

Identifying nonequilibrium degrees of freedom in high-dimensional stochastic systems

Catherine Ji,^{1,2,*} Ravin Raj^{1,*}, Benjamin Eysenbach^{2,†} and Gautam Reddy^{1,‡}

¹Department of Physics, Princeton University, Princeton, New Jersey 08544, USA

²Department of Computer Science, Princeton University, Princeton, New Jersey 08544, USA



(Received 15 August 2025; accepted 20 April 2026; published 8 June 2026)

Any coarse-grained description of a nonequilibrium system should faithfully represent its latent irreversible degrees of freedom. However, standard dimensionality reduction methods typically prioritize accurate reconstruction over physical relevance. Here, we introduce a model-free approach to identify irreversible degrees of freedom in stochastic systems that are in a nonequilibrium steady state. Our method leverages the insight that a black-box classifier, trained to differentiate between forward and time-reversed trajectories, implicitly estimates the local entropy production rate. By parametrizing this classifier as a quadratic form of learned state representations, we obtain nonlinear embeddings of high-dimensional state-space dynamics, which we term latent embeddings of nonequilibrium systems (LENS). LENS effectively identifies low-dimensional irreversible flows and provides a scalable, learning-based strategy for estimating entropy production rates directly from high-dimensional time series data.

DOI: [10.1103/1ncx-6cxk](https://doi.org/10.1103/1ncx-6cxk)

I. INTRODUCTION

Internal reactions and external forces drive active matter systems out of thermodynamic equilibrium, creating irreversible state space flows that describe the extent to which detailed balance is broken at the microscopic level [1–3]. In biological systems, these flows are typically driven by internal enzymatic processes and the action of molecular machines [4–8]. Working directly with detailed microscopic descriptions, however, quickly becomes intractable as the number of degrees of freedom in a system is increased, so any reasonable macroscopic description of a nonequilibrium system must appropriately discard nonessential information. Thus, a central problem of faithfully describing nonequilibrium systems lies in identifying relevant entropy-producing degrees of freedom.

Several prior works have developed methods for estimating entropy production. These methods include plug-in estimators and compression-based algorithms, which directly estimate the densities of forward and time-reversed dynamics [9–11]. Thermodynamic uncertainty relations enable the estimation of entropy production using fluctuations in finite-time estimates of the probability current [12–16]. However, the amount of data required for density or current estimation scales poorly with the dimensionality of the system. Recent works use statistical learning methods such as dimensionality reduction [17], neural networks [18–21], or score matching [22,23] to

directly infer entropy production, or the information bottleneck to learn low-dimensional dynamics [24] from which it may be possible to estimate entropy production. While these estimators have been shown to scale favorably with dimensionality, in many applications it is useful to identify which features of the dynamics contribute to irreversibility.

In this article, we introduce latent embeddings of nonequilibrium systems (LENS) to identify relevant irreversible degrees of freedom of a high-dimensional stochastic system in a nonequilibrium steady state (NESS). In contrast to commonly used dimensionality reduction methods [such as principal component analysis (PCA)] that prioritize reconstruction, LENS is a model-free approach that leverages contrastive representation learning [25] to directly extract entropy-producing degrees of freedom from high-dimensional time series data.

II. LENS CONSTRUCTION

For a Markovian system in an NESS, the entropy production rate (EPR) $\dot{S}$ is defined as the Kullback-Leibler (KL) divergence between the forward $p(\vec{x} \rightarrow \vec{x}')$ and the reverse dynamics $p(\vec{x}' \rightarrow \vec{x})$ [26]. Here, $\vec{x}$ and $\vec{x}'$ denote the states of the system at consecutive time steps t and $t + 1$, respectively, with $p(\vec{x} \rightarrow \vec{x}')$ as the probability of observing a transition from $\vec{x}$ to $\vec{x}'$ and vice versa for $p(\vec{x}' \rightarrow \vec{x})$. This KL divergence is given by

$$\begin{aligned}\dot{S} &= D_{\text{KL}}[p(\vec{x} \rightarrow \vec{x}') \| p(\vec{x}' \rightarrow \vec{x})] \\ &= \mathbb{E}_{p(\vec{x} \rightarrow \vec{x}')} \ln \left[\frac{p(\vec{x} \rightarrow \vec{x}')}{p(\vec{x}' \rightarrow \vec{x})} \right].\end{aligned}\quad (1)$$

To construct the LENS objective, we reformulate EPR estimation as learning a binary classifier $C_\theta(\vec{x}, \vec{x}')$ (parametrized by θ) that can distinguish forward paths from time-reversed paths. The classifier $C_\theta(\vec{x}, \vec{x}') = \sigma(S_\theta(\vec{x}, \vec{x}'))$, where

*These authors contributed equally to this work.

