

Constructing custom thermodynamics using deep learning

Received: 13 August 2023

Accepted: 30 November 2023

Published online: 29 December 2023

 Check for updates

Xiaoli Chen^{1,2,7}, Beatrice W. Soh^{3,7}, Zi-En Ooi³, Eleonore Vissol-Gaudin⁴, Haijun Yu^{5,6}, Kostya S. Novoselov²✉, Kedar Hippalgaonkar^{2,3,4}✉ & Qianxiao Li^{1,2}✉

One of the most exciting applications of artificial intelligence is automated scientific discovery based on previously amassed data, coupled with restrictions provided by known physical principles, including symmetries and conservation laws. Such automated hypothesis creation and verification can assist scientists in studying complex phenomena, where traditional physical intuition may fail. Here we develop a platform based on a generalized Onsager principle to learn macroscopic dynamical descriptions of arbitrary stochastic dissipative systems directly from observations of their microscopic trajectories. Our method simultaneously constructs reduced thermodynamic coordinates and interprets the dynamics on these coordinates. We demonstrate its effectiveness by studying theoretically and validating experimentally the stretching of long polymer chains in an externally applied field. Specifically, we learn three interpretable thermodynamic coordinates and build a dynamical landscape of polymer stretching, including the identification of stable and transition states and the control of the stretching rate. Our general methodology can be used to address a wide range of scientific and technological applications.

The modern scientific method adopts a universal approach that ensures stable and non-conflicting progression of our understanding of nature: new theories need to be hypothesized and tested on previously amassed data, be compatible with the basic scientific principles and be verifiable by experiments. Unfortunately, there is no general algorithmic recipe to do so in complex systems to facilitate discovery. Hence, up to now, only the most basic physical phenomena—often static, in equilibrium—are described by an intuitive set of equations. Many dynamic, non-equilibrium phenomena, which determine functionality in biology, soft-condensed matter and chemistry, are instead described via very approximate, empirical laws. The advancement of artificial intelligence and machine learning gives rise to the possibility of a data-driven solution to this challenge^{1,2}.

In this paper, we develop Stochastic OnsagerNet (S-OnsagerNet), an artificial intelligence platform that can discover an interpretable and closed thermodynamic description of an arbitrary stochastic dissipative dynamical system directly from observations of microscopic trajectories. There are essentially two types of approach to understand and predict the behavior of dynamical processes from data—unstructured and structured. Unstructured approaches parameterize dynamical equations by a generic set of building blocks, such as fixed polynomials³, trainable feature maps or kernels^{4,5}, and determines the associated parameters that best fit the observations. Physical insights can be incorporated as regularizers in the fitting process⁶. Their generality comes at a cost of long-time predictive accuracy, stability⁷ and, more importantly, interpretability. This is addressed

¹Department of Mathematics, National University of Singapore, Singapore, Singapore. ²Institute for Functional Intelligent Materials, National University of Singapore, Singapore, Singapore. ³Institute of Materials Research and Engineering, A*STAR (Agency for Science), Singapore, Singapore. ⁴School of Materials Science and Engineering, Nanyang Technological University, Singapore, Singapore. ⁵LSEC and ICMSEC, Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing, China. ⁶School of Mathematical Sciences, University of Chinese Academy of Sciences, Beijing, China. ⁷These authors contributed equally: Xiaoli Chen, Beatrice W. Soh. ✉e-mail: kostya@nus.edu.sg; kedar@ntu.edu.sg; qianxiao@nus.edu.sg

by the class of structured approaches, where physical insights directly guide the design of model architectures. Our approach belongs to this latter category. Previous work in this direction includes models based on Hamiltonian or symplectic dynamics^{8,9}, Poisson systems¹⁰ and quasi-potentials¹¹. However, so far, there lacks a general structured approach to model dissipative, non-equilibrium and noisy dynamics that often arise in soft matter, biophysics and other applications. Our methodology based on the classical Onsager principle^{12,13} is tailored to such problems.

Macroscopic thermodynamic descriptions of physical systems are highly sought after for the insights they provide. A prototypical example is the ideal-gas law as a macroscopic description of non-interacting gas systems. These systems guide the design of verification experiments and provide principled ways to manipulate macroscopic behavior. For a general complex dynamical system, however, constructing an intuitive thermodynamic description that enables subsequent analysis and control is a daunting task. Our approach addresses this challenge as follows. For a given microscopic dynamics, we learn a macroscopic thermodynamic description via the simultaneous construction of low-dimensional closure coordinates—ensured to be partially interpretable—and a time evolution law on these coordinates. Unlike general artificial intelligence approaches, our platform intrinsically limits the search to physically relevant evolution laws. In particular, we ensure compatibility with existing scientific knowledge by constructing our neural network architecture based on a generalized Onsager principle.

We demonstrate our method by learning the stretching dynamics of polymer chains containing up to 900 degrees of freedom, condensing it into a thermodynamic description involving only three macroscopic coordinates that governs polymer stretching dynamics in both computational and experimental data. We build an energy landscape of the macroscopic evolution, revealing the presence of stable and transition states. This can be viewed as a dynamic equation of state. Mastering such an equation allows the design of verification computational experiments, including the interpretation of the thermodynamic coordinates and the control of the stretching rate of the polymers. We extend this further to conduct single-molecule DNA stretching experiments and show that our thermodynamic description can be used to distinguish fast and slow stretching polymers, much beyond current human-labeling capabilities. Furthermore, the predicted fluctuation correlations derived from the free-energy landscape agree with experimental data.

Constructing low-dimensional physical models from high-dimensional dynamical data is an active area of research. Data-driven modeling of dynamical processes based on the Onsager's principle was proposed in ref. 7 to study Rayleigh–Bénard convection. Reference¹⁴ combined encoder–decoders and manifold learning to construct latent dynamical models directly from video data, including those of reaction-diffusion processes and pendulum motions. Here we make several advancements in terms of methodology and applications. First, unlike the deterministic models considered previously, we explicitly capture stochastic fluctuations—an important element of non-equilibrium processes at finite temperatures. In fact, the non-trivial heterogeneity of polymer stretching dynamics studied in this paper is directly caused by thermal fluctuations. Developing our method in the stochastic setting requires extensions of model reduction theory and training algorithms (see ‘Theoretical results and model implementation’ in Methods). Second, and more importantly, we go beyond dimensionality reduction^{7,14,15} and solve a closure problem: given a priori fixed macroscopic variables of interest (for example, polymer extension), we construct both the closure coordinates sufficient to govern the evolution of these macroscopic variables, and the dynamics that describes this evolution. Compared with more flexible parameterizations of reduced dynamics¹⁴, our approach inherently limits the search space to those satisfying a generalized Onsager principle, which sacrifices complete generality

but facilitates physical interpretation of the closure coordinates and the dynamical landscape.

Results

We now describe our approach. The most complete description of a complex, multi-component system is the coordinates of all the components as a function of time t ($\mathbf{X}(t)$). For the ideal gas, it would be the positions and momenta of all molecules and for a magnetic system—the spin state of each atom. An alternative to this expensive microscopic modeling approach is a thermodynamic one, where the full description is replaced by some macroscopic coordinates ($\mathbf{Z}^*(t)$, with dimensionality much smaller than $\mathbf{X}(t)$). This can be the pressure of an ideal gas or the magnetization of a magnetic system. The thermodynamic approach links these macroscopic coordinates to other macroscopic coordinates or closure variables ($\hat{\mathbf{Z}}(t)$), and external parameters (volume and temperature for ideal gas, magnetic field and temperature for magnetic systems) via an equation of state.

We propose a generic approach of building such custom thermodynamics for an arbitrary stochastic, dissipative dynamical system from data. We are given macroscopic coordinates \mathbf{Z}^* whose dynamics we wish to model. For polymer dynamics this can be a single variable—the extension of the polymer chain (Fig. 1). Then, we learn a set of closure variables $\hat{\mathbf{Z}}(t)$ and simultaneously an evolution law on the combined thermodynamic coordinates $\mathbf{Z}(t) = (\mathbf{Z}^*(t), \hat{\mathbf{Z}}(t))$ that enables scientific understanding, experimental verification and control. The evolution equation is a generalization of the classical Onsager principle^{12,13} that has been successfully applied to model a variety of non-equilibrium phenomena, including phase separation kinetics, gel dynamics and molecular modeling^{16,17}. It posits a time evolution law

$$\dot{\mathbf{Z}}(t) = -M\nabla V(\mathbf{Z}(t)), \quad (1)$$

for a given set of coordinates $\mathbf{Z}(t)$, where the dot denotes time derivative, the symmetric positive semi-definite matrix M models energy dissipation and V is a generalized potential. A limitation of the Onsager principle is its inability to capture dynamics far from equilibrium, or with substantial stochastic behavior. To this end, we propose an extension in the form of a generalized stochastic Onsager principle

$$\dot{\mathbf{Z}}(t) = -[M(\mathbf{Z}(t)) + W(\mathbf{Z}(t))]\nabla V(\mathbf{Z}(t)) + \sigma(\mathbf{Z}(t))\dot{\mathbf{B}}(t), \quad (2)$$

where $M(\cdot)$ and $W(\cdot)$ are now functions of the reduced coordinates \mathbf{Z} that output $d \times d$ matrices. $M(\cdot)$ is symmetric positive semi-definite to conform to stability requirements and Onsager's reciprocal relations¹², while $W(\cdot)$ is anti-symmetric and models conservative forces. The diffusion matrix $\sigma(\cdot)$ together with the white noise process $\dot{\mathbf{B}}(t)$ models the thermal fluctuations in the system. Equation (2) forms the basis of our dynamical model in reduced coordinates. We note that alternative generalizations of the Onsager principle have been proposed using large deviations theory¹⁸, but their forms are more complex and hence less amenable to computations. It can be shown that our model has long-time stability through energy dissipation up to the order of thermal fluctuations (Theorem 2) and the flexibility to represent many physical stochastic processes, including Langevin and generalized Poisson dynamics (see ‘Theoretical results’ in Methods). Our method departs from classical modeling paradigms, where the unknown equation parameters are few and can be fitted from few experiments. Instead, the unknowns here are functions M , W , V and σ . We leverage machine learning and represent these functions as trainable deep neural networks, while preserving the required physical constraints (for example, symmetric positive definiteness of M). Simultaneously, we generate a set of the closure coordinates by another deep neural network, which combines approximation flexibility of residual networks and approximate feature orthogonality through principal component analysis (PCA). This is to be contrasted with generic coarse-graining

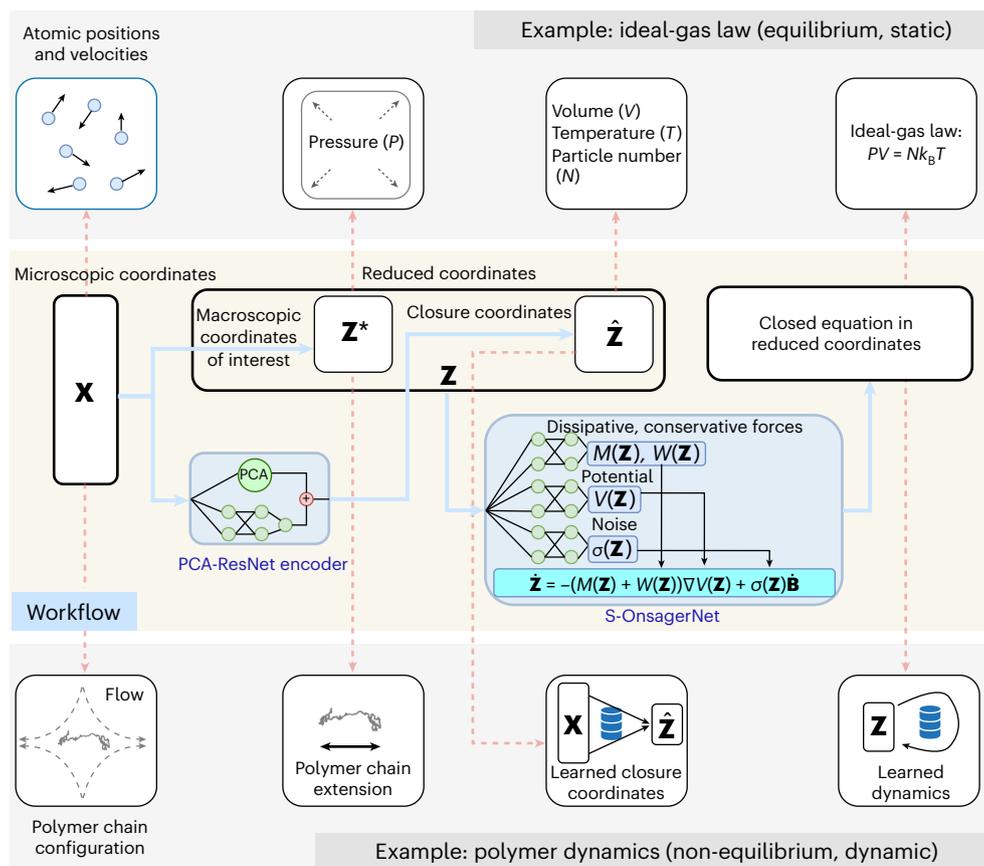


Fig. 1 | Overall workflow of the proposed approach. Given a complex system described by \mathbf{X} , the goal is to model the behavior of macroscopic coordinates of interest \mathbf{Z}^* . We construct closure coordinates $\hat{\mathbf{Z}}$ and closed (dynamical) equation on the combined reduced coordinates $\mathbf{Z} = (\mathbf{Z}^*, \hat{\mathbf{Z}})$. The classical ideal-gas law is an illustration of this process (top). For general non-equilibrium, dynamic systems (bottom), carrying out this workflow from theoretical analysis is

challenging. Our machine learning method (middle) addresses this by simultaneously constructing the closure coordinates using PCA-ResNet (Methods), and governing equations on reduced coordinates using the S-OnsagerNet with drift term $-(M(\mathbf{Z}) + W(\mathbf{Z}))\nabla V(\mathbf{Z})$ and noise term $\sigma(\mathbf{Z})$ (see equation (2)).

methods based on volume averages¹⁹, in that we seek a very small set of closure coordinates that are sufficient to describe the motion of the macroscopic states of interest. Our learning-based approach to discover hidden coordinates shares some similarity with recently proposed machine learning-based coarse-graining methods in molecular simulations^{15,20} but here we work with a closure problem, thus we may end up with macroscopic dynamics of substantially lower dimensions. We perform end-to-end training of the combined architecture on large-scale microscopic trajectory data to simultaneously learn the reduced coordinates and their dynamics. The overall workflow for creating custom thermodynamics, which we call S-OnsagerNet, is summarized in Fig. 1. Detailed network architectures and training algorithms are given in Methods.

Training and prediction of polymer stretching dynamics

We first demonstrate our approach by modeling the temporal evolution of polymer extension under elongational forces, which has long been of interest to the polymer physics community^{21–24}. Hallmark experimental^{25,26} and computational studies^{27–29} in elongational rheology of dilute polymers have examined the deformation of single DNA molecules in planar elongational flows and revealed the highly heterogeneous stretching dynamics among identical polymer chains. Due to the complex interactions within and stochastic nature of the system, it is challenging to identify macroscopic descriptors of the polymer chain (closure coordinates) and governing equations on these descriptors that are sufficient to determine the outcome of the stretching

dynamics. Yet, such a thermodynamic description is essential for understanding the origins of unfolding heterogeneity and paves the way to make desired modifications to the unfolding dynamics. Thus, our data-driven method offers a promising alternative to achieve this goal.

We simulate polymer chain stretching in a planar elongational flow. The polymer chain consists of 300 coarse-grained beads connected by rigid rods, resulting in 900 degrees of freedom if we ignore inertial effects (Fig. 2a,b). Snapshots of the shape of the chains under stretching conditions are shown in Fig. 2c, revealing highly heterogeneous dynamics of the chain extension (Fig. 2c,d), defined as the projected chain length along the elongational axis of the flow. This is our macroscopic coordinate of interest $\mathbf{Z}^*(t)$. Our aim is to model its stochastic evolution and understand the origin of its heterogeneity. We train the S-OnsagerNet on this dataset following the workflow in Fig. 1. The network architecture selection and training procedures are found in Methods. Our approach constructs two closure coordinates in addition to the chain extension $\mathbf{Z}^*(t)$, leading to a three-dimensional dynamical system—following equation (2) with learned functions M , W , V and σ —that governs the dynamics of stretching. We have empirically chosen the number of macroscopic coordinates: using more than three did not substantially improve predictive accuracy, whereas a two-variable system has modeling limitations due to physical symmetry. The detailed selection procedure of the reduced coordinate dimension is discussed in ‘Polymer dynamics analysis’ in Methods.

In Fig. 2e–m, we test the trained S-OnsagerNet on three unseen, different and representative initial polymer configurations. The selected

