures only part of the computation. It describes where the system tends to reside, but not how it moves, how irreversible its switching routes are, or how much dissipation is required to make a current precise.

For neural state-space analysis, this distinction matters because directed effective connectivity is not merely a nuisance term to be averaged away. Non-reciprocal interactions can make state-space routes directional, support cyclic computation, accelerate transitions and separate escape paths from return paths. These benefits come with an irreversible signature: circulating currents and positive observed-scale entropy production. The claim is not that every directed connection exists to dissipate energy. The claim is that whenever a resolved neural dynamics is time-asymmetric, densities and fixed points are insufficient observables. Currents and paths must be measured or inferred.

For neuromorphic hardware, the same framework provides a more precise language for energy efficiency. A device that performs a state transition is not merely “low

power” or “high power”; it has a latency, an error distribution, a current observable and a dissipative cost. The TUR relates these quantities in a task-relevant way. A device can improve its position on the Pareto front by changing its architecture, but it cannot arbitrarily suppress current fluctuations at fixed dissipation and fixed time in a Markovian steady state. Thus, thermodynamic bounds should complement conventional energy-per-operation benchmarks.

The study also clarifies what is and is not demonstrated. The data are generated from specified coarse-grained models. They are not direct measurements of biological dissipation and should not be interpreted as total metabolic costs of cortical computation. Entropy production estimated from coarse-grained variables is generally a lower bound on microscopic entropy production, because hidden variables can carry hidden currents. A biological application would need to infer latent state variables, drift and diffusion from recordings; compute transition paths in the inferred state space; and test whether observed rare transitions cluster around predicted instanton tubes.

Several extensions are natural. First, diffusion noise should be replaced by jump-process large deviations for spiking systems when individual spikes or synaptic release events dominate the rare transition. Second, coloured noise should be treated by augmenting the state space with slow noise variables. Third, high-dimensional attractor networks should be represented as directed graphs whose edges carry action barriers, path asymmetries and entropy production estimates. Fourth, empirical uncertainty in fitted latent dynamics should be propagated to uncertainty in the predicted paths and barriers.

The broader lesson is that neural or neuromorphic attractor computation should be analysed as irreversible




FIG. 11. **Numerical diagnostics.** (A) Fitted action barrier as a function of the smallest noise strength retained in the rate-scaling fit. The dashed line marks the analytic barrier. (B) Path-node convergence of the geometric action. (C) Euler-Maruyama time-step convergence for the first-passage rate at $\epsilon = 0.08$. (D) Bootstrap distribution of the TUR product Q for the housekeeping current at $\chi = 0.8$; the dashed line marks the thermodynamic bound.




FIG. 12. **End-to-end computational pipeline.** The workflow starts from a noisy neural SDE and produces both rare-event predictions and thermodynamic diagnostics. The same structure can be used for rate networks, latent neural dynamics, neural-field models or neuromorphic state variables.

motion through state space, not only as occupation of states. A complete state-space analysis should identify the metastable states, compute the paths that connect them, predict the exponential rarity of those paths and quantify the resolved dissipative cost of reliable currents. Large-deviation theory and stochastic thermodynamics provide the mathematical tools for doing this in a unified way.

Relation to equilibrium landscape theories.

Equilibrium landscape theories remain valuable because they provide a compact description of metastability. In a reciprocal Hopfield-like network, memory states can often be described as minima of an effective energy function, and barrier heights give a first approximation to storage stability. The present framework does not reject that picture; it specifies when the picture is incomplete.

If the drift has a solenoidal component, then the stationary density and the transition current carry different information. Two systems can share the same density and the same attractor locations while generating different transition paths. In that case, the energy landscape is a marginal description of occupancy, not a full description of computation.

This distinction is not semantic. Consider two circuits with the same two stable memory states and the same local linear stability. In an equilibrium model, their spontaneous switching should be summarized by the same barrier and the same reversed relaxation path. In a non-equilibrium model, the escape branch can pass through a different region of state space from the return branch. Intermediate neural populations activated during a lapse or perceptual switch can therefore differ depending on the direction of the transition. A theory that predicts only the source state, target state and barrier height will miss that asymmetry. This is why the path itself should be treated as a biologically meaningful object.