†Contact author: eyesenbach@princeton.edu

‡Contact author: greddy@princeton.edu




FIG. 1. LENS pipeline for learning representations. Observations $\vec{x}$ of high-dimensional stochastic dynamical systems are encoded using a learned function $\vec{\phi}_\theta$ that provides a low-dimensional description of the dynamics. The encoded observations are used to compute a logit $S(\vec{\phi}_t, \vec{\phi}_{t+1}; \mathbf{A}, \mathbf{B})$, where $\vec{\phi}_t \equiv \vec{\phi}(\vec{x}_t)$, $\vec{\phi}_{t+1} \equiv \vec{\phi}(\vec{x}_{t+1})$. This produces a prediction probability for whether the pair of observations $(\vec{x}_t, \vec{x}_{t+1})$ originate from the forward dynamics or the reverse dynamics.

$\sigma(z) = 1/(1 + e^{-z})$ is the logistic function, represents the probability that a randomly sampled pair of adjacent states $(\vec{x}, \vec{x}')$ originates from the forward dynamics ($\vec{x} \rightarrow \vec{x}'$) as opposed to the reverse dynamics ($\vec{x}' \rightarrow \vec{x}$), and $S_\theta(\vec{x}, \vec{x}')$ is the corresponding logit. The classifier is trained by maximizing the binary cross-entropy objective

$$J(\theta) = \mathbb{E}_{p(\vec{x} \rightarrow \vec{x}')} [\ln \sigma(S_\theta(\vec{x}, \vec{x}'))] + \mathbb{E}_{p(\vec{x}' \rightarrow \vec{x})} [\ln (1 - \sigma(S_\theta(\vec{x}, \vec{x}')))], \quad (2)$$

with respect to parameters θ . The expectations are computed over pairs of adjacent states sampled from a provided time series dataset. With infinite data, the ideal classifier that maximizes this objective (obtained from setting $\delta J / \delta S = 0$) is the optimal logit $S_{\text{opt}}(\vec{x}, \vec{x}') = \ln[p(\vec{x} \rightarrow \vec{x}') / p(\vec{x}' \rightarrow \vec{x})]$. Thus, a sufficiently expressive classifier can be used after convergence to estimate the EPR by averaging $S_\theta(\vec{x}, \vec{x}')$ over randomly sampled pairs $(\vec{x}, \vec{x}')$ from the dataset [Eq. (1)]. Maximizing the contrastive objective in Eq. (2) is equivalent (for continuous processes) to maximizing a lower bound on an f divergence [27], which relates LENS to an alternative learning-based approach for estimating the EPR [18].

Toward our goal of identifying irreversible degrees of freedom, we parametrize the logit $S_\theta(\vec{x}, \vec{x}')$ (Fig. 1) with an M -dimensional learned representation $\vec{\phi}_\theta(\vec{x})$, which is the output of a feedforward neural network. In particular, our LENS logit is parametrized as

$$S_\theta(\vec{x}, \vec{x}') = \vec{\phi}_\theta(\vec{x}')^T \mathbf{A}_\theta \vec{\phi}_\theta(\vec{x}) + \frac{1}{2} \vec{\phi}_\theta(\vec{x}')^T \mathbf{B}_\theta \vec{\phi}_\theta(\vec{x}') - \frac{1}{2} \vec{\phi}_\theta(\vec{x})^T \mathbf{B}_\theta \vec{\phi}_\theta(\vec{x}). \quad (3)$$

The learned matrix $\mathbf{A}_\theta$ has a block-diagonal form, where each block is a 2×2 skew-symmetric matrix. The block-diagonal form of $\mathbf{A}_\theta$ encourages LENS, without loss of generality, to learn the same representations across realizations of the stochastic system up to permutations of 2×2 subspaces and rotations within each subspace (except in situations where the eigenvalues of $\mathbf{A}_\theta$ are degenerate). $\mathbf{B}_\theta$ is symmetric, as any skew-symmetric component does not contribute to the logit [27]. Using deep nonlinear networks to parametrize $\vec{\phi}_\theta$ allows

for identifying features of arbitrary nonlinear stochastic systems that best approximate the EPR via $S_\theta(\vec{x}, \vec{x}')$.