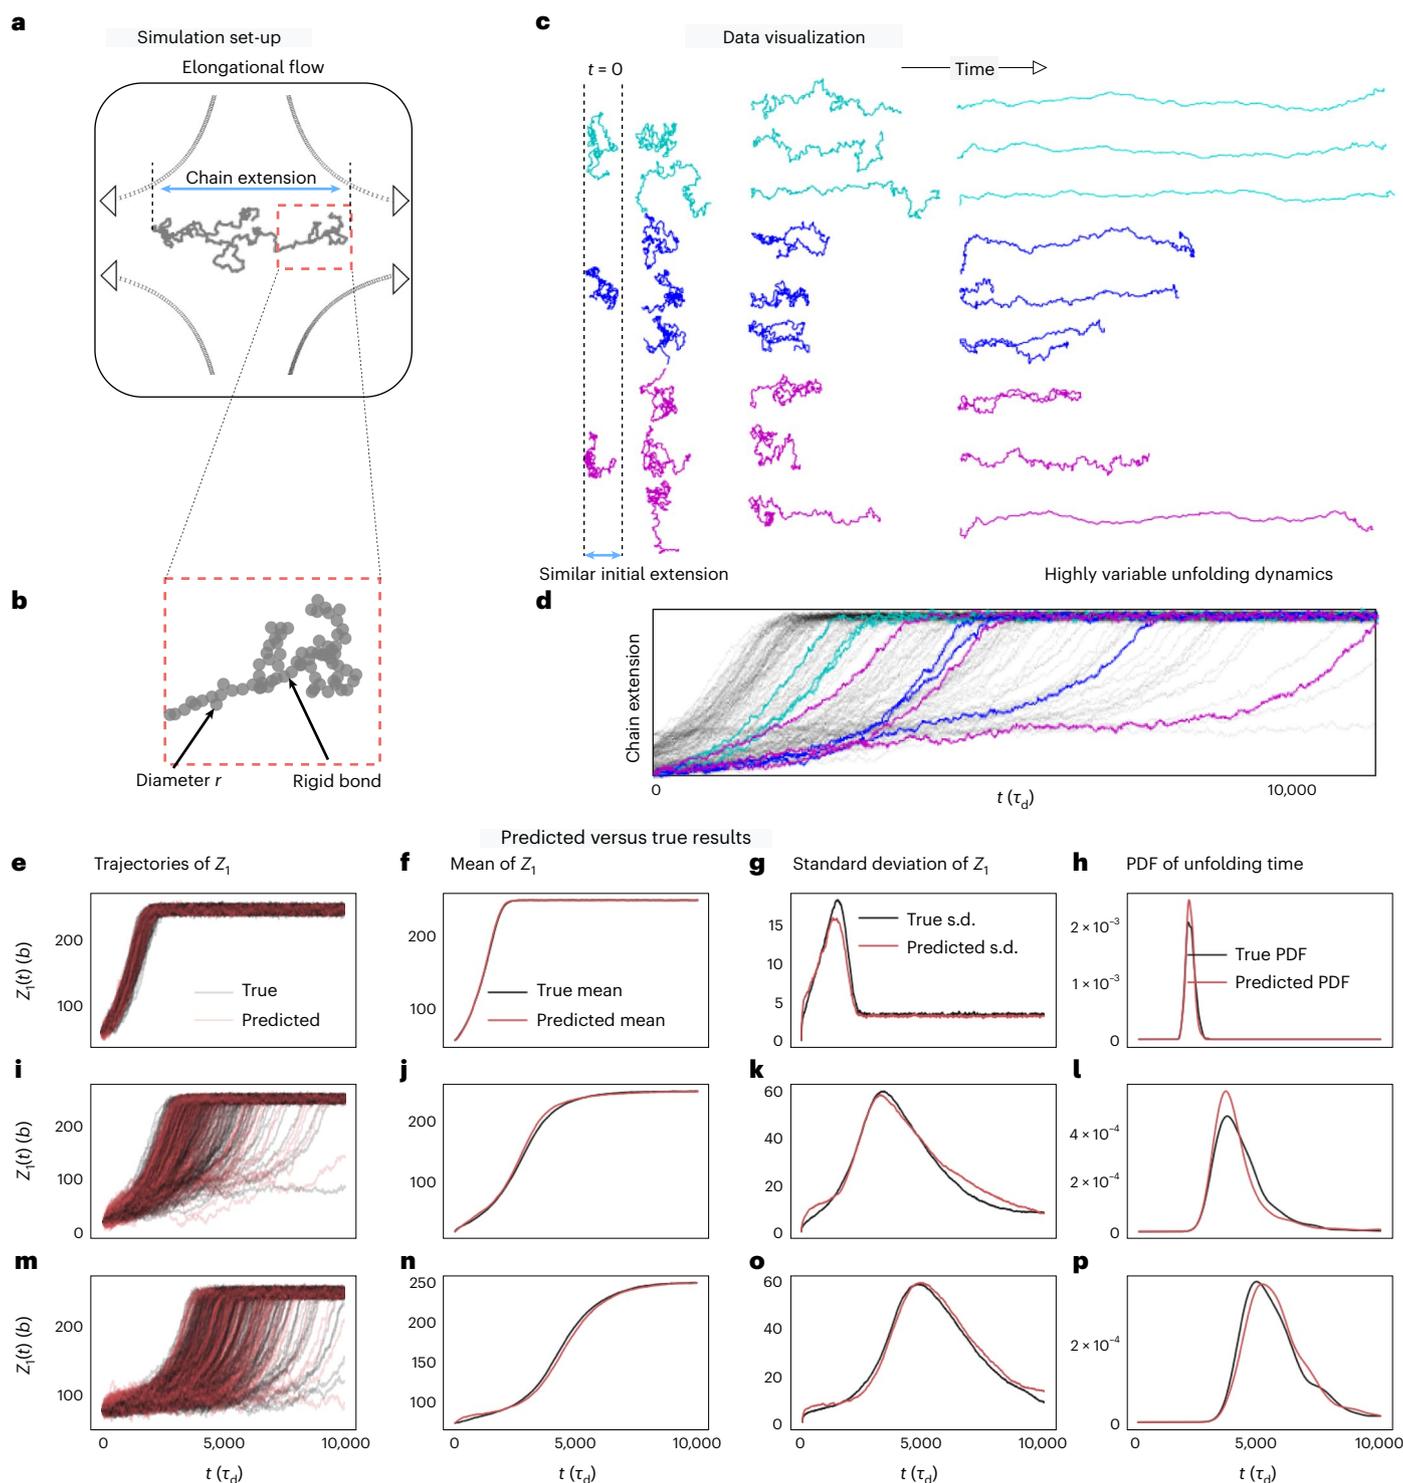


Fig. 2 | Simulation set-up, data visualization, and predicted versus true stretching trajectories and their statistics. **a,b**, The polymer chain is represented by a bead-rod model with bead diameter r (in units of $b = 10$ nm) and rigid bonds, subjected to hydrodynamic and Brownian forces. **b** shows a magnified portion of the polymer chain shown in **a**. **c,d**, The statistics of chain extension projected along the elongational axis are recorded. **c** shows the stretching polymers at different time instants and **d** depicts the extensions as a function of time for each trajectory. Times are reported in units of the characteristic rod diffusion time τ_d (see ‘Data preparation’ in Methods). Different initial conditions (colors) are chosen to have similar initial extension but varying

unfolding times. Identical initial configurations also have different unfolding dynamics due to thermal fluctuations. **e–p**, S-OnsagerNet can capture this heterogeneity. **e–g**, We plot 500 trajectories of polymer extensions (Z_1) from the same initial condition (**e**), together with their mean (**f**) and standard deviation (**g**). **h**, The probability density functions (PDFs) of unfolding times. **i–p**, Results for two other unseen initial chain configurations. **i** and **m** shows the extension length as a function of time for the medium and fast trajectories respectively. Similarly, **j** and **n** show the means, **k** and **o** the standard deviations, and **l** and **p** the probability density functions of unfolding times for the medium and fast trajectories, respectively.

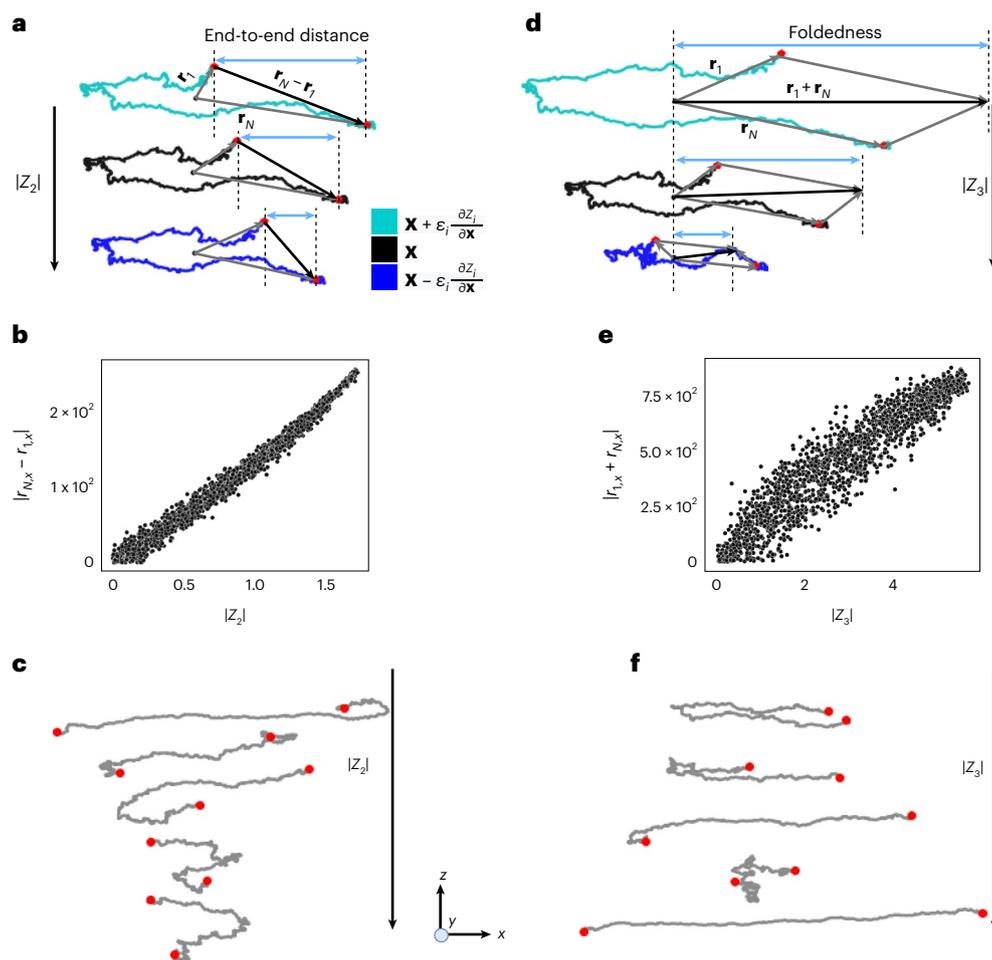


Fig. 3 | Physical interpretation of learned closure coordinates. a, Perturbation of the polymer chain $\mathbf{X} \pm \varepsilon_2 \frac{\partial Z_2}{\partial \mathbf{X}}$ from a given configuration, $\varepsilon_2 = 100 / \left\| \frac{\partial Z_2}{\partial \mathbf{X}} \right\|_2$. **b**, Plot of projected end-to-end distance $|r_{N,x} - r_{1,x}|$ as a function of $|Z_2|$ for the training data. **c**, Configurations of different polymer chains with decreasing $|Z_2|$ values. As $|Z_2|$ decreases, the projected distance between the chain ends decreases.

d, Perturbation of the polymer chain $\mathbf{X} \pm \varepsilon_3 \frac{\partial Z_3}{\partial \mathbf{X}}$ from a given configuration, $\varepsilon_3 = 260 / \left\| \frac{\partial Z_3}{\partial \mathbf{X}} \right\|_2$. **e**, Plot of degree of foldedness $|r_{1,x} + r_{N,x}|$ as a function of $|Z_3|$ for the training data. **f**, Configurations of different polymer chains with decreasing $|Z_3|$ values. As $|Z_3|$ decreases, the chains tend towards a more stretched state.

chains start with similar extension lengths, but subsequently stretch at vastly different rates. Figure 2e–g, i–k, m–p shows that the true statistics (black) can be accurately predicted (red). Moreover, the distributions of the time taken to reach a reference extension length are successfully captured (Fig. 2h, i, m).

Interpreting learned closure coordinates

Having shown that only two closure coordinates $\hat{\mathbf{Z}} = (Z_2, Z_3)$ are required to characterize the stochastic evolution of the extension length $Z^* = Z_1$, it is natural to probe the meaning of these discovered coordinates. Here we utilize an intrinsic property of neural networks—it represents the nonlinear reduction functions $\mathbf{X} \mapsto \mathbf{Z}$ as differentiable maps, as though we have learned their analytical forms. We compute via automatic differentiation the perturbations on a generic microscopic configuration \mathbf{X} in the directions of $\pm \partial Z_2 / \partial \mathbf{X}$ and $\pm \partial Z_3 / \partial \mathbf{X}$, corresponding to conformations with steepest changes in Z_2 and Z_3 , respectively. The resulting conformational changes suggest physical interpretation of these coordinates. For example, from Fig. 3a, we observe that perturbations in the direction of $\partial Z_2 / \partial \mathbf{X}$ tend to change the end-to-end distance in the elongational axis, or distance between the first and the last bead in the polymer chain along the elongational axis. We confirm this hypothesis by visualizing the correlation of the end-to-end distance and the magnitude of Z_2 in Fig. 3b, c. A similar

analysis reveals the correlation between Z_3 and a degree of foldedness of the chain in the elongational axis of the flow (Fig. 3d–f).

Free-energy landscape

The constructed potential V can be interpreted as a generalized free energy, allowing us to gain important insights into the dynamical landscape. The local minima of V represent stable or metastable states, while the saddle points correspond to transition states. The differentiable representation of V enables us to probe this landscape. Figure 4 shows two-dimensional projections of the three-dimensional free energy $V(Z_1, Z_2, Z_3)$. We identify the critical points of V by solving $\nabla V(\mathbf{Z}) = 0$ using the Broyden–Fletcher–Goldfarb–Shanno (BFGS) method. We found two stable fixed points and two saddle points of interest marked in Fig. 4. Using a simultaneously trained PCA-ResNet decoder, we can reconstruct the macroscopic spatial coordinates of the polymer chain at the critical points to identify their physical origin. Up to reflection symmetry in the elongational axis, the stable points correspond to fully stretched states, whereas the saddle points refer to completely folded states. The origin of the heterogeneity in unfolding times is now clear: a rapidly stretching polymer is the one that avoids the saddle point and goes directly to a stable minimum, whereas a slowly stretching chain gets trapped around the stable manifold (attractive part) of the saddle point for a long time, before finally escaping through the unstable

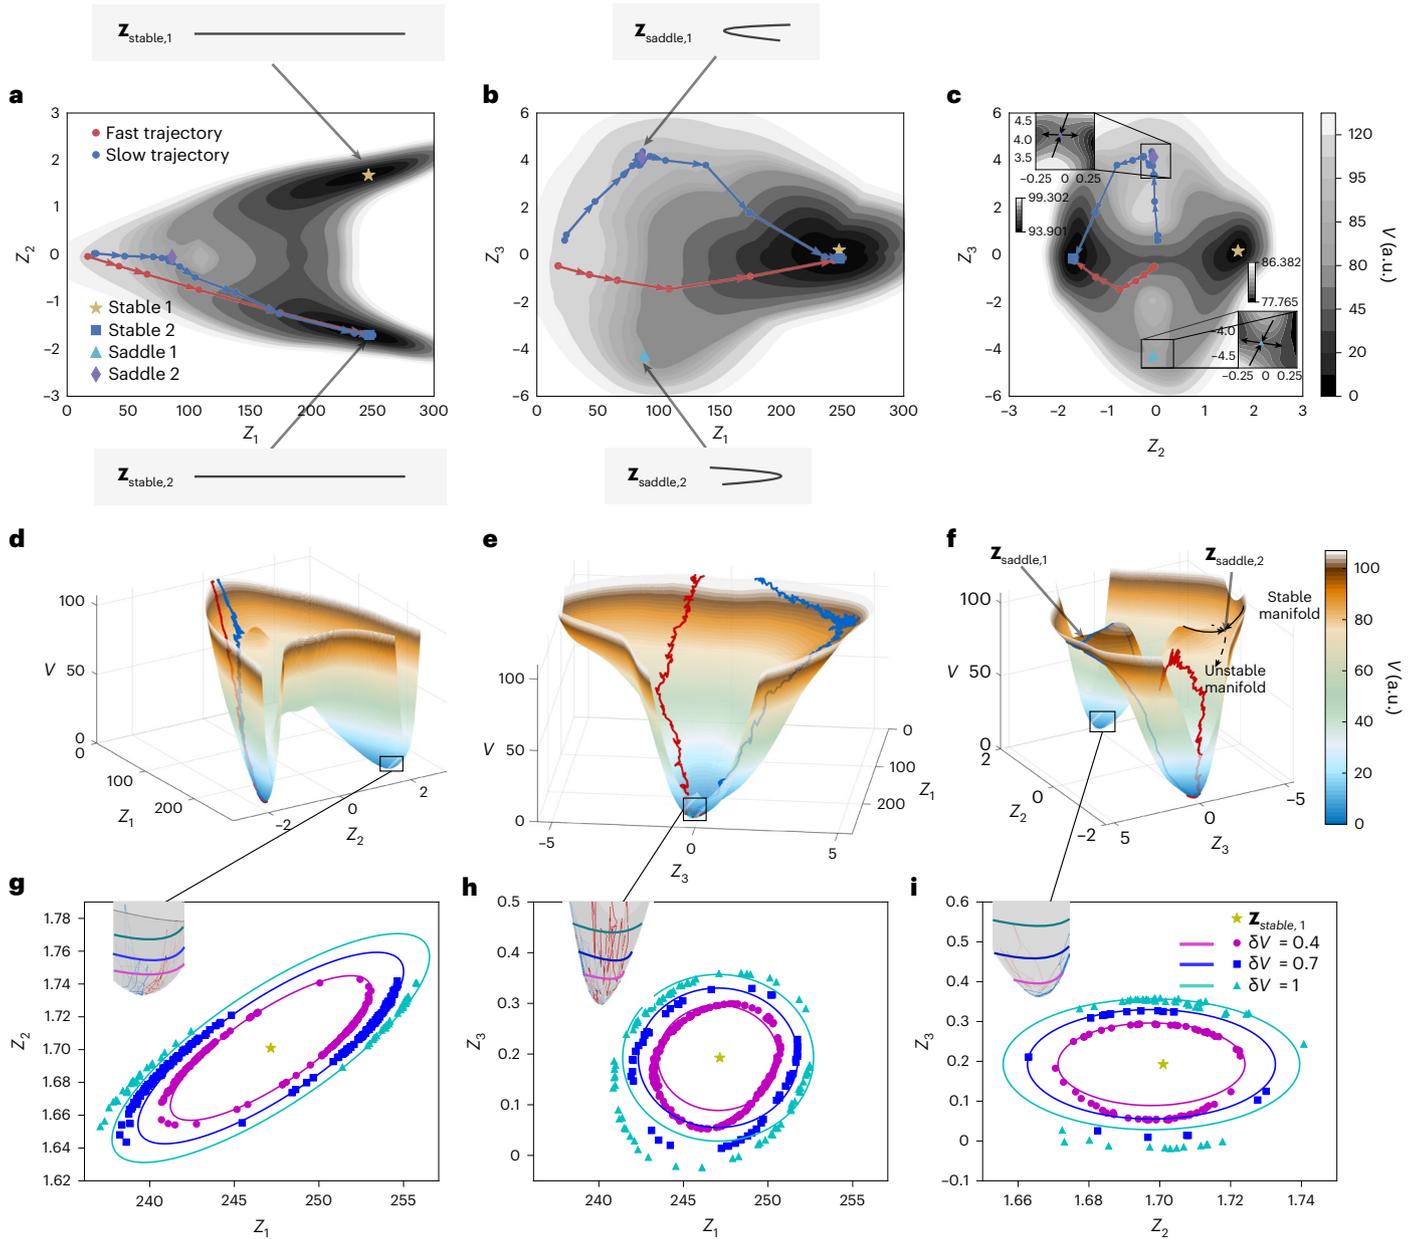


Fig. 4 | Learned potential energy landscape. a–f. We plot V projected onto the Z_1 – Z_2 (**a,d**), Z_1 – Z_3 (**b,e**) and Z_2 – Z_3 (**c,f**) planes. **a–c** are contour plots and **d–f** are surface plots to visualize the landscape. Insets: stable and unstable directions of the saddle points, corresponding to positive (top left inset) and negative (bottom right inset) Z_3 values. Projection is computed via minimization (for example, $V(Z_1, Z_2) = \min_{Z_3} V(Z_1, Z_2, Z_3)$), which at low temperatures closely approximates marginalizing with respect to the Boltzmann distribution. The stable and saddle points are marked on the energy landscape, and their corresponding reconstructed fully extended and folded states are shown. A pair of each exists due to reflection symmetry in the flow direction. Example ‘fast’ (red) and ‘slow’

(blue) trajectories from the training dataset are overlaid on the landscape. The fast trajectory avoids the saddle points and goes directly towards a stable minimum, whereas the slow trajectory gets trapped for long times near saddle 2, before finally escaping through its unstable manifold. For **b** and **e**, the stable manifolds of the saddles closely align with Z_2 , and hence are not visible due to minimization (marginalization). **g–i**, Scatterplots in the Z_1 – Z_2 (**g**), Z_1 – Z_3 (**h**) and Z_2 – Z_3 (**i**) planes, together with predicted isotherms (solid lines) capturing typical fluctuations around a fully stretched state $\mathbf{z}_{\text{stable},1}$. Insets: magnified views of the fluctuating trajectories around the stretched state over the energy landscape.

manifold (repulsive part) of the saddle towards the stable minimum (Supplementary Video 1). We confirm this by overlaying a fast and a slow unfolding trajectory with the potential landscape in Fig. 4c,f. Despite the similarity between the initial chain configurations, as demonstrated by the proximity between the initial points on the potential energy landscape, the chains show different stretching behaviors that can be rationalized by the constructed potential.

