

Characterizing Nonlinear Dynamics by Contrastive Cartography

Nicolas Romeo ^{1,2,*} Chris Chi ^{1,3} Aaron R. Dinner ^{1,3,†} and Elizabeth R. Jerison ^{1,2,‡}

¹Center for Living Systems, *University of Chicago, Chicago, Illinois 60637, USA*

²Department of Physics, *University of Chicago, Chicago, Illinois 60637, USA*

³Department of Chemistry, *University of Chicago, Chicago, Illinois 60637, USA*



(Received 22 April 2025; accepted 30 July 2025; published 11 September 2025)

The qualitative study of dynamical systems using bifurcation theory is key to understanding systems from biological clocks and neurons to physical phase transitions. Data generated from such systems can feature complex transients, an unknown number of attractors, and stochasticity. Characterization of these often-complicated behaviors remains challenging. Making an analogy to bifurcation analysis, which specifies that useful dynamical features are often invariant to coordinate transformations, we leverage contrastive learning to devise a generic tool to discover dynamical classes from stochastic trajectory data. By providing a model-free trajectory analysis tool, this method automatically recovers the dynamical phase diagram of known models and provides a “map” of dynamical behaviors for a large ensemble of dynamical systems. The method thus provides a way to characterize and compare dynamical trajectories without governing equations or prior knowledge of target behavior. We additionally show that the same strategy can be used to characterize the stochastic motion of bacteria, establishing that this approach can be used as a stand-alone analysis tool or as a component of a broader data-driven analysis framework for dynamical data.

DOI: [10.1103/jc9p-m3rn](https://doi.org/10.1103/jc9p-m3rn)

I. INTRODUCTION

Complex dynamics in many biological, physical, and engineered systems are difficult to analyze: Even relatively simple nonlinear dynamical systems can display markedly different behaviors as their parameters are changed, from oscillations and spikes to chaos [1–4]. With the emergence of powerful data-driven approaches, symbolic or generative models fit to data are capable of reproducing experimentally observed time series [5–11]. However, it remains challenging to identify and interpret the qualitative behaviors that such complex and often nonlinear models generate [12]. For example, suppose you observed a biochemical system that oscillated in some regimes and was multistable in others. How would you identify those behaviors?

The classical tool to study behavior regimes of nonlinear dynamical systems is bifurcation theory. Concerned with the study of qualitative changes in the flow of dynamical systems, it is the basis of theories across physical and biological disciplines, from the dynamical picture of phase transitions [13,14] and the analysis of biological regulatory networks [15,16] to the identification of ecological tipping points [17,18] and reduced order modeling in engineering [19].

Through the characterization of attractors in the dynamical flow, bifurcation theory provides a link between model parameters and behaviors. Defining observables (order parameters) characteristic of these attractors allows for quantification of changes in target behavior. This quantification in principle allows for determining which behavior regimes are present in experimental and observational data. However, in practice, differences between behaviors can be subtle, necessitating the development and use of specific statistical tests for each behavior of interest. For example, extensive work has been devoted to determining when oscillations are present in time-series data [20–22]. Additionally, most studies of bifurcations are limited to determining the invariant sets and linear analysis of flows around them.

However, global bifurcations involve the interaction of multiple flow structures and cannot be captured through linear analysis around invariant sets, making the definition of observables much more difficult [23,24]. Defining summary statistics that capture complex transient behaviors is also challenging [25,26]. The presence of noise further complicates analyses based on bifurcation theory, as the definition of invariant sets and flows requires the introduction of probabilistic criteria [21,27–29].

These challenges make developing statistics for complex behaviors a difficult task. These hurdles arise often in analyses of living systems, where complex, noisy dynamics are common, and there is usually no first-principles model of the dynamics. To define behavior without the need for prior knowledge, recent work has developed generic ways to characterize nonlinear dynamical systems by studying their approximations via linear models [30–32] or Markov chains [33,34]. While powerful and adaptable, these approaches have

*Contact author: nromeo@uchicago.edu

†Contact author: dinner@uchicago.edu

‡Contact author: ejerison@uchicago.edu

the drawback that they can obscure simple nonlinear effects in favor of complex linear models. For example, no linear model can be bistable.

Machine-learning methods for reducing the dimensionality of dynamical data to a few informative variables have proven effective to both simplify the dynamical space [35–37] and to characterize complex time series [12,38,39]. These methods take data generated by a particular process—either a model or an experiment—and train a neural network to produce a meaningful low-dimensional embedding. While powerful, the characterizations obtained from these methods do not transparently connect to bifurcation theory, making the study of transitions difficult.