This particular choice of $S_\theta(\vec{x}, \vec{x}')$ is motivated by two considerations. First, the form of $S_{\text{opt}}(\vec{x}, \vec{x}')$ requires that it is antisymmetric in its arguments, which to the lowest non-trivial order can be expressed as an antisymmetric quadratic form (with skew-symmetric $\mathbf{A}_\theta$) and the difference of two symmetric quadratic forms (with symmetric $\mathbf{B}_\theta$). This antisymmetric quadratic form with nonlinear $\vec{\phi}_\theta$ is universal: For a sufficiently large M , there exists a representation that approximates $S_{\text{opt}}(\vec{x}, \vec{x}')$ arbitrarily well [27]. Second, for linear stochastic systems, $S_{\text{opt}}(\vec{x}, \vec{x}')$ takes the form in Eq. (3) with linear $\vec{\phi}_\theta$ and the representations learned by LENS have a precise physical interpretation (discussed further below). We now demonstrate how LENS can be applied to various linear and nonlinear stochastic systems.

III. LINEAR STOCHASTIC DYNAMICS

To begin, we consider the paradigmatic example of N beads coupled by identical springs (of constant k) in a viscous medium [7, 28, 29]. Each bead is in contact with a local thermal reservoir as shown in Fig. 2(a). The system is maintained in a nonequilibrium steady state by imposing an inhomogeneous temperature profile along the chain of beads. Bead displacements $\vec{x}_t$ exhibit linear stochastic dynamics according to the overdamped Langevin equation

$$\vec{x}_{t+dt} = \vec{x}_t + \mathbf{G}\vec{x}_t dt + \mathbf{F}\vec{\xi}_t \sqrt{dt}. \quad (4)$$

The coupling matrix $\mathbf{G}$ has elements $G_{ij} = (-2\delta_{i,j} + \delta_{i,j+1} + \delta_{i,j-1})k/\gamma$ and each bead experiences uncorrelated thermal noise with scale $F_{ij} = \sqrt{2k_B T_i / \gamma} \delta_{i,j}$, where T_i is the temperature of the i th bead, γ is the viscous damping coefficient of the medium, and $\vec{\xi}_t$ is Gaussian white noise with zero mean and identity covariance.

We first apply LENS to such a system of $N = 2$ beads and maximize the LENS objective function over a time series of $T = 10^7$ observations with time step $dt = 10^{-2}$. In this setting, LENS accurately estimates the EPR to within 10% error across 6 orders of magnitude [Fig. 2(b)], regardless of whether the learned representations are parametrized as linear or nonlinear (see the Supplemental Material [27] for representation architecture).

To unpack the information contained in these linear representations, we apply LENS to high-dimensional systems of beads ($N = 16$) with four different temperature profiles for different M [Figs. 2(a) and 2(c)]. For $M = 2$, LENS identifies a circulating flow [Fig. 2(d)] that explains 25% of the total EPR. LENS compares favorably to two-dimensional representations produced by PCA [17] [Fig. 2(d)], time-lagged independent component analysis, and time-lagged canonical correlation analysis [27].

We also apply LENS to rendered systems of coupled beads, where the data are generated by convolving bead displacements $\vec{x}$ with a Gaussian point spread function to obtain “blurred” observations of N linearly arranged beads [top of Fig. 2(e)]. The relative EPR error is shown in Fig. 2(e) for various representation dimensions M , where we find that linear representations are able to capture the maximal possible EPR