Moreover, a Taylor expansion via automatic differentiation of the learned $V(\mathbf{Z})$ captures the leading-order fluctuations near a stable

stretched state. We denote by δV a typical energy fluctuation around the stretched state (proportional to temperature). Then, ignoring small terms we find

$$\delta V \approx a_1 (\delta Z_1 - a_4 \delta Z_2)^2 + a_2 \delta Z_2^2 + a_3 \delta Z_3^2, \quad (3)$$

where $a_1 = 153.1$, $a_2 = 205.5$, $a_3 = 36.96$, $a_4 = 1.54$ and $\delta Z_i = Z_i - [Z_{\text{stable},1}]_i$ is the fluctuations in the thermodynamic variables. Note that the coefficients a_j implicitly depend on the strength of the flow. Equation (3)

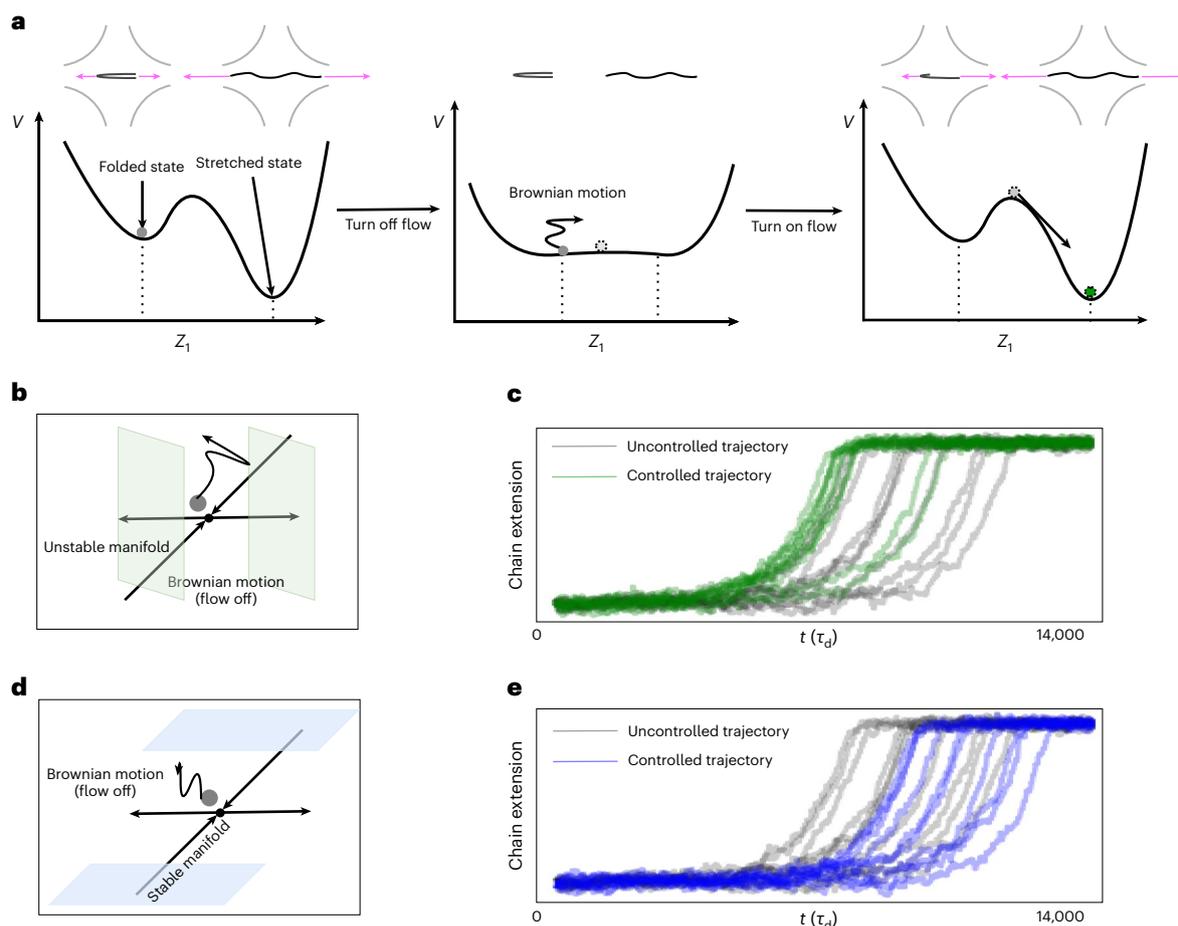


Fig. 5 | Data-driven control of the stretching dynamics. **a**, Control protocol to speed up unfolding. Left: projected along the extension direction, the polymer must overcome energy barriers to transition from the folded to the stretched state. Middle: in the control protocol, the flow is turned off if the reduced coordinate of the polymer is near the saddle point corresponding to the folded state, and the polymer drifts under the Brownian motion. Right: then the flow is turned back on if the reduced coordinates of the polymer becomes sufficiently aligned (green shaded region in **b**) with the unstable manifold of the saddle or when the the equilibration time reaches $100 \tau_d$. Without any control, a folded

state near the saddle point will unfold slowly (gray lines in **c**); with control, the chains unfold more rapidly (green lines in **c**). **b**, Schematic of the control protocol near the saddle point. **c**, Controlled and uncontrolled chain extension dynamics. For the slowest 10 trajectories shown, their mean unfolding time was reduced by 14.14%. **d, e**, Reversed control protocol (**d**) to slow down unfolding by turning on the flow when the reduced coordinates become aligned with the stable manifold instead (blue shaded region in **d**). **e** shows that the mean unfolding time increased by 14.96% (blue lines in **e**).

is an effective equation of state, from which we observe the positive correlation of Z_1 and Z_2 . The physical interpretation is that near the stretched state, the chain extension and the end-to-end distance tend to change simultaneously (see Fig. 4g–i and ‘Polymer dynamics analysis’ in Methods).

Controlling polymer stretching dynamics

Further, understanding the laws of the custom thermodynamics of polymer chain folding allows us to interact with the dynamics by designing controls over the polymer environment to initiate desired changes in its behavior. To this end, we perform another Taylor expansion of V —this time near a saddle point $Z_{\text{saddle},2}$ corresponding to a folded state to give another local equation of state

$$\delta V \approx b_1 \delta Z_1^2 - b_2 (\delta Z_2 - b_4 \delta Z_3)^2 + b_3 \delta Z_3^2, \quad (4)$$

where $b_1 = 102.96$, $b_2 = 31.13$, $b_3 = 24.16$ and $b_4 = 0.255$. Equation (4) suggests that to escape this saddle point leading to polymer unfolding, it is most effective to increase end-to-end distance (Z_2) while decreasing foldedness (Z_3) in a proportional way. This leads to a data-driven control protocol in Fig. 5a,b. We choose the external

elongational flow as the only control parameter (in real experiments it corresponds to switching on and off the flow of fluid or the electric field²⁶). We start with a polymer configuration near the saddle point of the energy landscape, corresponding to a folded state. Without any intervention, the subsequent unfolding is expected to follow a slow trajectory, staying near the folded state for a long time. From our landscape analysis above, the most effective escape from the saddle point is along its unstable manifold—approximately corresponding to increasing $|\delta Z_2 - b_4 \delta Z_3|$. Thus, to speed up unfolding, we can design the following control strategy: we turn off the external elongational flow, so that the polymer drifts randomly under Brownian forces around the saddle point. We track its reduced coordinates, and once we observe sufficient alignment with the unstable manifold, we turn on the externally applied elongational flow. We observe in Fig. 5c that this simple control system speeds up the unfolding dynamics substantially. We can also increase the unfolding time by reversing this protocol (Fig. 5d,e). These control strategies based on the learned thermodynamic description have notable advantages over classical model-free control regimes (for example, reinforcement learning), which may require large exploration times or small, finite state spaces³⁰.

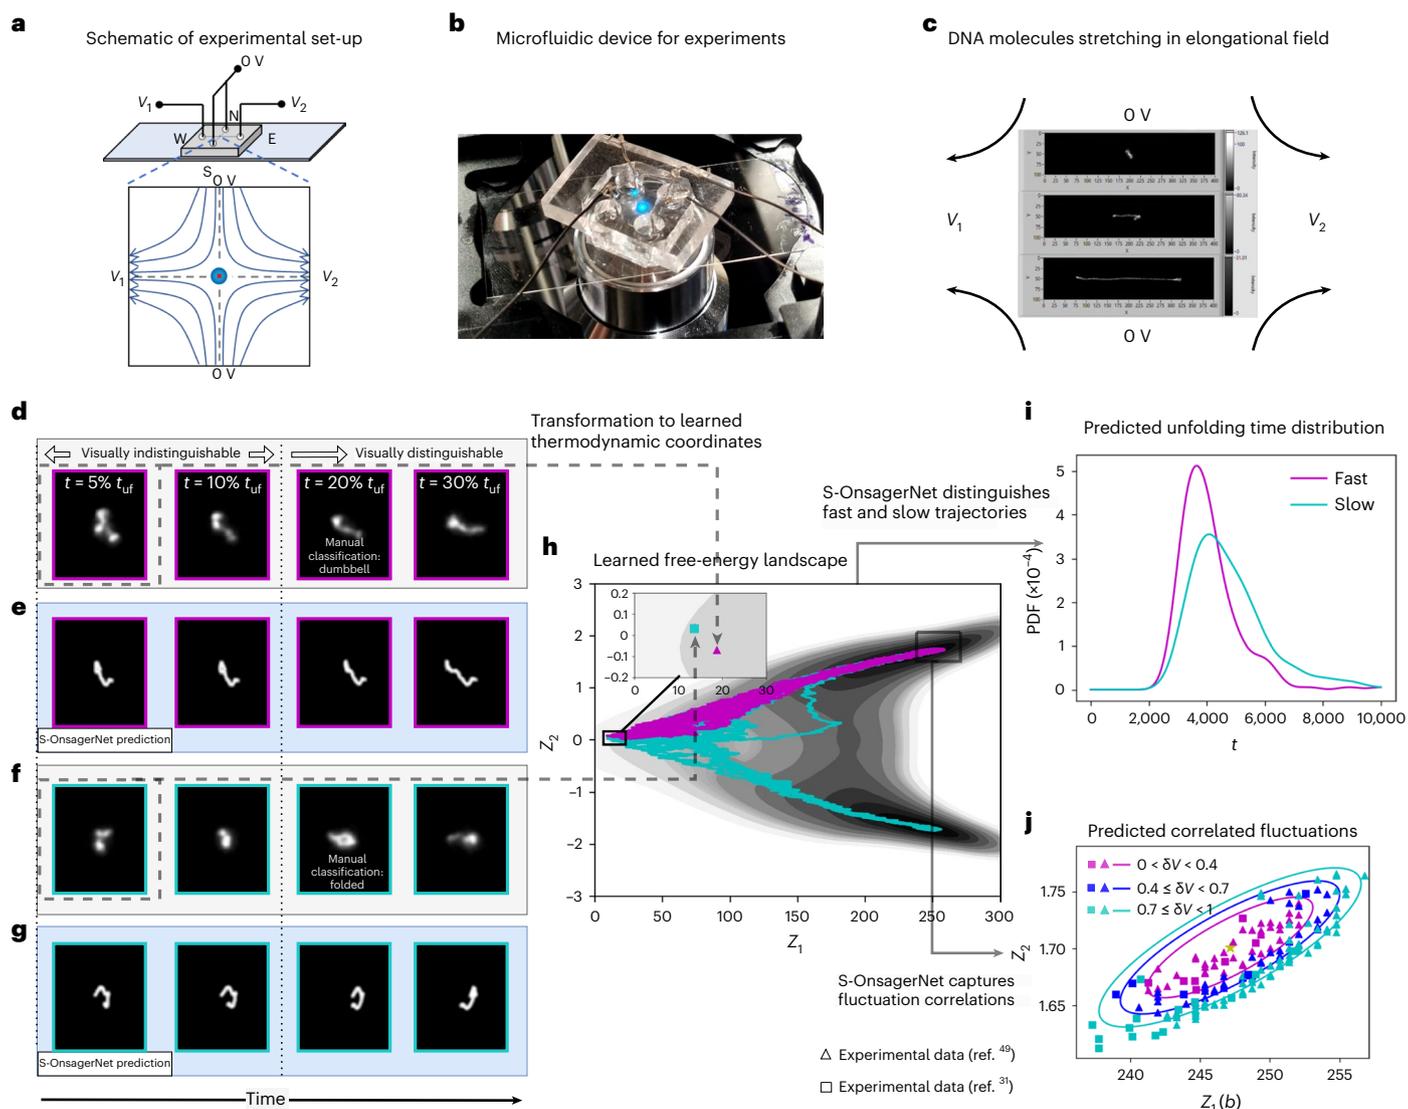


Fig. 6 | Analysis of experimental data. **a, b**, Schematic (**a**) and photograph (**b**) of the experimental set-up, consisting of a microfluidic cross-slot device and platinum electrodes in the reservoirs. Via the electrodes, positive voltage levels V_1 and V_2 are applied to the west and east reservoirs (W and E, respectively), while the north and south reservoirs (N and S, respectively) are kept at 0 V. During trapping, the negatively charged DNA will thus flow from N/S to W/E, until a molecule is trapped at the center of the channel—blue dot in the schematic. **c**, Snapshots of a DNA molecule stretching. **d, f**, Processed experimental images at various percentages of the unfolding time (t_{uf}) of a fast (**d**) and a slow (**f**) trajectory. The selected trajectories have similar initial configurations and are visually indistinguishable in terms of unfolding dynamics. **h**, Learned potential landscape and predicted slow and fast trajectories using S-OnsagerNet with the initial configurations at $t = 5\% t_{uf}$ from **d** and **f**. We note slight differences in the

initial Z_2 values only (inset). **e, g**, Reconstructed high-dimensional configurations of selected simulated trajectories with similar low-dimensional coordinates as the experimental configurations in **d** (corresponding to **e**) and **f** (corresponding to **g**). The S-OnsagerNet is capable of distinguishing between the manually classified dumbbell and folded states. **i**, Predicted PDFs of unfolding time with initial condition of fast and slow experimental trajectories at $t = 5\% t_{uf}$ using S-OnsagerNet. **j**, Fluctuations in Z_1 and Z_2 around the stable (stretched) state from experimental images. Data were obtained from ref. 49 (triangle) and ref. 31 (square). Reduced coordinates are constructed according to the procedure outlined in ‘Data preparation’ in Methods. Colors indicate predicted energy levels according to the learned potential. We observe that fluctuations in Z_1 and Z_2 are highly correlated and agree well with that predicted by the effective equation of state.

Experimental validation

Remarkably, some qualitative predictions of the constructed thermodynamic description are confirmed by physical experiments. Not only can one show that the constructed dynamical landscape allows for fine-grained classification of stretching behavior on simulation data (see ‘Polymer dynamics analysis’ in Methods), but also we demonstrate in Fig. 6 that this applies directly to physical experiments. Here we perform single-molecule experiments to observe the stretching trajectories of DNA molecules in a planar elongational field (see ‘Data preparation’ in Methods). We select two samples that initially appear similar (Fig. 6e, g), making it impossible to visually distinguish them in terms of

stretching behavior. We then cast them into the learned thermodynamic coordinates Z , which when superimposed on the free-energy landscape reveals that the Z_2 coordinates differ subtly, leading to different predicted stretching statistics (Fig. 6i). This substantially improves on human-level labeling, which can only occur much later in the dynamical evolution (Fig. 6h, j). Furthermore, we show in Fig. 6j that the effective equation of state equation (3) that captures the correlations of Z_1 and Z_2 around the stretched state also applies to experimental data from two sources, including the current experiments and previously available data³¹. These results demonstrate the promise of the current approach in enabling physical understanding and control of real polymer dynamics.

Modeling spatial epidemics

To further demonstrate general applicability, we employ our method to derive macroscopic dynamics of spatial epidemics. The classical spatial Susceptible-Infectious-Recovered (SIR) model³² governs microscopic evolution of infective and susceptible individuals on a spatial domain (Extended Data Fig. 1a–c). Using these microscopic trajectories, we construct a thermodynamic model that accurately models the evolution of the spatial averages of infective and susceptible individuals (Extended Data Fig. 1d,e) with an additional learned closure coordinate. Moreover, following the same approach as before, we can interpret this coordinate as the spatial overlap of infective and susceptible individuals (Extended Data Fig. 2), thereby rationalizing the dynamical landscape (Extended Data Fig. 3), where this overlap determines the onset and outcome of disease spread. Details are found in ‘Spatial epidemics analysis’ in Methods.

Discussion

The potential applicability of our method goes beyond polymer and epidemic dynamics, and includes general complex dissipative processes such as protein folding³³, self-assembly^{34,35} and glassy systems^{36,37}. Despite the importance of potential energy landscapes for material systems with functional properties, the challenge in constructing them has limited current approaches to systems with small degrees of freedom and/or requiring judicious selection of system descriptors based on expert knowledge³⁸. The method described in this work offers the potential of automating this process, creating pathways towards a multitude of opportunities for understanding and control over various complex systems and their scientific applications.

There are many worthwhile future research directions to further improve the robustness and generality of the proposed method. Here we inherently constrain the search space to macroscopic dynamics that conform to the generalized stochastic Onsager principle. Thus, it naturally has limited ability to model systems that may not readily admit such a description, such as chaotic systems. Moreover, the present model reduction and stability theory require the thermal noise to be small compared with the dissipative and conservative forces. While this is the case for the polymer dynamics studied here, our theory needs to be expanded to handle highly stochastic cases. In terms of training methodology, the current trial-and-error selection of the dimension of the closure variables can be made more systematic, for example, by building on manifold learning approaches¹⁴. Another potential improvement is the data sampling process. We observed in our numerical experiments that accurate construction of the dynamical landscape requires the trajectory data to sufficiently sample the regions of interest (stable and transition states). An adaptive or active learning algorithm^{39,40} that couples data sampling and S-OnsagerNet training can be developed to improve on the current random sampling strategy. On the scientific problem of polymer dynamics, we have considered motion under only a single stretching force. It is worthwhile to extend our study to varying stretching conditions to build a more comprehensive picture of polymer stretching. More broadly, one may apply our approach to learn macroscopic thermodynamics of other systems of scientific interest.

Methods

Theoretical results

In this section, we collect a number of theoretical results concerning the S-OnsagerNet approach. We first show that if a high-dimensional stochastic dynamical system satisfies the generalized stochastic Onsager principle (GSOP), then, any well-behaved reduction into a lower-dimensional system will result in one that obeys approximately the GSOP introduced in equation (2) (Theorem 1). An immediate consequence is that our model reduction approach is theoretically justified for a wide variety of dissipative and conservative systems, including molecular dynamics⁴¹, stochastic Hamiltonian systems⁴² and the stochastic Lotka–Volterra model⁴³. Next, we prove that dynamics described by the GSOP satisfy an energy dissipation law (Theorem 2) and thus our

machine learning approach produces stable dynamics at sufficiently low temperatures.

For convenience, we show that there are two equivalent forms of the GSOP. The formulation of the GSOP we use to construct our neural networks is

$$d\mathbf{Z}(t) = -(M(\mathbf{Z}(t)) + W(\mathbf{Z}(t)))\nabla V(\mathbf{Z}(t))dt + \sigma(\mathbf{Z}(t))d\mathbf{B}(t). \quad (5)$$

Here, B is the standard Brownian motion. Now, assuming $M(\cdot) + W(\cdot)$ is invertible, we define

$$\tilde{M}(\mathbf{Z}) = \frac{(M(\mathbf{Z}) + W(\mathbf{Z}))^{-1} + ((M(\mathbf{Z}) + W(\mathbf{Z}))^{-1})^T}{2},$$

$$\tilde{W}(\mathbf{Z}) = \frac{(M(\mathbf{Z}) + W(\mathbf{Z}))^{-1} - ((M(\mathbf{Z}) + W(\mathbf{Z}))^{-1})^T}{2},$$

$$\tilde{\sigma}(\mathbf{Z}) = (M(\mathbf{Z}) + W(\mathbf{Z}))^{-1}\sigma(\mathbf{Z}).$$

Observe that equation (5) can be rewritten in the form

$$(\tilde{M}(\mathbf{Z}(t)) + \tilde{W}(\mathbf{Z}(t)))d\mathbf{Z}(t) = -\nabla V(\mathbf{Z}(t))dt + \tilde{\sigma}(\mathbf{Z}(t))d\mathbf{B}(t). \quad (6)$$

A similar construction shows that we can also rewrite equation (6) in the form of equation (5) for any \tilde{M} , \tilde{W} , $\tilde{\sigma}$ assuming similar invertibility conditions. Thus, they are in fact equivalent. While the numerical implementation is based on equation (5), the form equation (6) is also useful, and in the following we refer to both as GSOP.

A similar construction shows that we can also rewrite equation (6) in the form of equation (5) for any \tilde{M} , \tilde{W} , $\tilde{\sigma}$ assuming similar invertibility conditions. Thus, they are in fact equivalent. While the numerical implementation is based on equation (5), the form equation (6) is also useful, and in the following we refer to both as GSOP.