Here our goal is to characterize the dynamical behavior of a set of stochastic trajectories without prior knowledge in a manner that recovers bifurcation theory when it is applicable. One earlier study also pursued this goal for deterministic trajectories [40]. Their approach, based on diffusion maps, yielded promising results on a small number of model systems but necessitated complex hand-construction of preprocessing and dimensional reduction steps and used methods that do not allow for out-of-sample comparisons. Given advances in machine learning, we were motivated to ask whether a different approach could enable application to noisy trajectories across a broad range of nonlinear systems.

Our strategy is to use contrastive learning, which is commonly used in modern computer vision and natural language processing [41–44] to learn a representation of trajectories characterizing the observed dynamics. By designing data augmentations based on some of the topological invariances of these systems, we recover a latent space that is consistent with bifurcation theory and can be used to classify and separate dynamical behaviors. We further analyze experimental data for a biological process for which low-dimensional dynamical bifurcations are less relevant because of intrinsic stochasticity and complex dynamics that are internal to cells, and we show that a similar self-supervised strategy to incorporate geometric invariance can still provide more informative latent spaces than generic dimensional reduction methods.

II. RESULTS

We present the construction of the latent space and evaluate its ability to discriminate between stereotypical dynamical behaviors. We then show that it reflects the topological structure of the dynamical flows that generated the input data and present additional features of the latent space that can make the method useful for general dynamical data exploration. We finally show, using an application to run-and-tumble bacterial tracks, that accounting for invariances using a similar contrastive strategy can be valuable to characterize strongly stochastic dynamics.

A. Contrastive learning framework

Dynamical systems can exhibit various behaviors. Consider a two-dimensional (2D) stochastic dynamical system $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \boldsymbol{\zeta}(t)$ that governs the evolution of a variable $\mathbf{x} = (x, y)$ under a deterministic flow $\mathbf{f}$ and standard Brownian noise $\boldsymbol{\zeta}(t)$ with variance D . A generic 2D smooth $\mathbf{f}$ can be

Taylor-expanded to order N as

$$\mathbf{f}(\mathbf{x}) = \sum_{k=1,2} \sum_{0 \leq i+j \leq N} c_{ij}^k x^i y^j \mathbf{e}_k + O(x^{N+1}, y^{N+1}) \quad (1)$$

with $\mathbf{e}_1 = (1, 0)^T$ and $\mathbf{e}_2 = (0, 1)^T$ the 2D Cartesian basis vectors. For $N = 3$, the resulting cubic-order differential equations cover many commonly encountered physical and biological models described by cubic normal forms and canonical models [1,2]. As we simulate [45] an ensemble of such systems with normally sampled coefficients c_{ij}^k up to $N = 3$, under the restriction that $\mathbf{f} \cdot \mathbf{x} \leq 0$ as $\mathbf{x} \rightarrow \infty$ to guarantee boundedness of the flow (Ref. [46], Sec. I), we observe a diversity of commonly encountered phenomena such as (multi-)stability, oscillations, or excitable spiking transients, with smooth transitions between them (Fig. 1(a); see Ref. [46], Fig. S1, for sampling statistics of different fixed point structures). While the different 2D dynamical behaviors are clear to the eye away from transition regimes, they are hard to characterize directly from trajectory data without prior knowledge of desired characteristics. Similar arguments apply to d -dimensional dynamical systems for which cubic polynomial systems can exhibit even more complex behaviors, including chaos in $d \geq 3$ [1,4].

A core insight from bifurcation theory is that the topological features of the flow $\mathbf{f}$ provide a way to classify dynamical systems, most commonly by establishing the presence, number, and stability of fixed points and attractors [2–4]. Such topological objects are invariant under any continuous change of coordinate system (or homeomorphism), and traditional normal form analysis seeks to find particular coordinates in which dynamics are reduced to well-characterized normal forms with particularly simple flows of the form given by Eq. (1). This perspective suggests that a neural network trained to recognize features invariant to coordinate changes should be able to detect what we mean by dynamical behavior.