FIG. 2. Learning nonequilibrium degrees of freedom in linear stochastic dynamical systems. (a) Coupled bead-spring model. (Top) N beads are coupled using identical springs, with each bead coupled to a local thermal reservoir at temperature T_i . (Bottom) Plots of the various temperature profiles under consideration. (b) Estimated entropy production rate (EPR) for $N = 2$ bead-spring model swept over T_c/T_h , linear and nonlinear representations of dimension $M = 16$. (c) Relative (fractional) EPR error as a function of linear representation dimension M for $N = 16$ bead-spring model for various temperature profiles. The corresponding theoretical predictions are obtained by optimizing the reduced LENS objective (5). (d) State space flow (white arrows) and stationary density (heatmap) shown in representation space for $M = 2$ for fully observed trajectories of the $N = 16$ bead-spring model with alternate temperature profile. (Inset) State space flow (blue arrows) and stationary density (heatmap) of top two principal components of the corresponding trajectory in coordinate space. (e) (Top) Representative one-dimensional (1D) render of Gaussian-convolved bead displacements presented as input to networks. (Bottom) Relative EPR error as a function of representation dimension for Gaussian-convolved trajectories of the $N = 16$ bead-spring model with alternate temperature profile, using linear and nonlinear representations. (f) State space flow (white arrows) and stationary density (heatmap) shown in representation space for $M = 2$ nonlinear representations of partially observed trajectories of the $N = 16$ bead-spring model with alternate temperature profile. All stationary density heatmaps use the same colorbar. All error bars represent the standard error of the mean (s.e.m.) over five random seed initializations.

available for a rank- M approximation to the linear dynamics. The benefit from learning a nonlinear approximation to the underlying linear stochastic dynamics is negligible, even when the observations are blurred with a nonlinearity. In Fig. 2(f), we similarly visualize the flow of nonlinear representations in the learned representational space for $M = 2$, which reveals cyclical irreversible flows similar to those in Fig. 2(d).

The representations learned by LENS for linear stochastic systems have a physical interpretation. To show this, we consider an N -dimensional linear system and M -dimensional linear representations $\tilde{\phi}_\theta(\vec{x}) = \mathbf{P}\vec{x}$. Maximizing the LENS objective Eq. (2) is equivalent to minimizing a reduced objective [27]

$$\mathcal{V}_M(\tilde{\mathbf{A}}) = 2\text{Tr}(\tilde{\mathbf{A}}\mathbf{G}\mathbf{C}) + \text{Tr}(\tilde{\mathbf{A}}^\top\mathbf{D}\tilde{\mathbf{A}}\mathbf{C}) \quad (5)$$

with respect to a rank M matrix $\tilde{\mathbf{A}} = \mathbf{P}^\top(\mathbf{A}_\theta + \mathbf{B}_\theta)\mathbf{P}$, where $\mathbf{D} = \mathbf{F}\mathbf{F}^\top/2$ is the diffusion tensor and $\mathbf{C}$ is the autocorrelation matrix. Computing the stationary points of this objective, we find that LENS implicitly learns an M -rank approximation

to the thermodynamic force matrix $\mathbf{A} = \mathbf{D}^{-1}\mathbf{G} + \mathbf{C}^{-1}$ and thereby the thermodynamic force field, $\mathbf{A}\vec{x}$. The EPR error decreases monotonically with M and converges to zero when $M = N$ [27].

IV. NONLINEAR STOCHASTIC DYNAMICS

Next, we consider a simplified model of molecular motors traversing cytoskeletal filaments by consuming adenosine triphosphate (ATP) [30,31]. Data obtained from imaging typically yield traces of motor locations along a one-dimensional track. In our model, a particle in a viscous medium is driven by a constant force h under the influence of a periodic potential U , as shown in Fig. 3(a). The particle's position x_t at time t is governed by

$$x_{t+1} = x_t + \left(h - \frac{dU}{dx}\right)dt + \sqrt{2D}d\xi_t, \quad (6)$$




FIG. 3. Learning entropy production in nonlinear stochastic dynamics. (a) Sketch of particle in a sinusoidal potential $U(x)$ subject to a constant driving force h . (b) Relative (fractional) EPR error as a function of representation dimension M for linear and nonlinear representations (with sine activation layer), with dimensionless drive strength $\kappa = hL/2\pi D = 1$. (c) Relative EPR error as a function of dimensionless drive strength κ for various representation dimensions M , for linear (with sine activation) and nonlinear representations. All error bars represent the standard error of the mean (s.e.m.) over five random seed initializations.

where the potential $U(x) = U_0 \cos(2\pi x/L)$ has amplitude U_0 and periodicity L . The system is nonlinear and yet admits an analytically tractable solution for the entropy production rate [27,32]. We simulate trajectories of $T = 10^7$ observations with time step $dt = 10^{-3}$ and estimate the EPR of these trajectories using LENS with both linear and nonlinear representations. Since the periodic nature of the simulated trajectories is known *a priori*, we include a sinusoidal activation function as an additional layer in our learned LENS encoder.