Now, we demonstrate the general applicability of the GSOP in the context of model reduction. We consider a microscopic (high dimensional) dynamics satisfying a GSOP of the form

$$d\mathbf{X}(t) = -(M_1(\mathbf{X}(t)) + W_1(\mathbf{X}(t)))\nabla V_1(\mathbf{X}(t))dt + \sqrt{\epsilon_1}\Sigma_1(\mathbf{X}(t))d\mathbf{B}_1(t), \quad (7)$$

where $\mathbf{X}(t) \in \mathbb{R}^D$, $M_1(\cdot)$, $W_1(\cdot)$ are symmetric positive semi-definite and anti-symmetric matrix valued functions respectively, $\Sigma_1(\cdot)$ is the $D \times p_1$ -dimensional diffusion matrix, \mathbf{B}_1 is a p_1 -dimensional Brownian motion and ϵ is a positive parameter related to temperature (for example, $\epsilon_1 \propto k_B T$ with k_B being Boltzmann’s constant and T the temperature).

We now show that many high-dimensional systems of physical interest indeed satisfy a version of GSOP. We first consider the well-known Langevin dynamics, which has been used to model many stochastic dynamical systems for example molecular dynamics⁴⁴.

Example 1. The Langevin equation

$$m\ddot{x} = -\nabla U(x) - m\gamma_1\dot{x} + \sqrt{2m\gamma_1 k_B T}R(t), \quad (8)$$

can be written in the form of equation (6), where the dot denotes a time derivative, \dot{x} is the velocity, \ddot{x} is the acceleration, m is the mass, $U(x)$ is the particle interaction potential, and so $-\nabla U(x)$ is the potential force; γ_1 is the damping constant (units of reciprocal time), $R(t) = \dot{B}(t)$ is a delta-correlated stationary Gaussian process with zero-mean, satisfying

$$\langle R(t) \rangle = 0, \quad \langle R(t)R(t') \rangle = \delta(t - t').$$

If we set $\dot{x} = v$, the Langevin equation can be written as

$$\begin{pmatrix} m\gamma_1 & m \\ -m & 0 \end{pmatrix} \begin{pmatrix} dx \\ dv \end{pmatrix} = - \begin{pmatrix} \nabla U(x) \\ mv \end{pmatrix} dt + \begin{pmatrix} \sqrt{2m\gamma_1 k_B T} \\ 0 \end{pmatrix} dB(t).$$

Denoting $\mathbf{X} = \begin{pmatrix} x \\ v \end{pmatrix}$, $\tilde{M} = \begin{pmatrix} m\gamma_1 & 0 \\ 0 & 0 \end{pmatrix}$, $\tilde{W} = \begin{pmatrix} 0 & m \\ -m & 0 \end{pmatrix}$, $\Sigma = \begin{pmatrix} \sqrt{2m\gamma_1 k_B T} \\ 0 \end{pmatrix}$ and

$V(x, v) = U(x) + \frac{m}{2}v^2$, the Langevin equation can be written in the form of the GSOP as follows:

$$(\tilde{M} + \tilde{W}) d\mathbf{X} = -\nabla V dt + \Sigma dB.$$

Another important class of dynamical systems are those described by Poisson brackets⁴⁵, whose stochastic extension encompasses many applications, including the stochastic Lotka–Volterra models and variants⁴³. In the following, we show that these dynamical systems can also be written in the form of the GSOP.

Example 2. The stochastic dynamics with generalized coordinates $(q_1, \dots, q_n, p_1, \dots, p_n)$ described by generalized Poisson brackets

$$dF = (\{F, H\} - [F, H])dt + \sigma(F)dB, \tag{9}$$

can be written in the form of equation (7), where $H(q_1, \dots, q_n, p_1, \dots, p_n)$ is the Hamiltonian of the system and F is an arbitrary function depending on the system variables. The Poisson bracket $\{ \cdot, \cdot \}$ and the dissipation bracket $[\cdot, \cdot]$ are defined as

$$\begin{aligned} \{F, H\} &= \sum_{i=1}^n \left(\frac{\partial F}{\partial q_i} \frac{\partial H}{\partial p_i} - \frac{\partial F}{\partial p_i} \frac{\partial H}{\partial q_i} \right) \\ [F, H] &= J_F M J_H^T, \quad J_F = \left[\frac{\partial F}{\partial q_1}, \dots, \frac{\partial F}{\partial q_n}, \frac{\partial F}{\partial p_1}, \dots, \frac{\partial F}{\partial p_n} \right], \end{aligned}$$

where M is symmetric positive semi-definite.

Denote $(h_1, h_2, \dots, h_{2n}) = (q_1, \dots, q_n, p_1, \dots, p_n)$. By the definition of $\{F, H\}$ and taking $W = \begin{pmatrix} 0 & -I_n \\ I_n & 0 \end{pmatrix}$, where I_n is the n -dimensional identity

matrix, we have

$$\{F, H\} = \left(\frac{\partial F}{\partial q}, \frac{\partial F}{\partial p} \right) \begin{pmatrix} 0 & I_n \\ -I_n & 0 \end{pmatrix} \begin{pmatrix} \frac{\partial H}{\partial q} \\ \frac{\partial H}{\partial p} \end{pmatrix} = \nabla_{\mathbf{h}} F \begin{pmatrix} 0 & I_n \\ -I_n & 0 \end{pmatrix} (\nabla_{\mathbf{h}} H)^T = -J_F W J_H^T.$$

Hence, equation (9) can be written as

$$dF = (\nabla_{\mathbf{h}} F)^T d\mathbf{h} = -J_F (W + M) J_H^T dt + \sigma(F)dB.$$

Taking $\mathbf{F} = (h_1, \dots, h_{2n})$ and $\nabla_{\mathbf{h}} \mathbf{F} = I_{2n}$, we obtain

$$d\mathbf{h} = -J_{\mathbf{F}} (W + M) J_{\mathbf{H}}^T dt + \sigma(\mathbf{F})dB = -(W + M) \nabla_{\mathbf{h}} H dt + \sigma(\mathbf{h})dB.$$

Next, let us consider the reduction of a microscopic dynamical system satisfying a GSOP $(\mathbf{X}(t))$ into a macroscopic dynamical system $(\mathbf{Z}(t))$. This is achieved by a differentiable reduction function $\phi : \mathbb{R}^D \rightarrow \mathbb{R}^d$ such that $\mathbf{Z}(t) \approx \phi(\mathbf{X}(t))$. Moreover, we consider a differentiable reconstruction function $\psi : \mathbb{R}^d \rightarrow \mathbb{R}^D$ such that $\mathbf{X}(t) \approx \psi(\mathbf{Z}(t))$. Our main result is that $\mathbf{Z}(t)$ also satisfies an approximate GSOP. In other words, the GSOP family is approximately invariant under dimensional reduction, or coordinate transformation in general.

In the following, we adopt the notation

$$\begin{aligned} \nabla \phi_i(\mathbf{X}(t)) &:= \left(\frac{\partial \phi_i(\mathbf{X}(t))}{\partial x_1}, \frac{\partial \phi_i(\mathbf{X}(t))}{\partial x_2}, \dots, \frac{\partial \phi_i(\mathbf{X}(t))}{\partial x_D} \right)^T, \\ \nabla \phi(\mathbf{X}(t)) &:= (\nabla \phi_1(\mathbf{X}(t)), \dots, \nabla \phi_d(\mathbf{X}(t))). \end{aligned}$$

We will also adopt the following technical assumptions.

Assumption 1. The functions $M_1, W_1, \nabla V_1 : \mathbb{R}^D \rightarrow \mathbb{R}^D$, $\Sigma_1 : \mathbb{R}^D \rightarrow \mathbb{R}^{D \times p_1}$, $\phi : \mathbb{R}^D \rightarrow \mathbb{R}^d$ and $\psi : \mathbb{R}^d \rightarrow \mathbb{R}^D$ satisfy:

- (1) Growth condition: there exists a positive constant $L > 0$ such that, for all $\mathbf{x} \in \mathbb{R}^D$ and $\mathbf{z} \in \mathbb{R}^d$

$$|(M_1(\mathbf{x}) + W_1(\mathbf{x})) \nabla V_1(\mathbf{x})| + |\Sigma_1(\mathbf{x})| + \sum_{i=1}^d |\nabla \phi_i(\mathbf{x})| \leq L^2(1 + |\mathbf{x}|),$$

$$\sum_{i=1}^D |\nabla \psi_i(\mathbf{z})| + |\psi(\mathbf{z})| \leq L^2(1 + |\mathbf{z}|),$$

- (2) Lipschitz condition: there exists a positive constant $L > 0$ such that, for all $\mathbf{x} \in \mathbb{R}^D$, and $\mathbf{z} \in \mathbb{R}^d$, the function $(M_1(\mathbf{x}) + W_1(\mathbf{x})) \nabla V_1(\mathbf{x})$, $\Sigma_1(\mathbf{x})$, $\{\nabla \psi_i(\mathbf{z})\}_{i=1}^D$ and $\psi(\mathbf{z})$ satisfy the Lipschitz condition with constant L .
- (3) Approximate reconstruction: there exists $\epsilon_0 > 0$ such that, $\sup_{\mathbf{x} \in \Omega} |\mathbf{x} - \psi(\phi(\mathbf{x}))| < \epsilon_0$ where $\Omega \subset \mathbb{R}^D$ is a domain such that $\mathbf{X}(t) \in \Omega$ for all $t \in [0, T]$ almost surely.

Here, $|\cdot|$ is the Euclidean norm for a vector and Frobenius norm for a matrix.

Theorem 1. Let $\mathbf{X}(t)$ satisfy equation (7) and $\mathbf{Z}(t)$ satisfy equation (5) with

$$\begin{aligned} M(\mathbf{Z}) &= \nabla \phi[\psi(\mathbf{Z})]^T M_1(\psi(\mathbf{Z})) \nabla \phi[\psi(\mathbf{Z})], \\ W(\mathbf{Z}) &= \nabla \phi[\psi(\mathbf{Z})]^T W_1(\psi(\mathbf{Z})) \nabla \phi[\psi(\mathbf{Z})], \\ V(\mathbf{Z}) &= V_1[\psi(\mathbf{Z})], \\ \sigma_1(\mathbf{Z}) &= \sqrt{\epsilon_1} \nabla \phi[\psi(\mathbf{Z})]^T \Sigma_1[\psi(\mathbf{Z})], \\ \sigma(\mathbf{Z}) &= [\sigma_1(\mathbf{Z}) \sigma_1(\mathbf{Z})^T]^{\frac{1}{2}}, \end{aligned}$$

Then, for each $u \in C^2(\mathbb{R}^d)$ (twice continuously differentiable function) there exists a constant $C > 0$, independent of ϵ_0 and ϵ_1 , such that

$$|\mathbb{E}u(\phi[\mathbf{X}(t)]) - \mathbb{E}u(\mathbf{Z}(t))| \leq C(\epsilon_0 + \epsilon_1).$$

Proof. Suppose $\mathbf{Y}(t)$ satisfy the following equation:

$$d\mathbf{Y}(t) = -(M(\mathbf{Y}(t)) + W(\mathbf{Y}(t))) \nabla V(\mathbf{Y}(t)) dt + \sigma_1(\mathbf{Y}(t)) dB_1(t), \tag{10}$$

we will prove $|\mathbb{E}[u(\phi[\mathbf{X}(t)]) - u(\mathbf{Y}(t))]| \leq C(\epsilon_0 + \epsilon_1)$. By Itô's formula, equation (7) and Assumption 1, we obtain

$$\begin{aligned} d\phi_i(\mathbf{X}(t)) &= [\nabla \phi_i(\mathbf{X}(t))]^T d\mathbf{X}(t) + \frac{1}{2} [d\mathbf{X}(t)]^T \nabla^2 \phi_i(\mathbf{X}(t)) d\mathbf{X}(t) \\ &= [\nabla \phi_i(\mathbf{X}(t))]^T [-(M_1(\mathbf{X}(t)) + W_1(\mathbf{X}(t))) \nabla V_1(\mathbf{X}(t)) dt \\ &\quad + \sqrt{\epsilon_1} \Sigma_1(\mathbf{X}(t)) dB_1(t)] \\ &\quad + \frac{\epsilon_1}{2} \text{Tr}[\Sigma_1(\mathbf{X}(t)) \Sigma_1^T(\mathbf{X}(t)) \nabla^2 \phi_i(\mathbf{X}(t))] dt, \end{aligned} \tag{11}$$

where $\nabla^2 \phi_i$ is the Hessian matrix of ϕ_i and Tr is the trace of a square matrix.

Now, by definition, $\mathbf{Y}(t)$ satisfies the following stochastic differential equation (SDE)

$$\begin{aligned} d\mathbf{Y}(t) &= -(M(\mathbf{Y}(t)) + W(\mathbf{Y}(t))) \nabla V(\mathbf{Y}(t)) dt + \sigma_1(\mathbf{Y}(t)) dB_1(t) \\ &= -\nabla \phi[\psi(\mathbf{Y}(t))]^T [M_1(\psi(\mathbf{Y}(t))) + W_1(\psi(\mathbf{Y}(t)))] \nabla \phi[\psi(\mathbf{Y}(t))] \nabla V(\mathbf{Y}(t)) dt \\ &\quad + \sqrt{\epsilon_1} \nabla \phi[\psi(\mathbf{Y}(t))]^T \Sigma_1(\psi(\mathbf{Y}(t))) dB_1(t) \\ &= -\nabla \phi[\psi(\mathbf{Y}(t))]^T [M_1(\psi(\mathbf{Y}(t))) + W_1(\psi(\mathbf{Y}(t)))] \nabla V_1(\psi(\mathbf{Y}(t))) dt \\ &\quad + \sqrt{\epsilon_1} \nabla \phi[\psi(\mathbf{Y}(t))]^T \Sigma_1(\psi(\mathbf{Y}(t))) dB_1(t). \end{aligned} \tag{12}$$

Subtracting equations (11) and (12), and integrating on $[0, t]$, we get

$$\begin{aligned} & \phi_i(\mathbf{X}(t)) - Y_i(t) \\ &= \int_0^t [\nabla\phi_i(\mathbf{X}(r))^T f_i(\mathbf{X}(r)) - \nabla\phi_i[\psi(\mathbf{Y}(r))]^T f_i(\psi(\mathbf{Y}(r)))] dr \\ &+ \epsilon_1 \int_0^t g_i(\mathbf{X}(r)) dr + \sqrt{\epsilon_1} \int_0^t [\nabla\phi_i(\mathbf{X}(r))^T \Sigma_1(\mathbf{X}(r)) \\ &- \nabla\phi_i[\psi(\mathbf{Y}(r))]^T \Sigma_1(\psi(\mathbf{Y}(r)))] dB_1(r), \end{aligned}$$

where $f_i(\cdot) = -(M_i(\cdot) + W_i(\cdot)) \nabla V_1(\cdot)$ and $g_i(\cdot) = \frac{1}{2} \text{Tr}[\Sigma_1 \Sigma_1^T \nabla^2 \phi_i(\cdot)]$ with $i = 1, 2, \dots, d$.

By Cauchy–Schwarz inequality and Itô isometry, we have

$$\begin{aligned} & \mathbb{E}|\phi_i(\mathbf{X}(t)) - Y_i(t)|^2 \\ & \leq 3\mathbb{E} \left| \int_0^t [\nabla\phi_i(\mathbf{X}(r))^T f_i(\mathbf{X}(r)) - \nabla\phi_i[\psi(\mathbf{Y}(r))]^T f_i(\psi(\mathbf{Y}(r)))] dr \right|^2 \\ & + 3\mathbb{E} \left| \int_0^t \epsilon_1 g_i(\mathbf{X}(r)) dr \right|^2 \\ & + 3\epsilon_1 \mathbb{E} \left| \int_0^t [\nabla\phi_i(\mathbf{X}(r)) \Sigma_1(\mathbf{X}(r)) - \nabla\phi_i[\psi(\mathbf{Y}(r))]^T \Sigma_1(\psi(\mathbf{Y}(r)))] dB_1(r) \right|^2 \quad (13) \\ & \leq 3t\mathbb{E} \int_0^t |\nabla\phi_i(\mathbf{X}(r))^T f_i(\mathbf{X}(r)) - \nabla\phi_i[\psi(\mathbf{Y}(r))]^T f_i(\psi(\mathbf{Y}(r)))|^2 dr \\ & + 3\epsilon_1^2 t \mathbb{E} \int_0^t |g_i(\mathbf{X}(r))|^2 dr \\ & + 3\epsilon_1 \mathbb{E} \int_0^t |\nabla\phi_i(\mathbf{X}(r)) \Sigma_1(\mathbf{X}(r)) - \nabla\phi_i[\psi(\mathbf{Y}(r))]^T \Sigma_1(\psi(\mathbf{Y}(r)))]|^2 dr. \end{aligned}$$

By Assumption 1, there exists positive constants C_1 and C_2 such that

$$\begin{aligned} & \mathbb{E} \int_0^t |\nabla\phi_i(\mathbf{X}(r))^T f_i(\mathbf{X}(r)) - \nabla\phi_i[\psi(\mathbf{Y}(r))]^T f_i(\psi(\mathbf{Y}(r)))|^2 dr \\ & \leq \mathbb{E} \int_0^t |\nabla\phi_i(\mathbf{X}(r))^T f_i(\mathbf{X}(r)) - \nabla\phi_i[\psi(\phi(\mathbf{X}(r)))]^T f_i(\mathbf{X}(r)) \\ & + \nabla\phi_i[\psi(\phi(\mathbf{X}(r)))]^T f_i(\mathbf{X}(r)) \\ & - \nabla\phi_i[\psi(\mathbf{Y}(r))]^T f_i(\mathbf{X}(r)) + \nabla\phi_i[\psi(\mathbf{Y}(r))]^T f_i(\mathbf{X}(r)) \\ & - \nabla\phi_i[\psi(\mathbf{Y}(r))]^T f_i(\psi(\phi(\mathbf{X}(r)))) \\ & + \nabla\phi_i[\psi(\mathbf{Y}(r))]^T f_i(\psi(\phi(\mathbf{X}(r)))) - \nabla\phi_i[\psi(\mathbf{Y}(r))]^T f_i(\psi(\mathbf{Y}(r)))|^2 dr \\ & \leq C_1 \mathbb{E} \int_0^t |\mathbf{X}(r) - \psi(\phi(\mathbf{X}(r)))|^2 dr + C_2 \mathbb{E} \int_0^t |\phi(\mathbf{X}(r)) - \mathbf{Y}(r)|^2 dr \\ & \leq C_1 t \epsilon_0^2 + C_2 \mathbb{E} \int_0^t |\phi(\mathbf{X}(r)) - \mathbf{Y}(r)|^2 dr. \end{aligned}$$

We employ a similar above argument of the third term of equation (13) and get

$$\mathbb{E}|\phi(\mathbf{X}(t)) - \mathbf{Y}(t)|^2 \leq C_3 \int_0^t \mathbb{E}|\phi(\mathbf{X}(r)) - \mathbf{Y}(r)|^2 dr + C_4(\epsilon_0 + \epsilon_1)^2.$$

Here we have used the Lipschitz conditions and the boundness of the first moment of $f_i(\mathbf{X})$ and $g_i(\mathbf{X})$, which is implied by the growth condition in Assumption 1. This shows that $\phi(\mathbf{X}(t))$ and $\mathbf{Y}(t)$ are close in the mean-square sense. By Gronwall’s inequality, we get

$$\mathbb{E}|\phi(\mathbf{X}(t)) - \mathbf{Y}(t)|^2 \leq C_5(\epsilon_0 + \epsilon_1)^2. \quad (14)$$

Now, we employ a similar argument to show that $u(\phi(\mathbf{X}(t)))$ and $u(\mathbf{Y}(t))$ are close for any sufficiently smooth u . We apply Itô formula to