There is also a practical inference consequence. Many data-driven neural landscape methods estimate an effective potential from state occupancy. Such methods can recover attractor basins when sampling is adequate, but they cannot by themselves determine whether detailed balance holds. To test irreversibility, one must estimate temporal currents, transition fluxes, cyclic probability flows or time-reversal asymmetries. A density-based method and a current-based method answer different questions. The former asks where the system spends time; the latter asks how it circulates through the states it visits.

Biological interpretation of non-reciprocity.

The non-reciprocity parameter χ in the benchmark is a simplified stand-in for several biological mechanisms. It may represent asymmetric effective connectivity, sequential excitation, adaptation-induced phase lags, delayed feedback, neuromodulatory drive or coupling to unobserved variables. In a real circuit these mechanisms need not combine into a literal two-dimensional rotation. What matters for the theory is that they create a component of the effective drift that is not expressible as the gradient of a scalar potential in the observed coordinates.

Such non-reciprocity can be computationally useful. It can make a circuit respond rapidly to a stimulus without destabilizing all memory states. It can route activity through a sequence of transient states. It can create hysteresis between the route into a state and the route out of it. It can also suppress some fluctuations while amplifying others. The thermodynamic result is that these benefits are accompanied by entropy production when they operate in a steady state. In other words, directed computation has a cost, but the cost is not merely waste; it is part of the physical implementation of directionality.

This interpretation aligns naturally with modern views of brain dynamics as a set of metastable and recurrently visited states rather than as a static code. If cognitive processing requires controlled movement among

metastable states, then the paths between states should be as important as the states themselves. A decision error may correspond to a rare escape through a particular transition tube. A perceptual switch may correspond to crossing a separatrix along a direction favoured by non-reciprocal drive. A working-memory lapse may correspond to a fluctuation that follows the least-action route out of a memory basin. These interpretations are hypotheses, but they are quantitative hypotheses: they can be tested by comparing predicted instanton tubes with observed transition ensembles.

Translation to empirical neural data.

An empirical version of this study would begin with state-space inference. Given population activity, one would infer low-dimensional coordinates using principal components, factor analysis, Gaussian-process latent-variable models, recurrent state-space models or neural stochastic differential equation models. The chosen coordinates should preserve the metastable states and transition events relevant to the task. Once the coordinates are defined, the drift and diffusion can be estimated from short-time conditional moments, maximum-likelihood fitting or Bayesian latent dynamics.

The next step is to infer the deterministic skeleton. Fixed points and slow manifolds can be identified in the fitted drift. Basins can be estimated by deterministic integration or by clustering long trajectories. Transition events in the data can then be aligned by source and target state. The minimum-action path predicted by the fitted model should be computed before looking at held-out transition events. The key validation is whether the held-out reactive trajectories concentrate around the predicted tube more strongly than around an equilibrium surrogate path.

The thermodynamic analysis requires additional care. Entropy production from coarse neural recordings is not total microscopic entropy production. It is the entropy production resolved by the inferred state variables, sampling rate and noise model. Hidden variables, unobserved inputs and measurement noise can all bias the estimate. Nevertheless, even a coarse-grained estimate can be useful if it is interpreted correctly. For example, task conditions that require faster decisions might show larger time-reversal asymmetry or larger resolved entropy production in the decision-state coordinates. Such comparisons would test relative changes in irreversibility rather than absolute metabolic cost.

Implications for neuromorphic hardware.

Neuromorphic hardware is often evaluated by energy per spike, energy per synaptic operation or total power. Those metrics are necessary but not sufficient for state-based computation. A memory-switching device must achieve a target transition probability, within a target time, with a target error tolerance. The relevant thermodynamic object is therefore a current: progress through a computation, circulation around a cycle or accumulated evidence along a decision coordinate. The TUR states

that the precision of such a current cannot be made arbitrarily high at fixed dissipation and fixed time.