LENS effectively captures the entropy production in this nonlinear dynamical system with two representation dimensions [Fig. 3(b)] and estimates the EPR to within 3% across a range of drive strengths [Fig. 3(c)]. The relative EPR error remains relatively constant from $M = 2$ to $M = 16$ regardless of whether the representations are linear with sine activations or nonlinear. Lower-dimensional linear representations suffice for summarizing the EPR of linear systems, while the nonlinear dynamics in this system require higher dimensional representations to capture entropy production.

V. HIGH-DIMENSIONAL COMPLEX DYNAMICS

Finally, we apply LENS to two high-dimensional systems to demonstrate its applicability. We first use data generated from the complex Ginzburg-Landau (CGL) model [33,34],

governed by

$$\frac{\partial \tilde{A}(\vec{r}, t)}{\partial t} = \tilde{A}(\vec{r}, t) + (1 + i\alpha) \nabla^2 \tilde{A}(\vec{r}, t) - (1 + i\beta) |\tilde{A}(\vec{r}, t)|^2 \tilde{A}(\vec{r}, t), \quad (7)$$

where $\tilde{A}(\vec{r}, t) = |A(\vec{r}, t)| \exp(i\varphi(\vec{r}, t))$ is a complex field with spatiotemporal phase field $\varphi(\vec{r}, t)$ and amplitude $|A(\vec{r}, t)|$. Here, α and β are the linear and nonlinear dispersions of the medium, respectively. The CGL captures statistical features of biochemical wave patterns generated by the Rho-guanosine triphosphate (GTP) system on the membranes of starfish oocytes [35,36].

In two dimensions, Eq. (7) admits stable defect solutions for small values of α, β , with defects characterized by points $\vec{r}_0$ where the amplitude vanishes ($|A(\vec{r}_0)| = 0$) and the phase field $\varphi(\vec{r})$ winds up or down by 2π around the defect point. Increasing β leads to defect turbulence, where defects rapidly proliferate and the phase spirals are no longer easily distinguishable [33]. Note that CGL dynamics are deterministic, and the EPR estimates diverge logarithmically since every forward step will have no corresponding observation of its reversal [36]. However, the finite grid resolution still leads to a finite EPR estimate, and it is possible to find representations that are indicative of the transition from the stable to chaotic regimes.

Fixing $\alpha = -0.8$, we simulate trajectories of length $T = 10^5$ observations with time step $dt = 10^{-3}$ and estimate the EPR using LENS. The phase fields $\varphi(\vec{r}, t)$ for all observations are fed into a convolutional neural network [27]. The EPR estimates computed at various values of β across the defect turbulence transition are shown in Fig. 4(a) for different M . The transition is clearly visible for low M , where the estimated EPR drops sharply past the transition line. The larger EPR estimate for larger M in the chaotic regime correlates with higher values of the LENS objective (lower binary cross-entropy loss), indicating that the larger M representations are better able to extract the presence of the forward-only dynamics in the chaotic regime. This is also corroborated by the state space flows and local (single time step) EPR estimates for $M = 2$ representations [Fig. 4(b)]. While the flows in both regimes are largely cyclical and closed, the local EPR predicted in the stable regime is relatively constant and entirely confined to the thin ring where the representation flows lie. In contrast, the learned local EPR map in the chaotic regime is characterized by small positive local EPR along the ring of representational flows with spurious negative local EPR patches within the ring, indicating that the learned representations are unable to capture the deterministic cyclic dynamics with small M (see the Supplemental Material [37]).