$u(\phi(\mathbf{X}(t)))$ and $u(\mathbf{Y}(t))$ to get

$$\begin{aligned} du(\phi(\mathbf{X}(t))) &= \nabla u \cdot d\phi(\mathbf{X}(t)) + \frac{1}{2} d\phi(\mathbf{X}(t))^T \nabla^2 u(\phi(\mathbf{X}(t))) d\phi(\mathbf{X}(t)) \\ &= \nabla u \cdot [(\nabla\phi(\mathbf{X}(t)))^T f_i(\mathbf{X}(t)) + \epsilon_1 g(\mathbf{X}(t))] dt \\ &+ \nabla\phi(\mathbf{X}(t)) \Sigma_1(\mathbf{X}(t)) dB_1(t) \\ &+ \frac{1}{2} \epsilon_1 \text{Tr}[\nabla\phi(\mathbf{X}(t)) \Sigma_1(\mathbf{X}(t)) (\nabla\phi(\mathbf{X}(t)) \Sigma_1(\mathbf{X}(t)))^T \nabla^2 u] dt, \\ du(\mathbf{Y}(t)) &= \nabla u \cdot d\mathbf{Y}(t) + \frac{1}{2} d\mathbf{Y}(t)^T \nabla^2 u(\mathbf{Y}(t)) d\mathbf{Y}(t) \\ &= \nabla u \cdot \nabla\phi[\psi(\mathbf{Y}(t))]^T f_i(\psi(\mathbf{Y}(t))) dt + \sqrt{\epsilon_1} \Sigma_1(\psi(\mathbf{Y}(t))) dB_1(t) \\ &+ \frac{\epsilon_1}{2} \text{Tr}[\nabla\phi[\psi(\mathbf{Y}(t))] \Sigma_1(\psi(\mathbf{Y}(t))) (\nabla\phi[\psi(\mathbf{Y}(t))] \Sigma_1(\psi(\mathbf{Y}(t))))^T \nabla^2 u] dt. \end{aligned}$$

As before, by Assumption 1, there exist positive constants C_6 and C_7 such that

$$\mathbb{E}|u(\phi(\mathbf{X}(t))) - u(\mathbf{Y}(t))|^2 \leq C_6 \mathbb{E} \int_0^t |\phi(\mathbf{X}(r)) - \mathbf{Y}(r)|^2 dr + C_7(\epsilon_0 + \epsilon_1)^2. \quad (15)$$

Combining equations (14) and (15), we obtain

$$\mathbb{E}|u(\phi(\mathbf{X}(t))) - u(\mathbf{Y}(t))|^2 \leq C^2(\epsilon_0 + \epsilon_1)^2.$$

By Jensen’s inequality, we can get

$$\mathbb{E}|u(\phi(\mathbf{X}(t))) - u(\mathbf{Y}(t))| \leq C(\epsilon_0 + \epsilon_1).$$

According to the equations (5) and (10), we can get $Z(t)$ and $\mathbf{Y}(t)$ has the same distribution, that is

$$\mathbb{E}u(\mathbf{Y}(t)) - \mathbb{E}u(Z(t)) = 0.$$

Finally, by triangle inequality, we have

$$\begin{aligned} & |\mathbb{E}u(\phi(\mathbf{X}(t))) - \mathbb{E}u(\mathbf{Y}(t))| \\ &= |\mathbb{E}u(\phi(\mathbf{X}(t))) - \mathbb{E}u(\mathbf{Y}(t)) + \mathbb{E}u(\mathbf{Y}(t)) - \mathbb{E}u(Z(t))| \\ &\leq \mathbb{E}|u(\phi(\mathbf{X}(t))) - u(\mathbf{Y}(t))| + |\mathbb{E}u(\mathbf{Y}(t)) - \mathbb{E}u(Z(t))| \\ &\leq C(\epsilon_0 + \epsilon_1). \end{aligned}$$

This completes the proof.

This results demonstrate the validity of the GSOP as a dimensionality reduction method. In short, it says that if the microscopic dynamics satisfies a GSOP, then the macroscopic dynamics will also satisfy a GSOP approximately. As a large amount of conservative and dissipative microscopic physical systems are shown to satisfy the GSOP, the S-OnsagerNet approach based on the GSOP is a principled model reduction ansatz for physical processes.

Next, we show the stability of a solution of the GSOP. More precisely, we prove in Theorem 2 below that the mean of the potential is non-increasing in t for sufficiently low temperatures (small $|\sigma|$). Consequently, the S-OnsagerNet produces dissipative dynamical systems that enjoy long-term stability.

Theorem 2. The solution of equation (5) satisfies the dissipation law

$$\begin{aligned} \mathbb{E}V(Z(t)) - \mathbb{E}V(Z(0)) &= - \int_0^t \mathbb{E} \|\nabla V(Z(r))\|_M^2 dr \\ &+ \frac{1}{2} \int_0^t \mathbb{E} \text{Tr}[\sigma(Z(r)) \sigma(Z(r))^T \nabla^2 V(Z(r))] dr. \end{aligned}$$

Here $\|\cdot\|_M^2$ denotes $|M^{1/2} \cdot|^2$ where $M^{1/2}$ is the non-negative square-root of M . If we assume further that there exists a positive constant α such that

$\mathbf{Z}^T M(\mathbf{Z})\mathbf{Z} \geq \alpha|\mathbf{Z}|^2$ and $\frac{1}{2}\text{Tr}[\alpha(\mathbf{Z})\sigma(\mathbf{Z})^T \nabla^2 V(\mathbf{Z})] \leq \alpha|\nabla V(\mathbf{Z})|^2$ for all \mathbf{Z} , then $\mathbb{E}[V(\mathbf{Z}(t))]$ is non-increasing in t .

Proof. By Itô formula, we obtain

$$\begin{aligned} dV(\mathbf{Z}) &= \nabla V \cdot d\mathbf{Z} + \frac{1}{2}(d\mathbf{Z})^T \nabla^2 V d\mathbf{Z} \\ &= -\nabla V \cdot [(M + W)\nabla V]dt + \nabla V \cdot \sigma d\mathbf{B}_1 + \frac{1}{2}\text{Tr}[\sigma\sigma^T \nabla^2 V]dt \quad (16) \\ &= -\|\nabla V\|_M^2 dt + \frac{1}{2}\text{Tr}[\sigma\sigma^T \nabla^2 V]dt + \nabla V \cdot \sigma d\mathbf{B}_1, \end{aligned}$$

where $\nabla V^T M \nabla V = \|\nabla V\|_M^2$ and $\nabla V^T W \nabla V = 0$ are used.

Integrating equation (16) from 0 to t and taking expectation, we obtain

$$\mathbb{E}V(\mathbf{Z}(t)) - \mathbb{E}V(\mathbf{Z}(0)) = -\int_0^t \mathbb{E}\|\nabla V\|_M^2 dr + \frac{1}{2}\int_0^t \mathbb{E}\text{Tr}[\sigma\sigma^T \nabla^2 V] dr.$$

Finally, according to the proposed condition, it is easy to arrive at $\mathbb{E}V(\mathbf{Z}(t)) - \mathbb{E}V(\mathbf{Z}(0)) \leq -\alpha\int_0^t \mathbb{E}\|\nabla V\|^2 dr + \frac{1}{2}\int_0^t \mathbb{E}\text{Tr}[\sigma\sigma^T \nabla^2 V] dr \leq 0$, thus the mean of the energy potential is non-increasing in t . This completes the proof.

Theorem 2 shows that energy is being dissipated if the temperature of the ambient reservoir is sufficiently low. Accordingly, if the free energy V has compact level sets, then the dynamics at low temperatures will be confined on average to these compact sets and is thus stable. This contrasts with unstructured methods that may learn dynamics that are accurate for short times but induce instability at long times. We numerically verify this energy dissipation law in the learned polymer dynamics in Supplementary Fig. 1.

Model implementation

We provide in this section the detailed implementation of S-OnsagerNet.

Model architecture. We begin with discussing the architecture design of the neural network approximators. Following the acquisition of the closure coordinates $\mathbf{Z}(t) = (\mathbf{Z}^s(t), \dot{\mathbf{Z}}(t))$, the S-OnsagerNet architecture implements equation (5). To ensure the symmetric positive definiteness of $M(\mathbf{Z})$ and the anti-symmetry of $W(\mathbf{Z})$, we use a neural network to approximate $A(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^{d^2}$ with dimension d^2 . Then, we take the lower-triangular part as $L_1(\mathbf{Z})$ and the upper-triangular part as $L_2(\mathbf{Z})$. $M(\mathbf{Z})$ and $W(\mathbf{Z})$ are represented by

$$\begin{aligned} M(\mathbf{Z}) &= L_1(\mathbf{Z})L_1(\mathbf{Z})^T + \alpha I, \\ W(\mathbf{Z}) &= L_2(\mathbf{Z}) - L_2(\mathbf{Z})^T, \end{aligned}$$

where α is a positive constant and I is an identity matrix.

The energy function $V(\cdot)$ is lower bounded, so we use the following structure

$$V(\mathbf{Z}) = \frac{1}{2} \sum_{i=1}^m \left(U_i(\mathbf{Z}) + \sum_{j=1}^d \gamma_{ij} \mathbf{z}_j \right)^2 + \beta |\mathbf{Z}|^2,$$

where $U(\mathbf{Z})$ is a neural network with d -dimensional input and m -dimensional output, $\{\gamma_{ij}\}$ is a trainable matrix and β is a positive parameter.

In the architecture used for the polymer dynamics application in this paper, we set $\alpha = 0.1$ and utilize a neural network with 2 hidden layers with 20 neurons each and the tanh activation function to approximate $M(\mathbf{z})$ and $W(\mathbf{z})$. To parameterize the potential $V(\mathbf{z})$, we decompose it into a sum of squares of the output layer (size $m = 50$) of 1 hidden layer neural network with 128 hidden neurons and the ReQUr

activation function^{7,46}. This is to ensure that the potential satisfies the correct growth conditions as outlined in Assumption 1.

For the diffusion matrix $\sigma(\mathbf{z})$, as it has no symmetry constraints other than a growth condition, we use a fully connected neural network to approximate it. In our polymer dynamics application, we found empirically that a diagonal, z -independent diffusion matrix (corresponding to a linear network with zero weight and trainable diagonal bias) performed the best, but our algorithm can handle general architectures for $\sigma(\mathbf{z})$.

Closure coordinate construction. We now provide details of the procedure to construct closure coordinates $\hat{\mathbf{Z}}(t)$ using the time series observation data of chain configuration coordinates at $\{t_k\}_{k=1}^{N_t}$, with $0 = t_0 < t_1 \dots < t_{N_t} = T$. The available data are $\{\mathbf{x}_j\}_{j=1}^M$ with $\mathbf{x}_j = \{\mathbf{X}(t_i)\}_{i=1}^{N_t} \in \mathbb{M}^{D \times N_t}$, where M is the number of trajectories and $\mathbf{X}(t_i)^{(j)}$ is the j th observation trajectory at $t = t_i$. We obtained 610 observational trajectories, and for each trajectory, the number of time snapshots is 1,001, that is $M = 610$ and $N_t = 1,001$. We reshape the observation data as $\Xi = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_M]$, where $\Xi \in \mathbb{M}^{M \times D \times N_t}$. We re-center the data, such that the mean of the training data is zero for each time snapshot, and set it as X . The covariance matrix of X is $S = \text{Cov}(X)$. Denote its eigenvalues by $\Lambda = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_D)$ (arranged in non-increasing order) and corresponding eigenvectors as $V = (\mathbf{V}_1, \mathbf{V}_2, \dots, \mathbf{V}_D)$.

We use the following PCA-ResNet encoder to find the closure coordinates

$$\hat{\mathbf{Z}}(t) = \mathbf{Z}_{2:d}(t) = \hat{\psi}(\mathbf{X}(t)) = P_d \mathbf{X}(t) + \text{NN}_e(\mathbf{X}(t)), \quad (17)$$

where $P_d = \frac{1}{\sqrt{\lambda_{1:d-1}}} V_{1:d-1}$ and $\text{NN}_e(\cdot)$ is a fully connected neural network with input dimension D and output dimension $d - 1$. We can reconstruct the high-dimensional coordinates via the decoder

$$\tilde{\mathbf{X}}(t) = \psi(\hat{\mathbf{Z}}(t)) = P_d^\dagger \hat{\mathbf{Z}}(t) + \text{NN}_d(\hat{\mathbf{Z}}(t)), \quad (18)$$

where $P_d^\dagger = V_{1:d-1}^T \sqrt{\lambda_{1:d-1}}$ and $\text{NN}_d(\hat{\mathbf{Z}}(t))$ is another neural network with input dimension $d - 1$ and output dimension D . Note that without NN_e, NN_d , this amounts to a PCA-based coordinate reduction. The combination of PCA and neural networks combines approximate feature orthogonality and approximation flexibility.

We construct the reconstruction loss function as $|\mathbf{X} - \tilde{\mathbf{X}}|^2$. We set the reconstruction error of PCA as $E_{\text{pca}} = |\mathbf{X} - P_d^\dagger P_d \mathbf{X}|^2$. To make the reconstruction error near but less than the reconstruction error of PCA alone, we add the regularization term $\text{ReLU}(\log |\mathbf{X} - \tilde{\mathbf{X}}|^2 - \log(E_{\text{pca}}))$ in the loss function, where $E_{\text{pca}} = |\mathbf{X} - P_d^\dagger P_d \mathbf{X}|^2$ is the reconstruction error of PCA and ReLU is the rectified linear unit, that is $\text{ReLU}(u) = \max(0, u)$. Thus, the combined reconstruction loss function is

$$\text{loss}_{\text{rec}} = |\mathbf{X} - \tilde{\mathbf{X}}|^2 + \rho_1 \text{ReLU}(\log |\mathbf{X} - \tilde{\mathbf{X}}|^2 - \log |\mathbf{X} - P_d^\dagger P_d \mathbf{X}|^2),$$

where ρ_1 is a regularization parameter. The regularization term penalizes the loss if we observe a reconstruction error that is higher than PCA.

Training algorithm based on maximum likelihood estimation. After constructing the structure of the drift term $f(\mathbf{Z}(t)) = -(M(\mathbf{Z}(t)) + W(\mathbf{Z}(t))) \nabla V(\mathbf{Z}(t))$ and diffusion term $\sigma(\mathbf{Z}(t))$, we consider how to construct the loss function to learn the stochastic dynamics. In deterministic dynamical systems, we can use the mean square error to learn f given the trajectory observation data. However, to deal with stochastic dynamics (in particular, learning the diffusion matrix σ), we have to devise more general methods based on maximum likelihood estimation.

We discretize equation (5) by the Euler-Maruyama scheme, giving

$$\mathbf{Z}(t_{i+1}) = \mathbf{Z}(t_i) + f(\mathbf{Z}(t_i))\Delta t + \sigma(\mathbf{Z}(t_i))\sqrt{\Delta t}\xi_i,$$

where Δt is the time step, $t_i = i\Delta t$, $i = 0, 1, \dots, N_t - 1$ and $T = N_t\Delta t$. Here, $\{\xi_i^j\}_{i=0}^{N_t-1}$ are independent random vectors following the standard normal distribution.

In the training dataset, we have access to the microscopic coordinates $\mathbf{X}(t)$, from which we construct the reduced coordinates $\{(Z(t_i)^{(j)}, Z(t_{i+1})^{(j)}, \Delta t)\}_{i,j=0,1}^{N_t-1, M}$ via the (to be trained) reduction function ϕ . Given the neural networks f_θ and σ_θ (we use the subscript θ to denote all trainable parameters) constructed previously, the conditional probability is given by

$$\begin{aligned} & p(Z(t_{i+1})^{(j)} | Z(t_i)^{(j)}) \\ &= \mathcal{N}(Z(t_{i+1})^{(j)} | Z(t_i)^{(j)} + f_\theta(Z(t_i)^{(j)})\Delta t, \sigma_\theta(Z(t_i)^{(j)})\sigma_\theta^T(Z(t_i)^{(j)})\Delta t) \\ &= \frac{1}{(2\pi\Delta t)^{d/2} \sqrt{\det(\sigma_\theta\sigma_\theta^T)}} \exp\left\{-\frac{1}{2\Delta t} [Z(t_{i+1})^{(j)} - Z(t_i)^{(j)} - f_\theta(Z(t_i)^{(j)})\Delta t]^T (\sigma_\theta\sigma_\theta^T)^{-1} [Z(t_{i+1})^{(j)} - Z(t_i)^{(j)} - f_\theta(Z(t_i)^{(j)})\Delta t]\right\}, \end{aligned}$$

where we use the short form $\sigma_\theta = \sigma_\theta(Z(t_i)^{(j)})$ and \det denotes the determinant of a matrix.

Taking the logarithm of the above equation, we obtain

$$\begin{aligned} & \log p(Z(t_{i+1})^{(j)} | Z(t_i)^{(j)}) \\ &= -\frac{1}{2} \log \det(\sigma_\theta\sigma_\theta^T) - \frac{\Delta t}{2} \left(\frac{Z(t_{i+1})^{(j)} - Z(t_i)^{(j)} - f_\theta(Z(t_i)^{(j)})}{\Delta t} \right)^T (\sigma_\theta\sigma_\theta^T)^{-1} \left(\frac{Z(t_{i+1})^{(j)} - Z(t_i)^{(j)} - f_\theta(Z(t_i)^{(j)})}{\Delta t} \right) + \text{constant}. \end{aligned}$$

As a result, we may obtain the loss function by maximizing the log-likelihood of the training data

$$\begin{aligned} & \text{loss}_{\text{MLE}} \\ &= \frac{1}{N_t M} \sum_{i=1}^{N_t} \sum_{j=1}^M \left(\frac{1}{2} \log \det(\sigma_\theta\sigma_\theta^T) + \frac{\Delta t}{2} \left(\frac{Z(t_{i+1})^{(j)} - Z(t_i)^{(j)} - f_\theta(Z(t_i)^{(j)})}{\Delta t} \right)^T (\sigma_\theta\sigma_\theta^T)^{-1} \left(\frac{Z(t_{i+1})^{(j)} - Z(t_i)^{(j)} - f_\theta(Z(t_i)^{(j)})}{\Delta t} \right) \right). \end{aligned}$$

The total loss is then

$$\text{loss} = \text{loss}_{\text{MLE}} + \rho \text{loss}_{\text{Rec}}, \tag{19}$$

where ρ is a parameter to balance the accuracy of the learned dynamics and the error from reconstruction. In our computation, we first train the loss function (19) with $\rho = 0.01$. After some training steps, we fix the encoder part (17) and the decoder part (18) of the neural network, and train loss_{MLE} ($\rho = 0$) to fine-tune accuracy of the stochastic dynamics. We use the Adam optimizer for training. The overall implementation, including the network architectures and loss computation, is shown in Supplementary Fig. 2.

Data preparation

Simulation data. We used a Brownian dynamics approach to simulate linear, touching-bead chains as polymer chains in a planar elongational flow (Fig. 2, left). Each polymer chain consisted of $N = 300$ ($D = 3N$) beads with diameter r at positions of bead i \mathbf{r}_i , connected by $N - 1$ rigid rods of length $b = r = 10$ nm. The governing stochastic differential equation was obtained by considering the following forces acting on the system: excluded volume, constraint, Brownian and hydrodynamic.

The excluded volume potential characterizes the short-range repulsions between beads and can be described by

$$E^{\text{ev}} = - \sum_{i,j} \mu r_{ij} \quad \text{if } r_{ij} < r \text{ nm}$$

where $\mu = 35$ pN has been demonstrated to result in a low frequency of chain crossings⁴⁷. The constraint force is given by

$$\mathbf{F}_i^c = T_i \mathbf{b}_i - T_{i-1} \mathbf{b}_{i-1}$$

where \mathbf{b}_i is the unit vector of bond i and T_i is the tension in rod i that imposes constant bond length. The Brownian forces are random forces that satisfy the fluctuation-dissipation theorem, represented as

$$\langle \mathbf{F}_i^{\text{br}}(t) \rangle = 0 \quad \text{and} \quad \langle \mathbf{F}_i^{\text{br}}(t) \mathbf{F}_j^{\text{br}}(t) \rangle = \frac{2k_B T \zeta \mathbf{I} \delta_{ij}}{\Delta t}$$

where δ_{ij} is the Kronecker delta, \mathbf{I} is the identity matrix and Δt is the simulation time step. By neglecting hydrodynamic interactions between the beads, the drag force on the i th bead is

$$\mathbf{F}_i^d = -\zeta \left(\mathbf{u}(\mathbf{r}_i) - \frac{d\mathbf{r}_i}{dt} \right)$$

where $\zeta \approx 3\pi\eta r$ is the drag coefficient of a single bead, η is the solvent viscosity and $\mathbf{u}(\mathbf{r}_i)$ is the unperturbed solvent velocity.