In a physical device, the abstract Stratonovich current $\mathcal{J}_T = \int_0^T u(X_t) \circ dX_t$ should be mapped to the task variable whose fluctuations matter. In a charge-based circuit it can represent the signed integrated charge transported through a synaptic crossbar or through a readout line. In an event-driven spiking chip it can represent the net count of state-machine advances minus reversals, or the accumulated signed spike flux around a recurrent computational cycle. In a memristive or phase-based implementation it can represent the progress variable associated with conductance updates, oscillator phase slips or attractor-to-attractor switches. The housekeeping current used in the simulations is the model-space analogue of these device-level currents: it measures sustained irreversible progress generated by the non-reciprocal force, not merely occupation of a state.

This gives a different way to benchmark hardware. Instead of asking only how much energy a device consumes, one can ask where the device lies on a speed–precision–dissipation diagram. A slow device may appear efficient because it integrates for a long time. A fast device may appear accurate because it dissipates heavily. A meaningful comparison should fix the task, the observable current, the allowed latency and the required precision, then measure the dissipation required to meet those targets. The closer a device operates to the TUR bound, the more thermodynamically efficient it is for that observable and that task.

The benchmark also suggests an architectural principle. Symmetric components of a recurrent network can stabilize attractor basins, while controlled non-reciprocal components can route transitions. Maximizing non-reciprocity is not automatically optimal because it increases entropy production and can distort the state-space geometry. The design problem is to tune non-reciprocity so that transitions are fast and reliable enough without unnecessary dissipation. This is a more nuanced objective than simply lowering power or maximizing gain.

The design map in Fig. 9 makes this statement operational. For a target relative variance r and latency T , the TUR demands $\langle \sigma \rangle \geq 2/(rT)$. In a neuromorphic circuit this inequality can be read as a lower bound on the dissipative budget required by a specific current-like computation, independently of whether the physical carrier is charge, spin, phase, memristive conductance or spike timing. This is the point at which the theory speaks to the von Neumann bottleneck: reducing data movement and event energy is valuable, but an irreversible computation that must be fast and precise still has a thermodynamic floor. Similar scale-bridging questions are now central in active matter and information thermodynamics, where macroscopic fluctuation theory and entropy-production bounds connect microscopic driving to observable currents [28, 35–39]. The neural-network contribution of the present work is to translate that logic into attractor switching, memory reliability and hardware design coordi-

ates.

Limitations.

The main limitation is that the benchmarks remain deliberately controlled. The two-dimensional model has an analytic stationary density, and the $N = 50$ –500 extensions preserve the same normal-form barrier structure while testing high-dimensional optimization and seed-to-seed robustness. Biological neural systems have many degrees of freedom, structured noise, spiking events, adaptation, delays and hidden inputs. The diffusion approximation may fail when rare transitions are driven by discrete spikes or large synaptic events. In such cases the action functional must be replaced by a jump-process or hybrid-process large-deviation functional.

Another limitation is that entropy production is model-dependent under coarse-graining. If the observed variables omit a hidden cyclic process, then the observed current underestimates the true microscopic irreversibility. This is especially important for the brain, where ionic pumps, synaptic vesicle cycling and metabolic processes are not directly observed in neural activity recordings. The framework should therefore be used to quantify the thermodynamic cost of the resolved effective dynamics, not to claim the total energetic cost of tissue-level computation without additional biophysical modelling.

Finally, rare-event validation is statistically difficult. The direct Monte Carlo validation used here is possible because the barrier is moderate. In more stable networks, direct transition sampling may be infeasible. Enhanced sampling methods, uncertainty quantification and held-out predictive tests will be essential for empirical applications. The minimum-action path should be regarded as a hypothesis about the transition ensemble, not as a decorative curve drawn after the fact.

Future directions.

The most direct next step is to apply the framework to a fitted latent dynamical system rather than to a prescribed normal form. Large-scale neural recordings now make it possible to infer low-dimensional state spaces during decision-making, working memory, sleep transitions and spontaneous cortical activity. A decisive test would fit drift and diffusion on one subset of data, compute predicted minimum-action paths between metastable states, and then ask whether held-out rare transitions concentrate around those paths. The same analysis should compare a non-equilibrium model with an equilibrium surrogate that matches the stationary density. If the non-equilibrium model predicts direction-specific transition intermediates better, then broken detailed balance is not just a descriptive statistic; it is needed for prediction.