Finally, we apply LENS to videos of active microtubule dynamics obtained at varying concentrations of ATP [38]. The estimated EPR is shown for various representation dimensions M in Fig. 4(c), revealing that the EPR generally increases with ATP concentration. Four or more degrees of freedom are necessary to capture irreversibility in these data. We used GradCAM [39] to visualize the features that contribute most to the EPR between pairs of frames, with a representative image shown in Fig. 4(d) for the GradCAM attention map of the difference between consecutive frames $(x_{t+1} - x_t)/2$,




FIG. 4. Learning low-dimensional representations of complex nonlinear dynamics. (a) (Top) Representative images of CGL phase field in the stable (left) and chaotic (right) regimes. (Bottom) Estimated EPR as a function of nonlinear dispersion β across the defect turbulence transition (gray dotted line), for various representation dimensions M with nonlinear representations $\tilde{\phi}_\theta$. Defect turbulence transition line from Chaté and Manneville [33]. (b) Representative state space flows (arrows) and local EPR estimates (heatmaps) in representation space for $M = 2$ in the stable ($\beta = 0$, top) and chaotic ($\beta = 2.0$, bottom) regimes. Both heatmaps share the same colorbar scale for the estimated local (single time step) EPR. (c) EPR estimate for microtubule dynamics obtained from LENS for varying concentrations of ATP. (d) Representative images of active microtubule dynamics shown for $24 \mu\text{M}$ ATP. Underlying averaged phase contrast image of a representative pair of frames $(x_t + x_{t+1})/2$ is shown in gray scale, with a heatmap overlay displaying the regions of nonzero GradCAM attention between the difference of frames $(x_{t+1} - x_t)/2$. All error bars represent the standard error of the mean (s.e.m.) over five random seed initializations.

overlaid atop the average of the two frames $(x_t + x_{t+1})/2$. The highly attended regions identify the specific portions of the image that LENS identifies as being relevant to entropy production.

VI. DISCUSSION AND OUTLOOK

Here, we propose a method for identifying relevant irreversible degrees of freedom in high-dimensional NESS systems. Importantly, our method is model-free, conceptually simple, and bypasses direct density or current estimation. In line with recent work to estimate EPR [18], our method uses black-box neural networks but differs in that it identifies which aspects of the dynamics are irreversible. The quadratic form of the LENS logit suggests that LENS will be particularly insightful in situations where high-dimensional, nonlinear dynamics are generated by a small set of latent, irreversible features with linear dynamics [40]. A mathematical and empirical characterization of LENS in such scenarios is an interesting direction for future work. We note, however, that the converse is not necessarily true. That is, while the universality of the quadratic form ensures that the log probability ratio of the forward and reverse dynamics can be approximated arbitrarily well (with enough data), this does not assume or imply that a system's nonlinear dynamics can always be transformed to a globally linear representation.

Moreover, LENS assumes a Markovian state representation; hidden slow variables and other non-Markovian effects could confound EPR estimates and may require augmenting the state with learned latent variables from time series data [41]. We anticipate that our method will open up avenues to analyze biological active matter systems and other complex dynamical systems with inaccessible internal states.

ACKNOWLEDGMENTS

The authors thank Linnea Lemma, Benjamin Figuerado, and Joshua Shaevitz for feedback and providing experimental active nematics data. The authors thank Trevor GrandPre, Junang Li, Bezia Lemma, and William Bialek for insightful discussions on the subject, as well as Hugues Chaté and Paul Manneville for providing supplementary data for the CGL phase diagram. This material is based upon work supported by the National Science Foundation (Graduate Research Fellowship Program or CISE Graduate Fellowship) under Grant No. DGE-2444107. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. The simulations presented in this article were performed on computational resources managed and supported by Princeton Research Computing, a consortium of groups including the

Princeton Institute for Computational Science and Engineering (PICSciE) and the Office of Information Technology's High Performance Computing Center and Visualization Laboratory at Princeton University. G.R. was partially supported by a joint research agreement between NTT Research Inc and Princeton University.

C.J. and R.R. contributed equally to this work. C.J. led project implementation and experiments; R.R. and G.R. performed theory calculations; R.R., G.R., and B.E. formulated proofs; R.R. and C.J. conducted data analysis; R.R. led paper writing. B.E. and G.R. equally advised this work. All authors

designed and performed research. All authors contributed to writing and revising the paper.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

DATA AVAILABILITY

The data that support the findings of this article are openly available [42].