Due to the small mass of the beads, it is common to neglect inertial effects and set the sum of forces on the beads to be zero. Hence, the Langevin equation that describes the motion of each bead along the chain is

$$\frac{d\mathbf{r}_i}{dt} = \mathbf{u}(\mathbf{r}_i) + \frac{1}{\zeta} (\mathbf{F}^{\text{ev}} + \mathbf{F}^c + \mathbf{F}^{\text{br}})$$

We employed a predictor-corrector scheme⁴⁸ to determine the position of each bead at every time step. The enforcement of rigid rod constraints leads to a system of nonlinear equations for the rod tensions T_i , which we solved for using Newton's method.

For each simulation run, the polymer chain was allowed to equilibrate for $10^4 \tau_d$, with $\tau_d = b^2 \zeta / k_B T$ being the characteristic rod diffusion time. During equilibration, the chain would adopt random configurations as governed by thermal fluctuations. At $t = 0$, the chain was subjected to a planar elongational flow of the form $\mathbf{u}(\mathbf{r}_i) = \dot{\epsilon}(\hat{\mathbf{x}} - \hat{\mathbf{y}}) \cdot \mathbf{r}_i$, where $\dot{\epsilon}$ is the strain rate and $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ are unit vectors parallel to the x and y axes, respectively. The simulations were run until $t = 10^4 \tau_d$, using a time step of $\Delta t = 5 \times 10^{-4} \tau_d$. To generate training data, we simulated 610 stretching trajectories. To test the predictions, we simulated 500 stretching trajectories each for 3 different initial chain configurations, which were deliberately selected from 3 vastly different trajectories. For each trajectory, we obtain the (x, y, z) positional coordinates \mathbf{r}_i of N beads every $10 \tau_d$. Given the observation data, we can get the chain extension (Fig. 2). We note that time reported hereon is in units of τ_d . Although the dataset in this work is generated based on known equations, it should be highlighted that our machine learning approach for constructing the reduced dynamical model and all resulting consequences are independent of the microscopic model used in the simulations, with only the positional coordinates as inputs into the S-OnsagerNet. In other words, the approach used is purely data driven and can therefore be generally applied to other non-equilibrium problems.

Experimental data. We provide the details of the experimental validation of the S-OnsagerNet results.

Experiments on electrokinetic stretching of DNA. To collect the experimental data leading to Fig. 6, we needed to create an automated single-molecule stretching trap. We describe the essential features of the experiments, while the reader can find additional details about the trap and the material preparation methodology in ref. 49. The polymer samples used for this study were T4 phage double-stranded DNA (165.6 kbp, Nippon Gene), chosen for high monodispersity and ready

availability. The DNA was diluted in buffer solution and fluorescently labeled (YOYO-1) to aid visualization.

The trap was based on a microfluidic cross-slot channel device with a wide central chamber, splitting into 40- μm -wide channels in each of the four cardinal directions (Supplementary Fig. 3). Each of the four channels was terminated by a macroscopic reservoir, in which we pipetted the DNA sample and inserted platinum wire electrodes. These electrodes were connected to a computer-controlled analog voltage source, such that the north and south electrodes were grounded, but east and west were positively biased $V_0 = +30\text{ V}$. This electric quadrupole arrangement set up a potential well in the north–south direction and a potential hill in the east–west direction. The saddle point was located in the central chamber.

In aqueous solution, DNA is naturally negatively charged, and in the microfluidic device, molecules drifted electrokinetically from the north and south reservoirs towards east and west. At the saddle point, the east–west potential hill could be exploited to stretch a DNA molecule, but being an unstable equilibrium point, molecules would approach its location and slow down, but could not remain there for long observation times without external intervention.

To actively trap and observe single DNA molecules at the saddle point, the microfluidic electrokinetic device was placed on an inverted fluorescence microscope (Nikon Eclipse Ti2U) with a $\times 60$ oil-immersion objective (1.4 numerical aperture). A live fluorescence image was captured by a scientific complementary metal–oxide–semiconductor camera (Teledyne Photometrics Prime 95b) at 50 ms intervals and sent to a desktop computer that analyzed the incoming images in real time. The image analysis included a fast clean-up step (detailed in the next section) that removed background noise and stray ‘passer-by’ molecules/fragments, followed by a step to calculate the intensity centroid of the target molecule and its projected length. The displacement x of the molecule’s centroid from the saddle point was input to a proportional feedback loop that output a voltage tilt $\Delta V \propto x$, which was then superimposed on the east–west electrode biases, such that $\Delta V = V_0 \pm Gx$. Setting $G = 2.2\text{ V } \mu\text{m}^{-1}$ confined the DNA centroid to within 1 μm of the saddle point even while a molecule stretched. This feedback process for tracking and trapping DNA molecules was automated through a custom LabVIEW program.

In a single stretch experiment, the platform was programmed to actively search for a molecule to trap as they flowed through the central chamber. Once a molecule was trapped, V_0 was set to zero temporarily to allow the molecule to relax into an unstretched, equilibrium state over 10 s (chosen based on experience from preliminary experiments). After this relaxation period, recording began with images being streamed to the computer’s solid-state drive. With every image, the associated centroid coordinates, projected length, voltages and other parameters were logged in real time. V_0 was then reset to +30 V, which re-established the east–west potential hill and stretched the DNA molecule. By monitoring the projected length history, the platform could recognize when the molecule was fully stretched, and in response stop the recording and release the molecule to escape naturally towards east or west. Using this protocol, we were able to capture a diverse set of stretching trajectories, including the dumbbell and folded conformations that were selected for analysis in Fig. 6.

Real-time image processing. A DNA molecule deforms over time, and the purpose of this section is to elucidate our method used to extract sequential image snaps that can be used to capture the DNA unfolding. To capture the exact location of the DNA molecule, we devised an algorithm to classify pixels according to the density of a pixel’s immediate surrounding. The algorithm exploited the fact that the targeted molecule would very likely be centered in the tracking region of interest (ROI), and this was used to discriminate against stray particles during tracking. To make this algorithm fast enough for real time, the algorithm operated only on the binarized version of the raw image, after application of a threshold. The result was a binary image containing the targeted DNA pixels at the exclusion of pixels belonging

to noise speckles and stray particles. See Supplementary Fig. 4 as an illustrated example of the real-time image processing pipeline.

One would imagine that calculating the intensity centroid is a natural method for defining the molecule location and tracking it across image frames in time. However, doing so with the raw DNA images is problematic for the two following reasons.

First, background noise from the camera biases the centroid towards the center of the frame. This occurred even though we subtracted a pre-recorded background frame. The problematic noise in this case arises from pixel shot noise (electronic, photon) and pixel readout, and manifests as low-intensity sparkles in the image. One possible solution is to increase the excitation intensity, but this caused the DNA to photocleave much more readily, too often breaking into fragments before stretching out fully.

Second, while a trapped DNA molecule was being stretched, stray molecules continued to drift into the camera’s field of view, which biased the centroid towards these stray molecules. This problem was partially solved by calculating the centroid from a smaller ROI that was just large enough to cover the fluctuating motion of a single molecule. This ROI was software based and centered on the molecule centroid for tracking, that is implemented as a subset of the image array. The reduced ROI size avoided most of the stray particles, but it was still common for some to enter the tracking ROI during a trap–stretch cycle.

Data selection. Among the stretching trajectories collected, a few were selected for further analysis. Specifically, examples of molecules adopting the ‘dumbbell’ and ‘folded’ conformations during the stretching process with similar initial extension lengths were chosen to benchmark to human labeling. The trajectories of the conformations as the molecules were stretched were plotted in the reduced coordinates space (Fig. 6). When in the fully stretched stable state, the DNA molecules still undergo conformational fluctuations due to Brownian motion, which result in fluctuations in the reduced coordinate space. Images of molecules in this stable state were analyzed to obtain the fluctuations in Z_1 , Z_2 and Z_3 , as plotted in Fig. 6. The method to extract the learned thermodynamic coordinates Z_2 , Z_3 from experimental images are described next.

Extracting reduced coordinates from filtered images. Different from the simulated data described earlier, given a filtered image, it is more challenging to obtain in a robust way a sequence of ordered coordinates representing the location of each part of the DNA molecule. However, identifying the two end points is easier. Thus, we first find the coordinates of the two end points and also the center of mass (weighted by intensity) of all illuminated pixels (Supplementary Fig. 5). Then, we estimate Z_2 , Z_3 by computing the end-to-end distance and foldedness as defined in Fig. 3, which are shown to strongly correlate with Z_2 and Z_3 , respectively. These computations require only the relative position vectors of the two end points with respect to the center of mass ($\mathbf{r}_1 = (r_{1,x}, r_{1,y}, 0)$ and $\mathbf{r}_N = (r_{N,x}, r_{N,y}, 0)$). The third coordinate is set to zero because the DNA molecule is confined to have limited motion in this direction. The center of mass \mathbf{r}_{cm} is computed by

$$\mathbf{r}_{\text{cm}} = \frac{1}{\bar{I}} \sum_{m,n} u_{m,n} I_{m,n},$$

where $\bar{I} = \sum_{m,n} I_{m,n}$, $u_{m,n}$ is the spatial position of the (m, n) -pixel and $I_{m,n}$ is the intensity of the image at pixel (m, n) . This allows for the computation of \mathbf{r}_1 and \mathbf{r}_N (Supplementary Fig. 5). According to the definition of end-to-end distance and foldedness and their linear correlation with Z_2 , Z_3 (Fig. 3), we may compute the reduced coordinates as

$$\begin{aligned} Z_1 &= C_1 L, \\ Z_2 &= C_1 C_2 (r_{1,x} - r_{N,x}), \\ Z_3 &= C_1 [C_3 (r_{1,x} + r_{N,x}) + C_4], \end{aligned}$$

where L is the extension length scale of the experimental data, and the parameters C_1 , C_2 , C_3 and C_4 are scaling parameters to account for the change in molecule configuration properties (for example, length scales) from simulation to experimental data. In particular

- C_1 is obtained by dividing the average extension of unfolding simulated data by that of experimental data.
- C_2 is the relationship between Z_2 and end-to-end distance (Fig. 3b) and is obtained by dividing the average value of Z_2 in simulated data by that of the end-to-end distance.
- C_3 and C_4 is obtained by the relationship between the Z_3 and fold- edness. We use least squares method to get C_3 and C_4 (Fig. 3e).

The method of extraction differs in the plot of Fig. 6j. Here we consider only the stretched state, for which we can extract 300 coordinates (assuming the third coordinate $z = 0$) equally spaced along the stretched polymer. A scaling along the x axis is performed so that the average full extension of the experimental polymer data matches that of the simulation data. These coordinates are then fed into the trained PCA-ResNet to extract the Z_2 and Z_3 values.

Polymer dynamics analysis

In this section, we detail the modeling of the polymer unfolding problem using the S-OnsagerNet.

Accurate prediction of the statistics of unfolding. In our training process, we have 610 training trajectories and 110 testing trajectories. Note that although the dataset is generated by known yet complex microscopic equations, our approach does not require, nor rely on, the knowledge of these equations. The chain extension evolution of the training data with different initial chain are shown in Supplementary Fig. 6a (black) and the test data are shown in Supplementary Fig. 6e (black). After training the model, we predict the extension of the polymer (red). We can see the extension of the polymer can be predicted well. We also compute the mean (Supplementary Fig. 6b,f), standard derivation (Supplementary Fig. 6c,g) and the distribution of unfolding time (Supplementary Fig. 6d,h) of the training and test results. We also compute the error of the mean (relative L^2 error), standard derivation and the probability distribution of unfolding time of the training and test results in Supplementary Table 1. We observe that our model successfully captures the statistical behavior of a polymer stretching with only a three-variable dynamical system.

Interpreting the learned closure coordinates. Supplementary Fig. 7 shows the evolution of chain extension and the two learned closure coordinates with time for the training data. The trajectories are colored by unfolding time. Based on the unfolding times, we observe that chains with similar initial extensions (determined by the y intercept, that is, Z_1 at $t = 0$) can take vastly different times to stretch; hence, it is not sufficient to consider only chain extension for the purposes of predicting dynamics. However, the successful prediction of the statistics of unfolding dynamics implies that Z_2 and Z_3 capture crucial information of the system. Thus, we seek to gain some physical understanding of the learned closure coordinates Z_2 and Z_3 .

Here we provide details on the analysis of the closure coordinates (Z_2, Z_3) that characterize the stochastic evolution of the extension length (Z_1). Recall that Z_2, Z_3 are deterministic functions of the microscopic configuration \mathbf{X} , and the functions are approximated by a trained PCA-ResNet. A useful property we can exploit is differentiability of the neural network. We can ask: given a certain configuration \mathbf{X} , what kind of small perturbations to \mathbf{X} will most drastically increase or decrease the value of Z_2 and Z_3 ? In other words, we can consider perturbations in the directions of $\pm \partial Z_2 / \partial \mathbf{X}$ and $\pm \partial Z_3 / \partial \mathbf{X}$, respectively. We first analyze the closure coordinate Z_2 . From Fig. 3a, we observe that perturbations in the direction of Z_2 tend to change the end-to-end distance in the elongational axis (distance between the first and the

last bead in the polymer chain along the elongational direction, that is $|r_{N,x} - r_{1,x}|$, where $r_{j,i}$ is the i th coordinate of the j th bead in the chain (cyan, $Z_2 = 0.453$; blue, $Z_2 = 0.194$; black, $Z_2 = 0.323$). We confirm this hypothesis by visualizing the correlation of the end-to-end distance and the magnitude of Z_2 in Fig. 3b,c. In general, we observe that as $|Z_2|$ decreases, the distance between the chain ends (marked by red points in the figure) decreases. Thus, we can interpret the first learned closure coordinate as an indicator of end-to-end distance.

We proceed with a similar analysis for the other closure coordinate Z_3 . Figure 3d shows a given chain configuration and perturbations in the positive and negative directions of $\frac{\partial Z_3}{\partial \mathbf{X}}$ (cyan, $Z_3 = 7.903$; blue, $Z_3 = 1.595$; black, $Z_3 = 4.749$). Here we observe that the end-to-end distance is largely unchanged, but the degree of foldedness of the chain in the elongational axis of the flow (x direction) appears to change. This leads us to hypothesize that the second learned coordinate represents a degree of foldedness with respect to the elongational flow. As a measure of the degree of foldedness of a chain, we compute $|r_{1,x} + r_{N,x}|$. During data pre-processing, the chain is centered such that its center of mass is 0. Hence, if $|r_{1,x} + r_{N,x}|$ is small, the polymer is symmetric around zero in the elongational x direction and tends to be in the elongated state. If $|r_{1,x} + r_{N,x}|$ is large, the polymer is likely to be in the folded state. We plot $|r_{1,x} + r_{N,x}|$ as a function of $|Z_3|$ for all configurations in the training dataset in Fig. 3e. The strong correlation between $|r_{1,x} + r_{N,x}|$ and $|Z_3|$ supports our interpretation that the second learned closure coordinate is an indicator of the degree of foldedness in the elongational direction. To demonstrate this, we plot in Fig. 3f visualizations of different chains with a range of $|Z_3|$ values, with the chain ends marked by red points. As $|Z_3|$ decreases, we observe that the chains generally shift from the folded to the elongated state. We note that the degree of foldedness is sufficiently described solely by considering the projection of chain coordinates onto the elongational axis of the flow, as the flow is stable in the compressional axis and thus the degree of foldedness is primarily relevant to the unstable elongational axis that drives the unfolding process.

Advancing classification methods for polymer stretching. With the new understanding of the closure variables, we now consider how our results improve the current understanding of polymer chain dynamics. In the landmark experimental study of dilute polymer chains stretching under elongational flow, the molecules were categorized into seven different conformations and the dynamics of dominant conformations were analyzed²⁵. Specifically, it was found that chains in the ‘folded’ conformation (which is one of the seven categories) took the longest time to stretch, while chains in the dumbbell conformation stretched relatively quickly (Supplementary Fig. 8).

Our analysis shows that instead of a categorical labeling, it is perhaps more useful to characterize the stretching dynamics of a polymer by three numbers representing the generalized coordinates: Z_1 (extension length), Z_2 (related to end-to-end distance) and Z_3 (related to foldedness in the flow direction). Our results show that these are sufficient to predict the dynamics of the chain extension. We show that this characterization is largely consistent with previous categorical ones, but improves on them in some intermediate cases. In Supplementary Fig. 9, we plot the values of the low-dimensional coordinates $|Z_2|$ and $|Z_3|$ of different chains with initial dumbbell and folded configurations at selected chain extension Z_1 values, colored by the predicted unfolding times. We observe segregation between the folded and dumbbell configurations in the Z_2 - Z_3 space, indicating that the qualitative differences between different conformations can be captured by our characterization. In general, the folded chains take a longer time to stretch compared with the dumbbell chains. This is consistent with experimental and computational observations reported in the literature²⁵⁻²⁷. However, we highlight that the region with high $|Z_2|$ and low $|Z_3|$ values encompasses a mix of folded and dumbbell chains with similar unfolding times. Therefore, while the broad categorization scheme allows for coarse discrimination of the stretching dynamics, the

qualitative classification does not allow for finer predictions. Instead of classifying the stretching trajectories based on qualitative, human judgement of chain conformation during the process, we present a robust, quantitative approach to interpreting the stretching dynamics that involves consideration of the initial chain configuration in reduced dimensions.

Free-energy landscape analysis. We now provide details on the analysis of the free-energy landscape. We begin with an important technical note. Our learned GSOP following equation (5) is in general not guaranteed to be a gradient system, unless $W(\mathbf{Z}) = 0$ and $M(\mathbf{Z}) = I$. However, as the drift term $f(\mathbf{Z}) = (f_1(\mathbf{Z}), f_2(\mathbf{Z}), f_3(\mathbf{Z}))^T = -(M(\mathbf{Z}) + W(\mathbf{Z}))\nabla V(\mathbf{Z})$, the stationary points of V are also critical points of the dynamics ($f(\mathbf{Z}) = 0$). The saddle points of V are the saddle foci of the non-gradient dynamics $\dot{\mathbf{Z}} = f(\mathbf{Z})$. For simplicity, we refer to a saddle point of V and saddle focus of f interchangeably.

We compute the critical points by numerically solving $\nabla V(\mathbf{Z}) = 0$ with the BFGS method from different initial conditions. We obtain four critical points: (247.15, 1.701, 0.193) (yellow star), (247.29, -1.700, -0.166) (blue square), (87.858, -0.041, -4.269) (cyan triangle) and (85.887, -0.050, 4.126) (purple diamond), which are shown in Fig. 4. The Jacobian matrix of the drift term is then computed

$$J(f) = \begin{pmatrix} \partial_{Z_1} f_1 & \partial_{Z_2} f_1 & \partial_{Z_3} f_1 \\ \partial_{Z_1} f_2 & \partial_{Z_2} f_2 & \partial_{Z_3} f_2 \\ \partial_{Z_1} f_3 & \partial_{Z_2} f_3 & \partial_{Z_3} f_3 \end{pmatrix}.$$

We calculate the eigenvalues and eigenvectors of J at the four critical points. It has three eigenvalues with negative real parts at the yellow and blue points, so these are the stable points (stable nodes). We name them $\mathbf{Z}_{\text{stable},1}$ and $\mathbf{Z}_{\text{stable},2}$. For the cyan triangle and magenta diamond points, J has one eigenvalue with positive real part, and two with negative real parts. These points are saddle points (saddle focus) of index 1. We call these two points $\mathbf{Z}_{\text{saddle},1}$ and $\mathbf{Z}_{\text{saddle},2}$. The direction of the eigenvector corresponding to the eigenvalue with positive real part is the unstable manifold, along which the trajectories escape from the saddle. On the other hand, the span of the real and imaginary parts of the other eigenvector (one must be a complex conjugate of the other) constructs the stable manifold, along which trajectories are attracted to the saddle.