To implement this idea, we use a self-supervised encoder neural network implemented in pytorch [47] which takes as inputs sets of regularly sampled trajectories $\{\mathbf{x}(t)\}$ and outputs a vector $h[\{\mathbf{x}(t)\}] \in \mathbb{R}^n$ characterizing the dynamical system. The central object of this article is the learned embedding map

$$h : \{\mathbf{x}(t)\} \mapsto h[\{\mathbf{x}(t)\}], \quad (2)$$

which we dub the *cartographer*. Our goal is to construct and train the cartographer to be interpretable in light of bifurcation theory by training it to recognize coordinate-invariant features of the input data. The construction of a neural network whose input is invariant to specific transformations, known as *augmentations*, is the central idea behind contrastive learning [42]: In the present study, the trained cartographer h must separate samples that are related by an augmentation transformation from ones that are not (Ref. [46], Sec. II). Such augmented samples should be clustered together in the embedding space, thus imposing a soft invariance constraint. However, arbitrary continuous coordinate transformations by construction erase all geometric data that could provide a basis to distinguish between otherwise topologically equivalent systems. We thus restrict ourselves to invertible linear coordinate changes along with permutations of the input trajectories (Ref. [46], Sec. IIB, with augmentations




FIG. 1. Contrastive learning to identify dynamical behavior. (a) A sample of trajectories generated from an ensemble of 2D dynamical systems with polynomial nonlinearities. The systems display a variety of dynamical behaviors. Color indicates time; white symbols indicate unstable fixed points; black symbols indicate stable fixed points. (b) Proposed pipeline for identifying dynamical behaviors, here applied to trajectories generated by the FitzHugh-Nagumo model with various parameter choices. During training on the dataset of (a) (dashed lines), linear transformations are applied to input trajectories to generate an augmented dataset which allows the optimization of a contrastive loss. Collections of trajectories $\{\mathbf{x}_i(t_j)\}$ are embedded into a smaller-dimensional space h , in which different dynamical behaviors emerge as clusters as seen through projection onto a 2D MDS basis.

generated according to Ref. [48]; nonlinear transformations do not impact performance within tested conditions, see Ref. [46], Sec. II C and Fig. S4).

Having identified a class of suitable augmentation transformations, we optimize the InfoNCE loss function with the Adam optimizer [49], a common choice for contrastive learning with strong ties to information theory [35,41,50–52]. Details of the neural network architecture, the loss function, and learning curves are provided in Ref. [46], Sec. II A, Figs. S2 and S3. Our last task to construct the map h is to define an appropriate training dataset that spans a diverse set of dynamics; for this, we integrate 2×10^4 systems sampled using the same protocol as for Fig. 1(a) (Ref. [46], Sec. I). With these training data, we find that our neural network can discriminate between different parameter regimes of stochastic FitzHugh-Nagumo systems with noise amplitude $D = 0.1$ (Fig. 1(b), Ref. [46], Sec. III A [53,54]). In the supplementary material [46], we investigate the effect of varying initial conditions (Ref. [46], Sec. VI) and compare our method to autocorrelation-based and Markov operator approaches [55,56], which, while simpler and training-free, are less accurate and do not provide an invariant characterization (Ref. [46], Sec. VII, Figs. S13 and S14). Compared to characterizing dynamics by finding sequences of linear propagators

[30], our approach is technically simpler and does not require finding a link between piecewise linear models and nonlinear behavior.

B. Cartographer classifies behaviors

What kind of understanding does the cartographer provide about dynamical systems? To answer this question, we now apply the cartographer trained on the data generated by cubic polynomial systems [i.e., the data described in the previous section and illustrated in Fig. 1(a)] to a series of examples. We emphasize that we do not further train the cartographer for any example.

As a first example, we consider the problem of identifying distinct dynamical behaviors without prior assumptions, for a system with a known analytical phase diagram. We thus consider a minimal model of nonreciprocal dynamics arising in active matter and ecological systems, dubbed the *SNIC-Hopf model*, given by [58]

$$\frac{d}{dt} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} k & 1 + \alpha \\ 1 - \alpha & k \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} - (x^2 + y^2) \begin{pmatrix} x \\ y \end{pmatrix}. \quad (3)$$

In the absence of noise $D = 0$, this system can display three different behaviors separated by well-characterized