-
- [1] C. Battle, C. P. Broedersz, N. Fakhri, V. F. Geyer, J. Howard, C. F. Schmidt, and F. C. MacKintosh, Broken detailed balance at mesoscopic scales in active biological systems, *Science* **352**, 604 (2016).
 - [2] J. Gladrow, C. P. Broedersz, and C. F. Schmidt, Nonequilibrium dynamics of probe filaments in actin-myosin networks, *Phys. Rev. E* **96**, 022408 (2017).
 - [3] C. W. Lynn, C. M. Holmes, W. Bialek, and D. J. Schwab, Emergence of local irreversibility in complex interacting systems, *Phys. Rev. E* **106**, 034102 (2022).
 - [4] U. Seifert, Stochastic thermodynamics, fluctuation theorems and molecular machines, *Rep. Prog. Phys.* **75**, 126001 (2012).
 - [5] H. Qian, S. Kjelstrup, A. B. Kolomeisky, and D. Bedeaux, Entropy production in mesoscopic stochastic thermodynamics: Nonequilibrium kinetic cycles driven by chemical potentials, temperatures, and mechanical forces, *J. Phys.: Condens. Matter* **28**, 153004 (2016).
 - [6] F. S. Gnesotto, F. Mura, J. Gladrow, and C. P. Broedersz, Broken detailed balance and non-equilibrium dynamics in living systems: A review, *Rep. Prog. Phys.* **81**, 066601 (2018).
 - [7] F. Mura, G. Gradziuk, and C. P. Broedersz, Nonequilibrium scaling behavior in driven soft biological assemblies, *Phys. Rev. Lett.* **121**, 038002 (2018).
 - [8] T. H. Tan, G. A. Watson, Y.-C. Chao, J. Li, T. R. Gingrich, J. M. Horowitz, and N. Fakhri, Scale-dependent irreversibility in living matter, [arXiv:2107.05701](https://arxiv.org/abs/2107.05701).
 - [9] Q. Wang, S. Kulkarni, and S. Verdu, Divergence estimation of continuous distributions based on data-dependent partitions, *IEEE Trans. Inf. Theory* **51**, 3064 (2005).
 - [10] J. Ziv and N. Merhav, A measure of relative entropy between individual sequences with application to universal classification, *IEEE Trans. Inf. Theory* **39**, 1270 (1993).
 - [11] R. Roldán and J. M. R. Parrondo, Entropy production and Kullback-Leibler divergence between stationary trajectories of discrete systems, *Phys. Rev. E* **85**, 031129 (2012).
 - [12] A. C. Barato and U. Seifert, Thermodynamic uncertainty relation for biomolecular processes, *Phys. Rev. Lett.* **114**, 158101 (2015).
 - [13] T. R. Gingrich, J. M. Horowitz, N. Perunov, and J. L. England, Dissipation bounds all steady-state current fluctuations, *Phys. Rev. Lett.* **116**, 120601 (2016).
 - [14] J. M. Horowitz and T. R. Gingrich, Thermodynamic uncertainty relations constrain non-equilibrium fluctuations, *Nat. Phys.* **16**, 15 (2020).
 - [15] J. Li, J. M. Horowitz, T. R. Gingrich, and N. Fakhri, Quantifying dissipation using fluctuating currents, *Nat. Commun.* **10**, 1666 (2019).
 - [16] E. Roldán, J. Barral, P. Martin, J. M. R. Parrondo, and F. Jülicher, Quantifying entropy production in active fluctuations of the hair-cell bundle from time irreversibility and uncertainty relations, *New J. Phys.* **23**, 083013 (2021).
 - [17] F. S. Gnesotto, G. Gradziuk, P. Ronceray, and C. P. Broedersz, Learning the non-equilibrium dynamics of Brownian movies, *Nat. Commun.* **11**, 5378 (2020).
 - [18] D.-K. Kim, Y. Bae, S. Lee, and H. Jeong, Learning entropy production via neural networks, *Phys. Rev. Lett.* **125**, 140604 (2020).
 - [19] Y. Bae, D.-K. Kim, and H. Jeong, Inferring dissipation maps from videos using convolutional neural networks, *Phys. Rev. Res.* **4**, 033094 (2022).
 - [20] J. Lyu, K. J. Ray, and J. P. Crutchfield, Learning entropy production from underdamped Langevin trajectories, *Phys. Rev. E* **110**, 064151 (2024).
 - [21] M. Vodret, C. Pacini, and C. Bongiorno, Functional decomposition and estimation of irreversibility in time series via machine learning, *Phys. Rev. E* **110**, 064310 (2024).
 - [22] N. M. Boffi and E. Vanden-Eijnden, Model-free learning of probability flows: Elucidating the nonequilibrium dynamics of flocking, *Phys. Rev. Lett.* **135**, 238301 (2025).
 - [23] N. M. Boffi and E. Vanden-Eijnden, Deep learning probability flows and entropy production rates in active matter, *Proc. Natl. Acad. Sci. USA* **121**, e2318106121 (2024).
 - [24] M. S. Schmitt, M. Koch-Janusz, M. Fruchart, D. S. Seara, M. Rust, and V. Vitelli, Information theory for data-driven model reduction in physics and biology, [arXiv:2312.06608](https://arxiv.org/abs/2312.06608).
 - [25] Z. Ma and M. Collins, Noise contrastive estimation and negative sampling for conditional models: Consistency and statistical efficiency, in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, edited by E. Riloff, D. Chiang, J. Hockenmaier, J. Tsujii (Association for Computational Linguistics, Brussels, Belgium, 2018), pp. 3698–3707.
 - [26] R. Kawai, J. M. R. Parrondo, and C. V. den Broeck, Dissipation: The phase-space perspective, *Phys. Rev. Lett.* **98**, 080602 (2007).
 - [27] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/1ncx-6cxk> for model and optimization details, and sufficiency proofs for the LENS formalism.