The local behavior of the learned potential $V(\mathbf{Z})$ around its critical points characterize the typical fluctuations around them. There lies important physical meaning that we can probe. For example, let us exploit automatic differentiation of neural networks and expand $V(\mathbf{Z})$ in a Taylor series around $\mathbf{Z}_{\text{stable},1}$ (fully stretched state), corresponding to (247.15, 1.701, 0.193). Neglecting small terms, we obtain the approximate formula

$$\delta V \approx 153.1(\delta Z_1 - 1.546\delta Z_2)^2 + 205.5\delta Z_2^2 + 36.96\delta Z_3^2, \quad (20)$$

where $\delta Z_i = Z_i - [Z_{\text{stable},1}]_i$ is the fluctuations in the thermodynamic variables. Now, assuming that the distribution of states around $\mathbf{Z}_{\text{stable}}$ follows a Boltzmann distribution $Z \sim \exp[-V(\mathbf{Z})/k_B T]$ (here we assume that $M \approx I$) where $|\sigma|^2 \propto k_B T$, the typical small fluctuations δV is proportional to $k_B T$. In other words, equation (20) is approximate ‘isotherms’ that captures the form of typical fluctuations. For example, we can infer from the formula that typical fluctuations of the extension length (Z_1) and the end-to-end distance (Z_2) are highly correlated. This is sensible, as in the fully stretched state, these two quantities are expected to change simultaneously. In Fig. 4g–i, we confirm these correlations.

Similarly, one can also expand $V(\mathbf{Z})$ around a saddle point (fully folded state) $\mathbf{Z}_{\text{saddle},2}$. We obtain the formula

$$\delta V \approx 102.96\delta Z_1^2 - 31.13(\delta Z_2 - 0.255\delta Z_3)^2 + 24.16\delta Z_3^2,$$

Here we immediately observe that to escape the saddle point (lowering the energy), one should increase end-to-end distance (Z_2) and decrease foldedness (Z_3). This approximately aligns with the unstable manifold described above, and forms the basis of our control protocols described in the main text.

Selection of the dimension of the reduced coordinates. In this part, we describe how we arrived at the selection of a three-variable reduced coordinate space. We tested various different numbers of reduction dimensions ($d = 2, 3, 4$) and the relative errors of predictions are summarized in Supplementary Table 1. We observed that going to a higher dimension ($d = 4$) did not result in noticeable gains. In fact, increasing dimension may cause increasing optimization error and hence worse results in standard deviation. Going to a lower dimension ($d = 2$) resulted in increased prediction error in general.

Interestingly, we can formulate a physical argument that suggests that a two-dimensional system (with only one additional closure coordinate) is not a suitable reduced model for the polymer dynamics we study. The argument is based on index theory for two-dimensional dynamical systems⁵⁰.

Let us assume for the sake of contradiction that the projected free-energy landscape of our learned three-dimensional system into Z_1 - Z_2 plane is the free energy of a system in two-dimensional (Supplementary Fig. 10). As Z_1 is the polymer extension, we expect that

- (1) There exists two stable states at large Z_1 corresponding to the fully extended state. There are two of them due to reflection symmetry in the flow direction. Around this Z_1 value, all trajectories in the reduced space should converge to one of the stable steady states.
- (2) There cannot be saddle points with the same Z_1 value, as it is close to the maximal extended chain length.

These conditions are enough to imply a contradiction in the following way. The simple concept we use from index theory is the definition of the index of a closed curve in the phase space of a two-dimensional dynamical system (see Ch. 6.8 in ref. 50 for details). Let Γ be a closed curve in \mathbb{R}^2 , and consider a dynamical system $\dot{\mathbf{z}} = f(\mathbf{z}) = (f_1(\mathbf{z}), f_2(\mathbf{z}))^T$, $\mathbf{z} = (z_1, z_2)^T \in \mathbb{R}^2$. The index of Γ with respect to the dynamics f is defined as

$$I_\Gamma(f) := \frac{1}{2\pi} \oint_\Gamma \frac{f_1 df_2 - f_2 df_1}{f_1^2 + f_2^2}. \quad (21)$$

Intuitively, this is the sum of the angles of the force vectors across the curve Γ .

From index theory, we know that the index of a closed curve must equal to the sum of indices of critical points it encloses. Moreover, the index of a stable critical point is +1, and the index of a saddle is -1. Now, we draw a curve enclosing the two stable points as shown in Supplementary Fig. 10. By condition (1) above, the vector fields should point inwards towards the interior of the curve, thus we can show (from equation (21)) that the index of this curve is +1. However, it can only enclose stable critical points due to condition (2), and consequently its index is at least +2. Thus we arrive at a contradiction, showing that the dynamical landscape in Supplementary Fig. 10 is not possible in two dimensions.

The only remaining possibility is that Z_2 does not distinguish the two symmetric stable extended states. However, in this case to allow a saddle point (which is the key to inducing heterogeneity in unfolding), we must have another stable steady state at a different Z_1 value. This is unlikely from a physical viewpoint for non-self-interacting polymers, as we expect the only asymptotically stable state to be the stretched, fully extended state.

The impact of dataset size on predictive accuracy of S-OnsagerNet

To study the impact of the number of training data on the computational accuracy of the model, the original training dataset containing 610 trajectories was split into two subsets containing 25% and 75% of the data. A third subset was produced by selecting at random 50% of the data from the original dataset. The three datasets and the learned potential landscapes are presented in Supplementary Fig. 11. The 25% and 75% datasets have no common trajectories, while the 50% dataset contains some trajectories from each of the other two datasets. The quantitative results reported in Supplementary Table 2, where the trained models are used to predict 500 unseen trajectories with fast, medium and slow unfolding times, suggest that with a smaller amount of training data, the S-OnsagerNet model loses some predictive accuracy. However, the more important factor is the diversity of the data contained in the training datasets. The full dataset has a high proportion of trajectories with fast, followed by middle (slower) and then slow unfolding times; however, it is the most balanced of all the training datasets (25%, 50%, 75%, 100%), in addition to containing the most data. Training with the 25% dataset, which contains the largest proportion of trajectories with fast unfolding times results in the lowest L^2 error, but the model overfits to that type of trajectory, and has the largest L^2 error for the medium and slow unfolding times. However, the 50% dataset is the most balanced of the reduced datasets, which is reflected in the relative L^2 errors. It can also be observed in Supplementary Fig. 11 (bottom) that the potential landscapes resulting from training with smaller datasets do not contain a clear stagnation (saddle) point. Overall, for training datasets with different numbers of trajectories and their respective proportions of fast, medium and slow trajectories, the results, both in terms of prediction error and potential landscape characteristics, suggest that the model is relatively robust, and the amount of training data cannot be reduced much without affecting predictive performance.

Spatial epidemics analysis

In this section, we provide details of our analysis method for an alternative application of S-OnsagerNet—modeling the macroscopic dynamics of the spread of epidemics. This highlights the general applicability of our method.

We focus on the most well-known model for disease spread in a spatial domain—the spatial SIR model³². Let us consider a two-dimensional square domain (representing a city, say) discretized into $n \times n$ sectors. We use $I_{i,j}$ and $S_{i,j}$ to represent the number (density) of infective and susceptible individuals at spatial location (i, j) . The basic mechanism of the model is as follows: each infective individual may infect a susceptible individual in the same spatial location. At the same time, each infective individual recovers (or is removed) at a rate, after which they are no longer infective. Finally, both infective and susceptible individuals move on the spatial domain randomly. Mathematically, the temporal evolution of I, S (understood as $n \times n$ matrices, or length n^2 vectors) are governed by the following dynamics

$$\begin{aligned} \dot{S}_{i,j} &= -\beta I_{i,j} S_{i,j} + \frac{\delta}{\delta x^2} (S_{i-1,j} - 2S_{i,j} + S_{i+1,j}) \\ &\quad + \frac{\delta}{\delta y^2} (S_{i,j-1} - 2S_{i,j} + S_{i,j+1}) + \sigma \dot{B}_1(t), \\ \dot{I}_{i,j} &= \beta I_{i,j} S_{i,j} - \gamma I_{i,j} + \frac{\delta}{\delta x^2} (I_{i-1,j} - 2I_{i,j} + I_{i+1,j}) \\ &\quad + \frac{\delta}{\delta y^2} (I_{i,j-1} - 2I_{i,j} + I_{i,j+1}) + \sigma \dot{B}_2(t). \end{aligned} \tag{22}$$

As usual, the dot denotes time derivative. The parameter β is a measure of the transmission efficiency of the disease from infectives to susceptibles, and $1/\gamma$ is the life expectancy (or expected recovery time) of an infective. The constant δ is the diffusion coefficient, and this

term in the equation models the spatial movement of individuals as a diffusion process over the domain. The last terms of the equation models the stochastic fluctuations of the number densities, with σ as the noise intensity. The parameters δ_x and δ_y are the spatial discretization sizes in the two spatial directions. In our simulations, we take $n = 16, \beta = 0.3, \gamma = 0.13, \delta = 0.5, \sigma = 0.03$ and $\delta_x = \delta_y = \frac{2}{3}$. Equation (22) governs the microscopic dynamics of disease spread, and non-trivial outcomes can result from different initial spatial configurations and parameters (for example, infection rate, recovery rate). See Extended Data Fig. 1.

While this microscopic model and its variants has been subject to intense study (see ref. 32 and references therein), a macroscopic understanding of the dynamics of disease spread is challenging due to the complex spatial interactions. For example, one may be interested to model the dynamics of average (or total) number of infective and susceptible individuals over the spatial domain. Observe in Extended Data Fig. 1 that configurations with identical initial spatial averages of infective and susceptible individuals can have drastically different subsequent evolution. More precisely, one spatial configuration Extended Data Fig. 1a can lead to initial disease spread (epidemic), where the mean number of infected individuals initially increases sharply, whereas another spatial configuration in Extended Data Fig. 1b leads to the disease dying out monotonically. Thus, it is of interest to develop a thermodynamic description to capture and elucidate the driving factors of such variations.

Modeling the thermodynamics of the spatial SIR model using S-OnsagerNet. Following our framework, we now set the macroscopic variables of interest as the respective spatial averages $Z_1 = \delta_x \delta_y \sum_{i,j} I_{i,j}$ and $Z_2 = \delta_x \delta_y \sum_{i,j} S_{i,j}$. Recall from Extended Data Fig. 1 that these variables alone are insufficient to determine their subsequent evolution. Our goal is to learn closure variable(s), and a stochastic dynamics that describes the evolution of these variables.

Training data are generated through integrating equation (22) using the Euler–Maruyama method with time step size $dt = 0.03$. The initial conditions are selected as $I_{i,j}(0) = 5e^{-(x_i-x_0)^2-(y_j-y_0)^2}$ and $S_{i,j}(0) = 5e^{-(x_i-1)^2-(y_j-1)^2} + 5e^{-(x_i-1)^2-(y_j+2)^2}$, where x_0 and y_0 are randomly generated from a uniform distribution in the unit square, and $x_i = -5 + i\delta_x$ and $y_j = -5 + j\delta_y$. This initial condition corresponds to the scenario where the initial susceptible population is fixed, but the initial infective population is a cluster that is uniformly and randomly distributed in the domain. Then we carry out the S-OnsagerNet workflow as shown in Fig. 1, with $\beta = 0.01, \alpha = 0.001, m = 50$ and $d = 3$ (that is, one closure coordinate). We employ a PCA–encoder network to obtain the closure coordinate. The training process involves initially training the PCA–encoder and the S-OnsagerNet simultaneously. Subsequently, we fix the former and continue training the latter to a desired error tolerance. Note that in the SIR model case, we do not utilize the decoder network as there is no need to obtain a reconstruction of I and S .

Capturing the stochastic dynamics. First, we show that with just one additional learned closure coordinate Z_3 , we can capture the statistics of the macroscopic dynamics of the spatial averages of infective and susceptible individuals. The results are shown in Extended Data Fig. 1d,e, where the true mean and standard deviation of Z_1 and Z_2 are obtained from equation (22), while the predicted results are derived from S-OnsagerNet. Four representative test initial conditions are shown: two with disease spread and the other two without. We observe that we can successfully capture the macroscopic dynamics of disease spread with just one additional closure coordinate.

Interpreting the closure coordinate. Next, as we have shown that only one closure coordinate is required for a thermodynamic description, it is natural to probe its physical meaning. We use the same technique described in the polymer stretching case, where we investigate

the effect of perturbation of a microscopic state in the directions that cause the sharpest changes in Z_3 . The microscopic state we probe is chosen as a pair of partially overlapping clusters of infective and susceptible individuals (Extended Data Fig. 2a). We observe that the perturbations induced by $dZ_3/d\mathbf{X}$ (where $X = (I, S)$) correspond to increasing/decreasing the overlap of clusters of susceptible and infected individuals. Hence, this suggests that Z_3 is a macroscopic descriptor that correlates with such effective spatial overlap. We confirm this hypothesis via a scatterplot in Extended Data Fig. 2b, where the overlap is defined by $I_{S_{\text{mean}}} = \delta_x \delta_y \sum_{i,j}^n I_{i,j} S_{i,j}$. Thus, we can interpret the closure coordinate Z_3 as an indicator of the overlap of clusters of susceptible and infective individuals. This is physically sensible, as the measure of overlap of spatial clusters should determine the outcome of an epidemic. Nevertheless, we must emphasize that the learned Z_3 is a quantitative measure and can be applied to more complex configurations than a pair of clusters, for which one may not be able to easily define a notion of effective overlap by empirical observation.

Free-energy landscape. Finally, we study the dynamical landscape of the learned S-OnsagerNet model. Using the same projection technique in the polymer case, we plot two-dimensional projections of the learned three-dimensional free-energy landscape in Extended Data Fig. 3, overlaid with two representative trajectories (with disease spread in blue and disease dying out in red).

A number of interesting features can be gleaned from the landscape. First, we can clearly see the origins of the divergence of the two different types of trajectory. While they have identical initial Z_1 (average infective) and Z_2 (average susceptible) values, their initial Z_3 (infective/susceptible overlap) values differ (Extended Data Fig. 3b,c). In particular, the initial disease spread seen in the red and blue trajectories is approximately in accordance to the steepest descent of the energy landscape. That is, the initial Z_3 value is a determining factor for the onset of epidemics. Second, we compute using the steady states of the dynamics corresponding by solving $\nabla V(\mathbf{Z}) = 0$. Instead of isolated steady states as in the case of polymer dynamics, we find a one-dimensional manifold of stable steady states in the Z_2 - Z_3 plane (Extended Data Fig. 3c,f). This implies, in particular, that the terminal state (the remaining number of susceptible individuals) is not unique, but rather depends on the initial configuration, and in particular on the initial overlap described by the learned coordinate Z_3 . This rationalizes the observed heterogeneity in the terminal configurations as shown in Extended Data Fig. 1.

Statistics and reproducibility

The train–test splits in this paper are performed by random uniform subsampling. Repeats of numerical experiments are performed by running the same code with different random seeds. No data were excluded from the analyses, and the investigators were not blinded to allocation during experiments and outcome assessment.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

The simulation and experimental datasets used are publicly available in the Harvard Dataverse public repository⁵¹. The simulation data were generated according to the methods introduced in ‘Data preparation’ in Methods. Source data are provided with this paper.

Code availability

The code to reproduce the analysis generated within the study is provided at <https://github.com/MLDS-NUS/DeepLearningCustomThermodynamics> ref. 52.

References

- Noh, J., Gu, G. H., Kim, S. & Jung, Y. Machine-enabled inverse design of inorganic solid materials: promises and challenges. *Chem. Sci.* **11**, 4871–4881 (2020).
- Hippalgaonkar, K. et al. Knowledge-integrated machine learning for materials: lessons from gaming and robotics. *Nat. Rev. Mater.* **8**, 241–260 (2023).
- Brunton, S. L., Proctor, J. L. & Kutz, J. N. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc. Natl Acad. Sci. USA* **113**, 3932–3937 (2016).
- Hamzi, B. & Owhadi, H. Learning dynamical systems from data: a simple cross-validation perspective, part I: parametric kernel flows. *Physica D* **421**, 132817 (2021).
- Dietrich, F. et al. Learning effective stochastic differential equations from microscopic simulations: combining stochastic numerics and deep learning. *Chaos* **33**, 023121 (2023).
- Raissi, M., Perdikaris, P. & Karniadakis, G. E. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707 (2019).
- Yu, H., Tian, X., E, W. & Li, Q. OnsagerNet: learning stable and interpretable dynamics using a generalized Onsager principle. *Phys. Rev. Fluids* **6**, 114402 (2021).
- Hesthaven, J. S., Pagliantini, C. & Ripamonti, N. Rank-adaptive structure-preserving model order reduction of Hamiltonian systems. *ESAIM: Math. Model. Numer. Anal.* **56**, 617–650 (2022).
- Valperga, R., Webster, K., Turaev, D., Klein, V. & Lamb, J. Learning reversible symplectic dynamics. In *Learning for Dynamics and Control Conference* (Eds. Firoozi, R. et al.) 906–916 (PMLR, 2022).
- Jin, P., Zhang, Z., Kevrekidis, I. G. & Karniadakis, G. E. Learning Poisson systems and trajectories of autonomous systems via Poisson neural networks. *IEEE Trans. Neural Netw. Learn. Syst.* **34**, 8271–8283 (2023).
- Lin, B., Li, Q. & Ren, W. A data driven method for computing quasipotentials. In *Mathematical and Scientific Machine Learning* (Eds. Dong, B. et al.) 652–670 (PMLR, 2022).
- Onsager, L. Reciprocal relations in irreversible processes. I. *Phys. Rev.* **37**, 405 (1931).
- Onsager, L. Reciprocal relations in irreversible processes. II. *Phys. Rev.* **38**, 2265 (1931).
- Chen, B. et al. Automated discovery of fundamental variables hidden in experimental data. *Nat. Comput. Sci.* **2**, 433–442 (2022).
- Noé, F., Tkatchenko, A., Müller, K.-R. & Clementi, C. Machine learning for molecular simulation. *Annu. Rev. Phys. Chem.* **71**, 361–390 (2020).
- Doi, M. Onsager’s variational principle in soft matter. *J. Phys. Condensed Matter* **23**, 284118 (2011).
- Doi, M. Onsager principle as a tool for approximation. *Chin. Phys. B* **24**, 020505 (2015).
- Mielke, A., Peletier, M. A. & Renger, D. R. M. On the relation between gradient flows and the large-deviation principle, with applications to Markov chains and diffusion. *Potential Anal.* **41**, 1293–1327 (2014).
- Tanaka, S., Watanabe, T. & Nagata, K. Multi-particle model of coarse-grained scalar dissipation rate with volumetric tensor in turbulence. *J. Comput. Phys.* **389**, 128–146 (2019).
- Friederich, P., Häse, F., Proppe, J. & Aspuru-Guzik, A. Machine-learned potentials for next-generation matter simulations. *Nat. Mater.* **20**, 750–761 (2021).
- Bird, R. B., Curtiss, C. F., Armstrong, R. C. & Hassager, O. *Dynamics of Polymeric Liquids, Volume 2: Kinetic Theory* (Wiley, 1987).
- Doi, M. & Edwards, S. F. *The Theory of Polymer Dynamics* Vol. 73 (Oxford Univ. Press, 1988).