FIG. 2. Dynamical phase identification in example 2D systems. (a) The SNIC-Hopf model system displays multiple possible behaviors as the parameters k and α are changed. Analytical results show that the different phases are separated by pitchfork (PF), Hopf, and saddle-node on an invariant circle (SNIC) bifurcations. (b) A 3-means clustering in the contrastive embedding space recovers the main distinct dynamics. Inserts show typical trajectories in the (x, y) plane. (c) Inferred phase diagram using the identified phase labels from (b). Mean-field results are overlaid in white. (d) A six-dimensional cell cycle model [57]. (e) Varying the degradation rate of the cdc2-cylin-P complex (denoted M) leads to distinct dynamical behaviors in the presence of noise. Concentrations are expressed in units of total cdc2 concentration. (f) Agglomerative clustering in latent space recovers the boundaries of the distinct behaviors in accordance with previous analytical results (white lines, see Refs. [46] and [57]). (g) As intake velocity U is increased, the flow past a channel displays a von Kármán vortex street and becomes oscillatory in time. (h) A noisy Galerkin projection of the flow $\mathbf{u}$ onto $\mathbf{w}_1$, the first dynamic mode at Reynolds number $\text{Re} = UL/\nu = 100$, shows the emergence of a limit cycle. (i) Contrastive embedding of reduced time series shows the emergence of two clusters separated at $\text{Re} \approx 50$, providing a characterization of the dynamics.

bifurcations (Fig. 2(a), see Ref. [46] Sec. IIIB). We numerically integrate the corresponding stochastic differential equations (SDEs) for varying k and α and fixed nonzero noise $D = 5 \times 10^{-4}$ and embed the trajectories using our contrastive map [Fig. 2(b)]. The embedding visually reveals three dis-

tinct regions. Using K -means clustering on the embedded data $\{h\}$ with $K = 3$, we recover the dynamical phases and their expected boundaries, albeit slightly shifted [Fig. 2(c)]. These shifts are expected: Oscillations are observed even in the monostable regime in the vicinity of Hopf bifurcations

in noisy systems, a phenomenon known as noise-induced oscillations [59], while the presence of noise tends to delay the visibility of bistability, as seen by the misidentification of bistable systems as monostable near the transition [60]. An additional example featuring more complex dynamical flow features is presented in Ref. [46] Sec. III C, Fig. S6: We find there that global bifurcations, which are discontinuous in flow space, are often easier to accurately locate than local bifurcations.

As a more complex example, we consider a six-dimensional dynamical model of the cell cycle presented in Ref. [57]; we focus on the dynamics of the cyclin and cdc2 proteins as they get phosphorylated and bind each other (Fig. 2(d); see Ref. [46] Sec. III D for model details). Under variation of the degradation rate k_6 of the phosphorylated complex M, the model can display distinct steady states separated by oscillatory and excitable regimes [Fig. 2(e)]. Given a two-dimensional projection of the integrated trajectories, the cartographer is able to identify these distinct regimes [Fig. 2(f)] even though the underlying dynamics are not in the training data (Fig. 1). For applications to data for which the number of distinct behaviors is not visually clear, clustering techniques can be used [61]. Note that we always cluster in the high-dimensional latent space and use projections only for visualization (see Ref. [46] Sec. V D for further discussion). Furthermore, it is important to acknowledge that there is no direct way to assess the quality of a clustering, and in practice validation of clusters relies on *a posteriori* evaluation of their consistency with independent knowledge.

The cartographer can also be applied to data generated by a higher-dimensional process. Consider incompressible hydrodynamic flow in a channel past a circular obstacle of diameter L (Fig. 2(g); implementation details in Ref. [46], Sec. III E, following Ref. [62] and using FeniCSx [63–66]). For a fluid of kinematic viscosity ν , as the intake flow velocity U increases, the flow undergoes a laminar-to-oscillatory transition when the Reynolds number $\text{Re} = UL/\nu$ crosses the critical value $\text{Re}_c \approx 50$. Although this partial differential equation (PDE) problem is high-dimensional, projecting the flow on the standard basis from dynamic mode decomposition [67] allows us to obtain a 2D time series which captures the two dynamical phases, and the cartographer is able to identify the critical Reynolds number [Figs. 2(h)–2(i); Ref. [46], Fig. S7]. We thus see that h is able to identify dynamical behaviors on data originally generated from a high-dimensional nonstationary system, suggesting potential applicability for the characterization of other PDE solutions and imaging data after suitable dimensionality reduction.

C. Cartographer recovers flow topology

The cartographer is thus able to discriminate between common stereotypical dynamical behaviors. However, given our learning strategy, we can expect more powerful structures to emerge in the latent space encoded by h . In this section, we show that the latent space trained only on trajectory data is consistent with bifurcation theory.