A second direction is to move beyond Gaussian diffusion. Spiking networks and synaptic-release models are often better described by jump processes. Their large-deviation actions are not quadratic and their optimal paths may differ from diffusion instantons even when the local variance is similar. This is especially impor-

tant when rare transitions are triggered by discrete events rather than by many small fluctuations. Extending the present workflow to jump processes would make it more directly applicable to spiking neural circuits and event-based neuromorphic hardware.

A third direction is to construct transition graphs in high-dimensional attractor networks. Instead of a single pair of memory states, one would infer a graph whose nodes are metastable neural states and whose directed edges are weighted by action barriers, transition rates, path asymmetries and entropy production. Such a graph would generalize the energy landscape into a non-equilibrium kinetic landscape. It would describe not only which states are stable but also which routes among them are dynamically preferred and thermodynamically costly.

Finally, the hardware implications should be tested in physical devices. A neuromorphic benchmark could require a device to switch between attractor states within a fixed time while maintaining a fixed current precision. The measured energy dissipation could then be compared with the TUR lower bound. This would turn the abstract speed–precision–dissipation relation into an engineering metric. The best devices would not be those with the lowest energy in isolation, but those closest to the thermodynamic bound for a specified computational current and latency.

Scope of the claims.

The claims of this paper are intentionally narrow and mechanistic. The work does not claim to have measured total biological dissipation. It does not claim that a two-dimensional normal form captures the anatomical or biophysical richness of cortical circuits. It claims that, once a noisy coarse-grained neural dynamics is specified, large-deviation theory and stochastic thermodynamics provide a coherent set of quantities that should be computed together: attractors, action barriers, transition paths, stationary currents, observed-scale entropy production and current-precision bounds.

This scope matters because it avoids a common problem in theoretical neuroscience: a model is often presented as a metaphor for the brain without a clear account of what the model can and cannot establish. Here the benchmark is used as a calibrated demonstration. Its analytic structure allows the action barrier to be checked, its probability current allows entropy production to be computed, and its stochastic simulations allow first-passage and TUR predictions to be validated. The model is simple by design, but the workflow is portable.

The most important conceptual claim is therefore not that the particular normal form is biologically complete. It is that rare neural transitions and thermodynamic precision should be analysed in the same state-space model. If future empirical work infers a different drift, a different diffusion tensor or a different noise law, the equations will change, but the logic remains: identify the rare paths, measure the currents and ask what cost is required for reliable computation.

VII. CONCLUSION

The physical conclusion is threefold. First, a scalar landscape is not enough. Two attractor systems can share the same stationary density and the same Freidlin–Wentzell quasipotential barrier while having different transition geometries. In the solvable normal form, the quasipotential remains U , but the non-reciprocal drift bends escape and relaxation branches into directed loops. Occupancy and barrier height therefore do not determine the route by which a rare switch occurs.

Second, geometric irreversibility is controlled by the solenoidal component of the drift and is quantitatively linked to entropy production. The path-asymmetry index obeys $\mathcal{A}_{\text{path}}(\chi) \sim \chi$, whereas $\langle \sigma \rangle \sim \chi^2/\epsilon$. Consequently, $\mathcal{A}_{\text{path}}^2 \propto \epsilon \langle \sigma \rangle$ in the small- χ regime. This identifies path irreversibility as a geometric square-root signature of dissipation in the benchmark model.

Third, current precision in non-reciprocal attractor dynamics is bounded by stochastic thermodynamics. Once a task-relevant current is specified, the TUR limits how small its relative fluctuations can be at fixed resolved entropy production and finite observation time. This gives a design constraint for neural and neuromorphic state-space computation: fast, reliable directed motion through attractor space has a thermodynamic precision cost.

The present results should not be interpreted as a measurement of total biological dissipation. They apply to resolved coarse-grained Markovian dynamics, and hidden variables can carry additional entropy production. In empirical neural or hardware data, the relevant claim is therefore about observed-scale irreversibility and lower bounds on microscopic dissipation, not about a complete accounting of all energetic processes.