23. Larson, R. G. *The Structure and Rheology of Complex Fluids* (Oxford Univ. Press, 1999).
24. McKinley, G. H. & Sridhar, T. Filament-stretching rheometry of complex fluids. *Annu. Rev. Fluid Mech.* **34**, 375–415 (2003).
25. Perkins, T. T., Smith, D. E. & Chu, S. Single polymer dynamics in an elongational flow. *Science* **276**, 2016–2021 (1997).
26. Smith, D. E. & Chu, S. Response of flexible polymers to a sudden elongational flow. *Science* **281**, 1335–1340 (1998).
27. Larson, R. G. Brownian dynamics simulations of a DNA molecule in an extensional flow field. *J. Rheol.* **43**, 267 (1999).
28. Jendrejack, R. M., De Pablo, J. J. & Graham, M. D. Hydrodynamic interactions in long chain polymers: application of the Chebyshev polynomial approximation in stochastic simulations. *J. Chem. Phys.* **116**, 2894 (2002).
29. Hsieh, C. C., Li, L. & Larson, R. G. Modeling hydrodynamic interaction in Brownian dynamics: simulations of extensional flows of dilute solutions of DNA and polystyrene. *J. Nonnewton. Fluid Mech.* **113**, 147–191 (2003).
30. Sutton, R. S. & Barto, A. G. *Reinforcement Learning: An Introduction* (MIT Press, 2018).
31. Soh, B. W., Narsimhan, V., Klotz, A. R. & Doyle, P. S. Knots modify the coil–stretch transition in linear dna polymers. *Soft Matter* **14**, 1689–1698 (2018).
32. Murray, J. D. *Mathematical Biology II: Spatial Models and Biomedical Applications* Vol. 3 (Springer, 2001).
33. Dobson, C. M. Protein folding and misfolding. *Nature* **426**, 884–890 (2003).
34. Whitesides, G. M. & Grzybowski, B. Self-assembly at all scales. *Science* **295**, 2418–2421 (2002).
35. Capito, R. M., Azevedo, H. S., Velichko, Y. S., Mata, A. & Stupp, S. I. Self-assembly of large and small molecules into hierarchically ordered sacs and membranes. *Science* **319**, 1812–1816 (2008).
36. Cipelletti, L. & Ramos, L. Slow dynamics in glassy soft matter. *J. Phys. Condensed Matter* **17**, 253 (2005).
37. Stillinger, F. H. & Debenedetti, P. G. Glass transition thermodynamics and kinetics. *Annu. Rev. Condens. Matter Phys.* **4**, 263–285 (2013).
38. Krivov, S. V. & Karplus, M. Hidden complexity of free energy surfaces for peptide (protein) folding. *Proc. Natl Acad. Sci. USA* **101**, 14766–14770 (2004).
39. Settles, B. *Active Learning Synthesis Lectures on Artificial Intelligence and Machine Learning* Vol. 6 (Morgan & Claypool, 2012).
40. Zhao, Z. & Li, Q. Adaptive sampling methods for learning dynamical systems. In *Proc. Mathematical and Scientific Machine Learning* (Eds. Dong, B. et al.) 335–350 (PMLR, 2022).
41. Nadler, W., Brünger, A. T., Schulten, K. & Karplus, M. Molecular and stochastic dynamics of proteins. *Proc. Natl Acad. Sci. USA* **84**, 7933–7937 (1987).
42. Milstein, G. N., Repin, Y. M. & Tretyakov, M. V. Symplectic integration of hamiltonian systems with additive noise. *SIAM J. Numer. Anal.* **39**, 2066–2088 (2002).
43. Goel, N. S., Maitra, S. C. & Montroll, E. W. On the Volterra and other nonlinear models of interacting populations. *Rev. Mod. Phys.* **43**, 231 (1971).
44. Tuckerman, M. *Statistical Mechanics: Theory and Molecular Simulation* (Oxford Univ. Press, 2010).
45. Beris, A. N. & Edwards, B. J. *Thermodynamics of Flowing Systems: with Internal Microstructure* Vol. 36 (Oxford Univ. Press, 1994).
46. Li, B., Tang, S. & Yu, H. Better approximations of high dimensional smooth functions by deep neural networks with rectified power units. *Commun. Comput. Phys.* **27**, 379–411 (2020).
47. Vologodskii, A. Brownian dynamics simulation of knot diffusion along a stretched DNA molecule. *Biophys. J.* **90**, 1594–1597 (2006).
48. Liu, T. W. Flexible polymer chain dynamics and rheological properties in steady flows. *J. Chem. Phys.* **90**, 5826–5842 (1989).
49. Soh, B. W., Ooi, Z.-E., Vissol-Gaudin, E., Leong, C. J. & Hippalgaonkar, K. Automated electrokinetic stretcher for manipulating nanomaterials. *Lab Chip* **23**, 3716–3726 (2023).
50. Strogatz, S. H. *Nonlinear Dynamics and Chaos: with Applications to Physics, Biology, Chemistry, and Engineering* (CRC Press, 2018).
51. Soh, B. Data for constructing custom thermodynamics using deep learning. *Harvard Dataverse* <https://doi.org/10.7910/DVN/NRIX7Y> (2023).
52. Chen, X. et al. Constructing custom thermodynamics using deep learning. *Zenodo* <https://doi.org/10.5281/zenodo.10212239> (2023).

Acknowledgements

This research is supported by the Ministry of Education, Singapore, under its Research Centre of Excellence award to the Institute for Functional Intelligent Materials (I-FIM, project no. EDUNC-33-18-279-V12 KSN). K.S.N. is grateful to the Royal Society (UK, grant number RSRP \R \190000 KSN) for support. Q.L. acknowledges support from the National Research Foundation, Singapore, under the NRF fellowship (project no. NRF-NRFF13-2021-0005 QL). H.Y. acknowledges support from the National Natural Science Foundation of China under Grant No. 12171467 HY and 12161141017 HY. K.H., B.W.S. and Z.-E.O. acknowledge support from the Accelerated Materials Development for Manufacturing Program at A*STAR via the AME Programmatic Fund by the Agency for Science, Technology and Research under Grant No. A1898b0043 KH. The computational work for this article was partially performed on resources of the National Supercomputing Centre, Singapore. The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript. We thank C. Jie Leong for help with experimental preparation.

Author contributions

X.C., B.W.S., K.S.N., K.H. and Q.L. conceptualized the project and developed the methodologies. X.C. and B.W.S. conducted the numerical experiments. B.W.S., E.V.-G. and Z.-E.O. set up and conducted the physical experiments. X.C., H.Y. and Q.L. conducted the theoretical investigations. All authors co-wrote the paper.

Competing interests

The authors declare no competing interests.

Additional information

Extended data is available for this paper at <https://doi.org/10.1038/s43588-023-00581-5>.

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s43588-023-00581-5>.

Correspondence and requests for materials should be addressed to Kostya S. Novoselov, Kedar Hippalgaonkar or Qianxiao Li.

Peer review information *Nature Computational Science* thanks the anonymous reviewer(s) for their contribution to the peer review of this work. Primary Handling Editor: Jie Pan, in collaboration with the *Nature Computational Science* team.

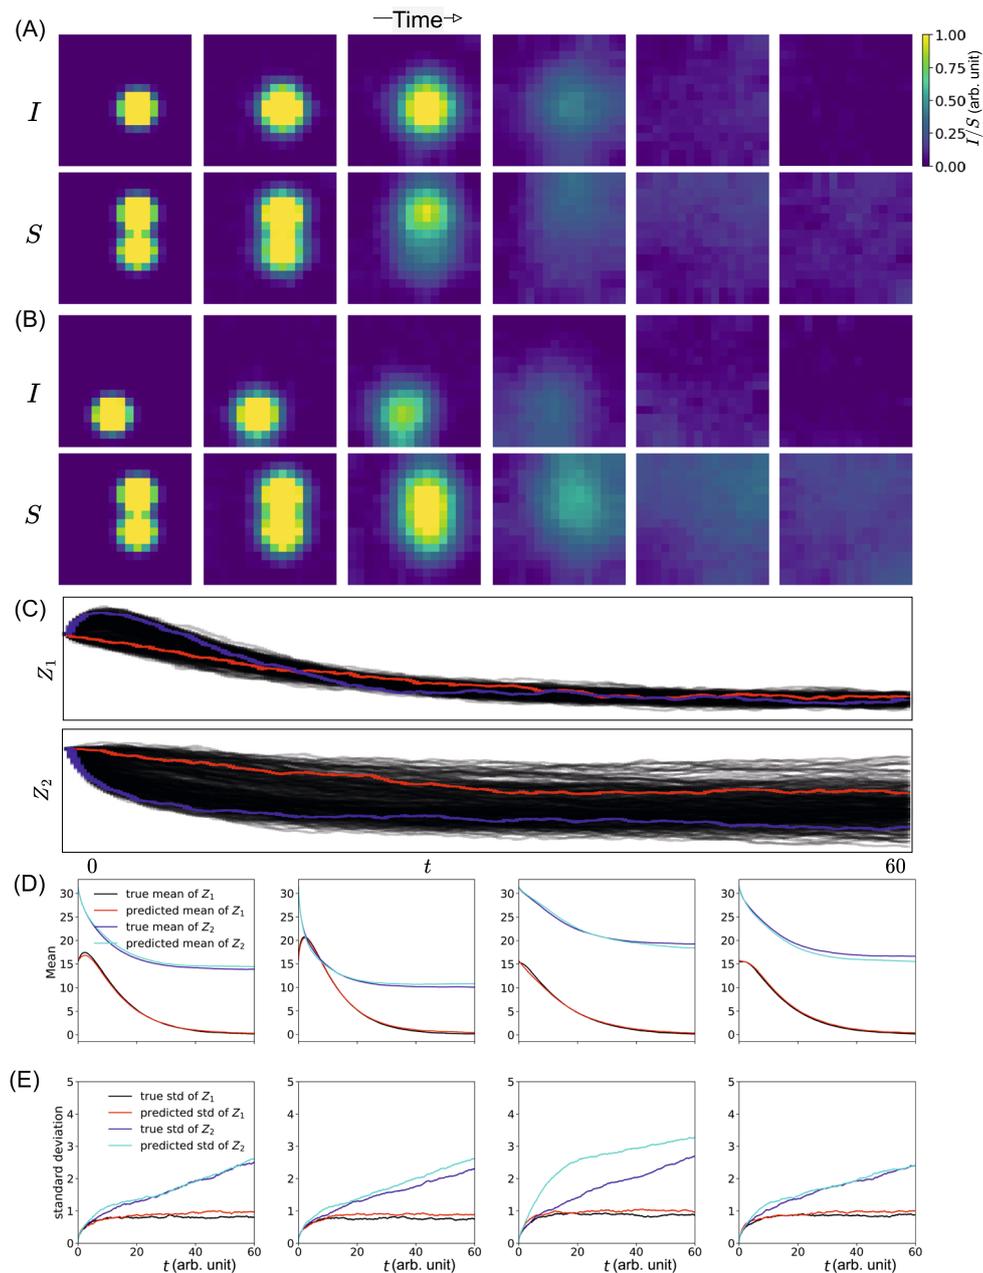
Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated

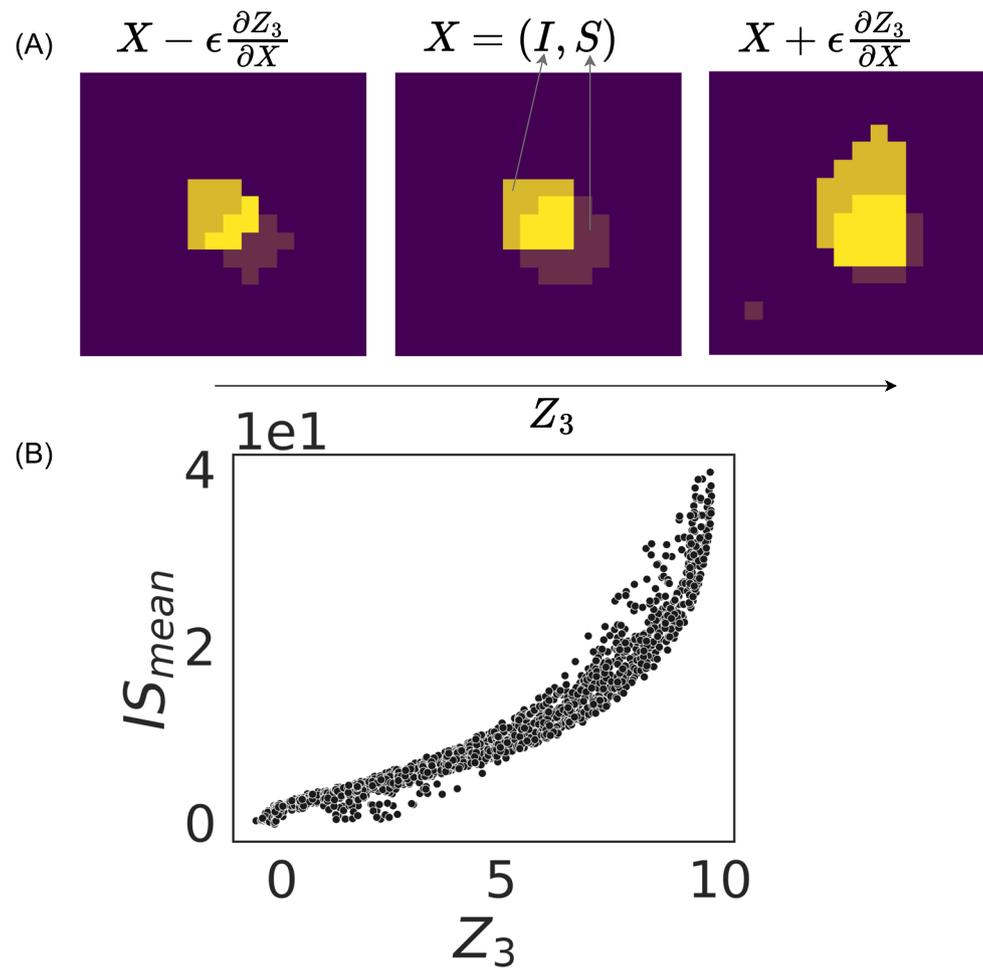
otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023



Extended Data Fig. 1 | Data visualization and prediction results. A) and (B): Two trajectories depicting the spatial evolution of I (infective) and S (susceptible) with distinct initial conditions are plotted. They have identical spatial averages initially but differing subsequent evolution. In particular, in (A) the disease spreads (Z_1 , the spatial average of I increases initially) but in (B) the disease dies out monotonically. (C) Scatter of Z_1 and Z_2 (spatial average of S) trajectories,

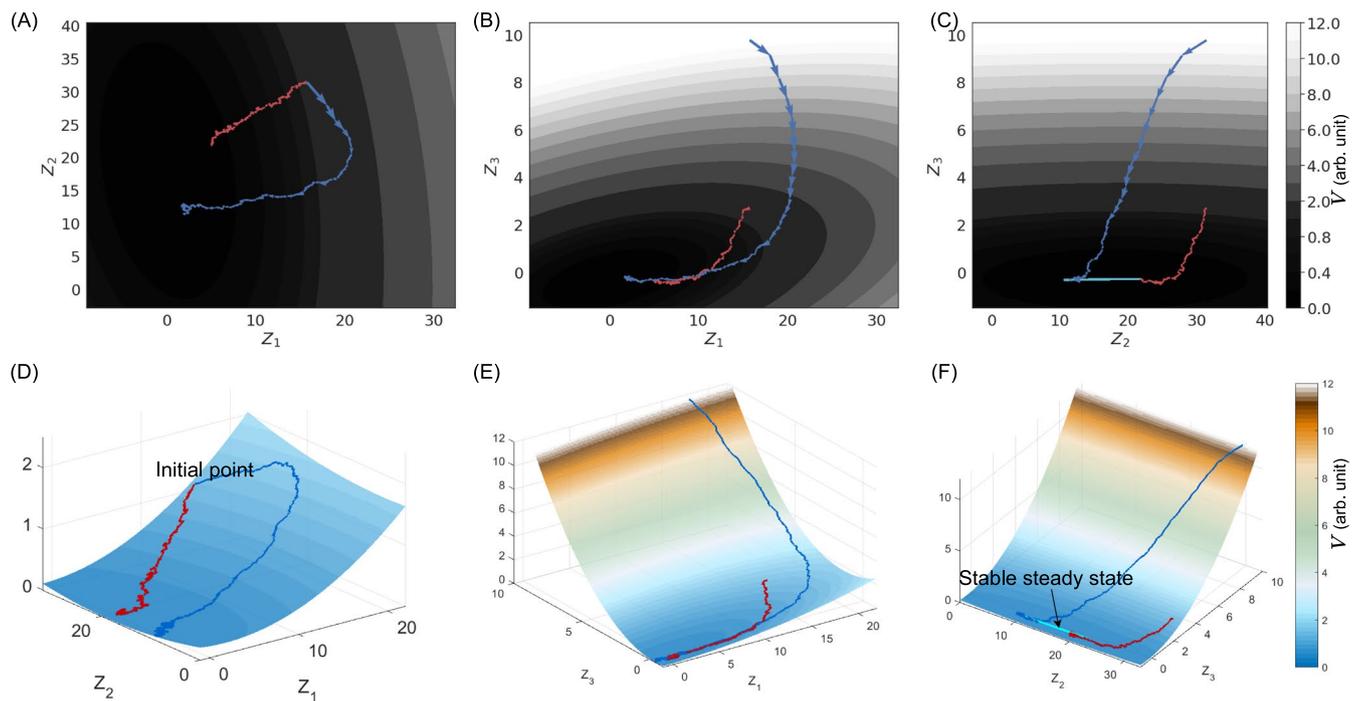
showing a high degree of variability despite identical initial values. Note that there is variability in both the presence of disease spread (Z_1 increasing initially) and the terminal value of Z_2 , corresponding to the remaining uninfected population after the epidemic. The blue (resp. red) trajectory corresponds to (A) (resp. (B)). (D, E) True vs predicted statistics using S-OnsagerNet with four different test initial conditions, showing good agreement.



Extended Data Fig. 2 | Physical interpretation of learned closure coordinate.

(A) Perturbations $X \pm \epsilon \frac{\partial Z_3}{\partial X}$ from a given configuration $X = (I, S)$ with $\epsilon = 2$ are shown. The are overlapping clusters, in the form of Gaussians centered at $(0, 0)$ for S and $(-1, 1)$ for I . To illustrate the overlap clearly, we binarize the values as follows: if the value of I and S is greater than 1, we truncate it to 1; otherwise, if the

value is less than 1, we truncate it to 0. Observe from (A) that increasing Z_3 corresponds to increasing spatial overlap of the clusters, and vice versa. We confirm this in (B), where we plot IS_{mean} (spatial overlap) vs Z_3 , showing a positive correlation.



Extended Data Fig. 3 | Potential landscape of the SIR model. The learned V is projected onto $Z_1 - Z_2$ (A,D), $Z_1 - Z_3$ (B, E) and $Z_2 - Z_3$ (C, F) planes. Projection is computed via minimization (for example $V(Z_1, Z_2) = \min_{Z_3} V(Z_1, Z_2, Z_3)$). Example of disease spread (blue) and disease dying out (red) trajectories with the same initial Z_1 and Z_2 from the training dataset are shown. We observe from (B) that Z_3

determines the onset of disease spread and differentiates the two trajectories. Moreover, (C) shows that Z_3 also differentiates the final outcome of the epidemics, where the final Z_2 value depends on the initial Z_3 value, and belongs to a 1D manifold of stable steady states as shown in (F).

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a Confirmed

- The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
- A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.
- A description of all covariates tested
- A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
Give P values as exact values whenever suitable.
- For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection

Data analysis

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our [policy](#)

The simulation and experimental datasets used are publicly available in the Harvard Dataverse public repository <https://doi.org/10.7910/DVN/NRIX7Y>.

Human research participants

Policy information about [studies involving human research participants and Sex and Gender in Research](#).

Reporting on sex and gender	<input type="text" value="N/A"/>
Population characteristics	<input type="text" value="N/A"/>
Recruitment	<input type="text" value="N/A"/>
Ethics oversight	<input type="text" value="N/A"/>

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

Life sciences Behavioural & social sciences Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Ecological, evolutionary & environmental sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	<input type="text" value="N/A"/>
Research sample	<input type="text" value="N/A"/>
Sampling strategy	<input type="text" value="N/A"/>
Data collection	<input type="text" value="N/A"/>
Timing and spatial scale	<input type="text" value="N/A"/>
Data exclusions	<input type="text" value="N/A"/>
Reproducibility	<input type="text" value="N/A"/>
Randomization	<input type="text" value="N/A"/>
Blinding	<input type="text" value="N/A"/>

Did the study involve field work? Yes No

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

- | n/a | Included in the study |
|-------------------------------------|--|
| <input checked="" type="checkbox"/> | <input type="checkbox"/> Antibodies |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> Eukaryotic cell lines |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> Palaeontology and archaeology |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> Animals and other organisms |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> Clinical data |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> Dual use research of concern |

Methods

- | n/a | Included in the study |
|-------------------------------------|---|
| <input checked="" type="checkbox"/> | <input type="checkbox"/> ChIP-seq |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> Flow cytometry |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> MRI-based neuroimaging |