- [28] Z. Rieder, J. L. Lebowitz, and E. Lieb, Properties of a harmonic crystal in a stationary nonequilibrium state, *J. Math. Phys.* **8**, 1073 (1967).
- [29] F. Bonetto, J. L. Lebowitz, and J. Lukkarinen, Fourier's law for a harmonic crystal with self-consistent stochastic reservoirs, *J. Stat. Phys.* **116**, 783 (2004).
- [30] A. Yildiz, J. N. Forkey, S. A. McKinney, T. Ha, Y. E. Goldman, and P. R. Selvin, Myosin V walks hand-over-hand: Single fluorophore imaging with 1.5-nm localization, *Science* **300**, 2061 (2003).
- [31] A. Yildiz, M. Tomishige, R. D. Vale, and P. R. Selvin, Kinesin walks hand-over-hand, *Science* **303**, 676 (2004).
- [32] H. Risken, *The Fokker-Planck Equation: Methods of Solution and Applications*, edited by H. Haken, Springer Series in Synergetics Vol.18 (Springer, Berlin, 1996).
- [33] H. Chat   and P. Manneville, Phase diagram of the two-dimensional complex Ginzburg-Landau equation, *Physica A* **224**, 348 (1996).
- [34] I. S. Aranson and L. Kramer, The world of the complex Ginzburg-Landau equation, *Rev. Mod. Phys.* **74**, 99 (2002).
- [35] T. H. Tan, J. Liu, P. W. Miller, M. Tekant, J. Dunkel, and N. Fakhri, Topological turbulence in the membrane of a living cell, *Nat. Phys.* **16**, 657 (2020).
- [36] J. Li, C.-W. J. Liu, M. Szurek, and N. Fakhri, Measuring irreversibility from learned representations of biological patterns, *PRX Life* **2**, 033013 (2024).
- [37] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/Incx-6cxc> for Movies M1 and M2.
- [38] C. Joshi, S. Ray, L. M. Lemma, M. Varghese, G. Sharp, Z. Dogic, A. Baskaran, and M. F. Hagan, Data-driven discovery of active nematic hydrodynamics, *Phys. Rev. Lett.* **129**, 258001 (2022).
- [39] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, Grad-CAM: Visual explanations from deep networks via gradient-based localization, in *Proceedings of the 2017 IEEE International Conference on Computer Vision (ICCV)* (IEEE, Piscataway, NJ, 2017), pp. 618–626.
- [40] M. Watter, J. T. Springenberg, J. Boedeker, and M. Riedmiller, Embed to control: A locally linear latent dynamics model for control from raw images, in *Proceedings of the 29th International Conference on Neural Information Processing Systems*, edited by C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett (MIT Press, Cambridge, MA, 2015), Vol. 2, pp.2746–2754.
- [41] A. v. d. Oord, Y. Li, and O. Vinyals, Representation learning with contrastive predictive coding, [arXiv:1807.03748](https://arxiv.org/abs/1807.03748).
- [42] Code repository for LENS, Catherine Ji, 2025, <https://github.com/cj7280/LENS>.