Appendix A: Perturbative derivation of the path–dissipation scaling

This appendix gives the perturbative argument behind Eq. (12). Let $\gamma_\chi(s)$ denote the minimum-action transition curve at non-reciprocity χ , parametrized by $s \in [0, 1]$. For small χ we expand the curve around the equilibrium minimum-action path,

$$\gamma_\chi(s) = \gamma_0(s) + \chi\eta(s) + O(\chi^2). \quad (\text{A1})$$

The gradient component of the drift is invariant under time reversal, whereas the solenoidal component $\chi R \nabla U$ changes sign when the direction of the transition is reversed. Therefore the displacement between the forward transition loop and the time-reversed backward loop has the form

$$\Delta\gamma(s) = \gamma_\chi^f(s) - \gamma_\chi^b(1-s) = \chi\Delta\eta(s) + O(\chi^2). \quad (\text{A2})$$

The root-mean-square loop asymmetry is consequently

$$\mathcal{A}_{\text{path}}^2 = \int_0^1 |\Delta\gamma(s)|^2 ds = \chi^2 \int_0^1 |\Delta\eta(s)|^2 ds + O(\chi^4). \quad (\text{A3})$$

For the same normal form, the steady current is $J_{\text{ss}} = \chi R \nabla U \rho_{\text{ss}}$ and the entropy production rate is

$$\langle \sigma \rangle = \frac{\chi^2}{\epsilon} \langle |\nabla U|^2 \rangle_{\rho_{\text{ss}}}. \quad (\text{A4})$$

Combining Eqs. (A3) and (A4) gives

$$\mathcal{A}_{\text{path}}^2 = C_{\text{diss}} \epsilon \langle \sigma \rangle + O(\chi^4), \quad (\text{A5})$$

with

$$C_{\text{diss}} = \frac{\int_0^1 |\Delta\eta(s)|^2 ds}{\langle |\nabla U|^2 \rangle_{\rho_{\text{ss}}}}. \quad (\text{A6})$$

The proportionality constant is model- and metric-dependent, but the small- χ scaling follows from the fact that path displacement is linear in the solenoidal drift while entropy production is quadratic.

Appendix B: Freidlin–Wentzell action and Hamilton–Jacobi equation

For a diffusion $dX_t^\epsilon = F(X_t^\epsilon)dt + \sqrt{2\epsilon}\Sigma(X_t^\epsilon)dW_t$, the weak-noise path probability obeys the large-deviation principle in Eq. (37). The associated Hamiltonian is

$$H(x, p) = p^T F(x) + p^T A(x)p, \quad (\text{B1})$$

and the quasipotential solves the stationary Hamilton–Jacobi equation

$$H(x, \nabla V) = 0. \quad (\text{B2})$$

For $F = -\nabla U + \chi R \nabla U$ with $R^T = -R$ and $\nabla U^T R \nabla U = 0$, the solution remains $V = U$ even though the transition paths and steady currents change.

Appendix C: Geometric minimum-action method

The geometric action in Eq. (21) removes the unknown transition time by optimizing over curves rather than time-parametrized trajectories. Curves are discretized,

updated by projected descent, and redistributed to maintain arclength resolution. In high-dimensional scans the path is optimized in the full ambient state space; PCA is used only for diagnostic visualization and never to define the variational problem.

Appendix D: Numerical convergence and bootstrap procedures

Rate estimates use first-passage simulations over a finite noise window and are checked by refitting the barrier after removing the smallest-noise points. For TUR estimates, independent current segments are generated after burn-in, and confidence intervals for $\mathcal{Q}$ are obtained by bootstrap resampling of the segment currents. The high-dimensional Fig. 7 uses four random antisymmetric matrices for each N and reports mean values with 95% confidence intervals.

Appendix E: Additional high-dimensional network statistics

The high-dimensional normal form uses sparse antisymmetric matrices to keep the attractor structure stable while introducing non-reciprocal transverse channels. The $N = 100$ realization contains 148 nonzero coupling pairs; the same construction is scaled to $N = 50, 200, 500$. The full data table `highdim_scaling_statistics.csv` reports dimension, seed, non-reciprocity, coupling density, path asymmetry, entropy production and initial/final gMAM actions for every realization.

ACKNOWLEDGEMENTS

The author acknowledges the theoretical literatures on rare events, non-equilibrium statistical physics, stochastic thermodynamics and neural attractor dynamics, which provide the foundations for the present computational framework.

COMPETING INTERESTS

The author declares no competing interests.

-
- [1] D. Attwell and S. B. Laughlin, “An energy budget for signaling in the grey matter of the brain,” *Journal of Cerebral Blood Flow & Metabolism*, vol. 21, no. 10, pp. 1133–1145, 2001.
- [2] P. Lennie, “The cost of cortical computation,” *Current Biology*, vol. 13, no. 6, pp. 493–497, 2003.

- [3] C. W. Lynn, E. J. Cornblath, L. Papadopoulos, M. A. Bertolero, and D. S. Bassett, “Broken detailed balance and entropy production in the human brain,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 118, no. 47, p. e2109889118, 2021.

- [4] Y. Sanz Perl, C. Pallavicini, I. Perez Ipiña, A. Demertzi, V. Bonhomme, C. Martial, R. Panda, J. Annen, A. Ibañez, M. L. Kringelbach, S. Laureys, G. Deco, and E. Tagliazucchi, “Nonequilibrium brain dynamics as a signature of consciousness,” *Physical Review E*, vol. 104, p. 014411, 2021.
- [5] L. A. de la Fuente *et al.*, “Temporal irreversibility of neural dynamics as a signature of consciousness,” *Cerebral Cortex*, vol. 33, no. 5, pp. 1856–1865, 2023.
- [6] M. Gilson, A. El-Gazzar, Q. J. M. Huys, G. Deco, and K. Friston, “Entropy production of multivariate ornstein–uhlenbeck processes correlates with consciousness levels in the human brain,” *Physical Review E*, vol. 107, p. 024121, 2023.
- [7] S. Geli, E. Tagliazucchi, G. Deco, M. L. Kringelbach *et al.*, “Non-equilibrium whole-brain dynamics arise from pairwise interactions,” *Cell Reports Physical Science*, vol. 6, no. 3, p. 102464, 2025.
- [8] R. Nartallo-Kaluarachchi, M. L. Kringelbach, G. Deco *et al.*, “Nonequilibrium physics of brain dynamics,” *Physics Reports*, vol. 1152, pp. 1–43, 2026.
- [9] J. J. Hopfield, “Neural networks and physical systems with emergent collective computational abilities,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 79, no. 8, pp. 2554–2558, 1982.
- [10] D. J. Amit, H. Gutfreund, and H. Sompolinsky, “Spinglass models of neural networks,” *Physical Review A*, vol. 32, pp. 1007–1018, 1985.
- [11] X.-J. Wang, “Probabilistic decision making by slow reverberation in cortical circuits,” *Neuron*, vol. 36, no. 5, pp. 955–968, 2002.
- [12] M. I. Freidlin and A. D. Wentzell, *Random Perturbations of Dynamical Systems*, 3rd ed. Springer, 2012.
- [13] H. Touchette, “The large deviation approach to statistical mechanics,” *Physics Reports*, vol. 478, no. 1–3, pp. 1–69, 2009.
- [14] L. Bertini, A. De Sole, D. Gabrielli, G. Jona-Lasinio, and C. Landim, “Macroscopic fluctuation theory,” *Reviews of Modern Physics*, vol. 87, pp. 593–636, 2015.
- [15] R. Chetrite and H. Touchette, “Nonequilibrium markov processes conditioned on large deviations,” *Annales Henri Poincaré*, vol. 16, pp. 2005–2057, 2015.
- [16] W. E, W. Ren, and E. Vanden-Eijnden, “Minimum action method for the study of rare events,” *Communications on Pure and Applied Mathematics*, vol. 57, no. 5, pp. 637–656, 2004.
- [17] M. Heymann and E. Vanden-Eijnden, “The geometric minimum action method: A least action principle on the space of curves,” *Communications on Pure and Applied Mathematics*, vol. 61, no. 8, pp. 1052–1117, 2008.
- [18] U. Seifert, “Stochastic thermodynamics, fluctuation theorems and molecular machines,” *Reports on Progress in Physics*, vol. 75, no. 12, p. 126001, 2012.
- [19] A. C. Barato and U. Seifert, “Thermodynamic uncertainty relation for biomolecular processes,” *Physical Review Letters*, vol. 114, p. 158101, 2015.
- [20] T. R. Gingrich, J. M. Horowitz, N. Perunov, and J. L. England, “Dissipation bounds all steady-state current fluctuations,” *Physical Review Letters*, vol. 116, p. 120601, 2016.
- [21] J. M. Horowitz and T. R. Gingrich, “Thermodynamic uncertainty relations constrain non-equilibrium fluctuations,” *Nature Physics*, vol. 16, pp. 15–20, 2020.
- [22] C. Dieball and A. Godec, “Direct route to thermodynamic uncertainty relations and their saturation,” *Physical Review Letters*, vol. 130, p. 087101, 2023.
- [23] G. Lan, P. Sartori, S. Neumann, V. Sourjik, and Y. Tu, “The energy-speed-accuracy trade-off in sensory adaptation,” *Nature Physics*, vol. 8, pp. 422–428, 2012.
- [24] T. Van Vu and K. Saito, “Thermodynamic unification of optimal transport: Thermodynamic uncertainty relation, minimum dissipation, and thermodynamic speed limits,” *Physical Review X*, vol. 13, p. 011013, 2023.
- [25] E. Nitzan, A. Ghosal, and G. Bisker, “Universal bounds on entropy production inferred from observed statistics,” *Physical Review Research*, vol. 5, p. 043251, 2023.
- [26] D. H. Wolpert, A. Kolchinsky, J. A. Owen, and J. A. Grochow, “Is stochastic thermodynamics the key to understanding the energy costs of computation?” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 121, no. 45, p. e2321116121, 2024.
- [27] R. Sekizawa, Y. Hoshino, and T. Sagawa, “Decomposing thermodynamic dissipation of linear langevin systems via oscillatory modes,” *Physical Review X*, vol. 14, p. 041003, 2024.
- [28] R. Bebon, J. F. Robinson, and T. Speck, “Thermodynamics of active matter: Tracking dissipation across scales,” *Physical Review X*, vol. 15, p. 021050, 2025.
- [29] É. Roldán and J. M. R. Parrondo, “Estimating dissipation from single stationary trajectories,” *Physical Review Letters*, vol. 105, p. 150607, 2010.
- [30] M. Esposito and C. Van den Broeck, “Three detailed fluctuation theorems,” *Physical Review Letters*, vol. 104, p. 090601, 2010.
- [31] J. van der Meer, B. Ertel, and U. Seifert, “Thermodynamic inference in partially accessible markov networks: A unifying perspective from transition-based waiting time distributions,” *Physical Review X*, vol. 12, p. 031025, 2022.
- [32] M. Aguilera, J. Igarashi, and H. Shimazaki, “Nonequilibrium thermodynamics of the asymmetric sherrington–kirkpatrick model,” *Nature Communications*, vol. 14, p. 3685, 2023.
- [33] T. Ishihara and H. Shimazaki, “State-space kinetic ising model reveals task-dependent entropy flow in sparsely active nonequilibrium neuronal dynamics,” *Nature Communications*, vol. 16, p. 10852, 2025.
- [34] P. Hänggi, P. Talkner, and M. Borkovec, “Reaction-rate theory: fifty years after kramers,” *Reviews of Modern Physics*, vol. 62, pp. 251–341, 1990.
- [35] M. J. Bowick, N. Fakhri, M. C. Marchetti, and S. Ramaswamy, “Symmetry, thermodynamics, and topology in active matter,” *Physical Review X*, vol. 12, p. 010501, 2022.
- [36] T. Agranov, M. E. Cates, and R. L. Jack, “Macroscopic fluctuation theory and current fluctuations in active lattice gases,” *SciPost Physics*, vol. 14, p. 045, 2023.
- [37] R. Zakine and E. Vanden-Eijnden, “Minimum-action method for nonequilibrium phase transitions,” *Physical Review X*, vol. 13, p. 041044, 2023.
- [38] T. Markovich, É. Fodor, E. Tjhung, and M. E. Cates, “Thermodynamics of active field theories: Energetic cost of coupling to reservoirs,” *Physical Review X*, vol. 11, p. 021057, 2021.
- [39] S. Chen, D. S. Seara, A. Michaud, S. Kim, W. M. Beament, and M. P. Murrell, “Energy partitioning in the cell cortex,” *Nature Physics*, vol. 20, pp. 1824–1832, 2024.