To this end, we examine the embedding of a sample of data generated from the random dynamical ensemble (but not used to train the network). Labeling trajectories by the

numerically determined fixed-point structure of the generating equations [68], we find that samples in latent space are clustered according to the number and stability of their fixed points (Fig. 3(a); alternative visualizations [69] and effects of different input datasets are presented in Ref. [46], Sec. V D, Figs. S11 and S12). Here we restrict ourselves to systems with up to two fixed points of a given type. Note that our cubic systems can have up to nine fixed points; we present sampling statistics and results for more complex topologies in Ref. [46] (Sec. V B, Figs. S1 and S9). Note that a complete topological characterization would also include the number and type of limit cycles, which is numerically challenging [70].

To explore the boundaries between dynamical regimes drawn by h , we follow the embedding $h(s)$ of trajectories from the SNIC-Hopf model as its parameters vary along the parametrized closed curve $(k(s), \alpha(s))$ with $s \in [0, 2\pi]$ (Fig. S5). By measuring the variation in embedding vector between successive parameter values through the differential similarity measure $h(s_i) \cdot h(s_{i+1})$, we observe stronger variations in the vicinity of bifurcations, with transition points shifted as before from noise-induced effects [Fig. 3(b)]. Moreover, in the 2D projection shown, we indeed observe piecewise continuous paths [Fig. 3(c)].

The latent space thus allows the definition of regions associated with different dynamical phases, in the sense of fixed-point topologies (Ref. [46], Sec. V, Fig. S10). In effect, our contrastive learning approach allows the construction of an approximate phase diagram for generic 2D dynamical systems, allowing for the identification of bifurcations by monitoring the variations in embedding.

That said, we note that there is substantial overlap between certain classes in Fig. 3(a), in apparent disagreement with the idea that flow topology determines behavior. As we now show, this is due to stochasticity, which breaks the clear-cut classification of deterministic dynamical systems by the nature of their asymptotic invariants. We also note that even in deterministic systems, the presence of non-normality [72] or “ghosts” [1,25,26,73] can lead to distinct transient behaviors not directly captured by fixed-point topology.

The effect of noise is best understood by an example. Consider the deterministic dynamics $\dot{x} = rx - x^3$, $\dot{y} = -y$. This system exhibits a single fixed point at the origin for $r < 0$, while it has two stable fixed points at $(x, y) = (\pm\sqrt{r}, 0)$ for $r > 0$. As we increase the noise, the difference between these two phases is progressively blurred [Fig. 3(d)], and this is reflected in the latent space: For small noise, as we vary r , the embedding presents two distinct clusters corresponding to mono- and bistable regimes, while, for large noise, there are no longer distinct clusters but the first principal component still relates to r [Fig. 3(e), Fig. S8]. The effect of variable time integration length is discussed in Ref. [46], Sec. V.

Coming back to the systems generated from the training ensemble, for a given trained cartographer, the accuracy of a classifier trained and tested on data integrated with noise amplitude D decreases as the noise increases [Fig. 3(f); confusion matrix shown in Ref. [46], Fig. S10(b)].

Finally, we note that many experimental data of interest are projections of higher-dimensional dynamics. Because there is a wider range of possible behaviors in higher dimensions, such data will typically be out of the distribution used for




FIG. 3. The contrastive map recovers bifurcation theory at low noise. (a) UMAP [71] embedding of ~ 2300 samples from the random ensemble, showing clustering by number and type of fixed points ($n_{\text{stable}}, n_{\text{unstable}}$). (b) The SNIC-Hopf system is taken along the path shown in the inset, and we examine the embedding $h(s)$ across different phases: bistable (B), oscillatory (O), and monostable (M). The embedding h shows discontinuities in similarity at bifurcations. (c) Comparing the embedding of the data in (b) with boundaries between topological regions shows that the bifurcations match expected region boundaries. (d) A simple potential system exhibits a pitchfork bifurcation as the control parameter r is changed, with examples at $r = -2.5$ (left) and $r = 2.6$ (right). (e) For small noise (bottom, $D = 0.006$), the two distinct phases are easily identified in the latent space, while, for large noise (top, $D = 0.39$), the phases are no longer separated in the latent space but the first principal component is still related to r . (f) The accuracy of a linear classifier to determine the topological structure through h decreases as noise increases, in accordance with results in (d) and (e).

training. To see how the cartographer performs in such situations, in Ref. [46], Sec. XI we sample data from a selected set of 3D chaotic systems and project them to 2D for processing by the cartographer trained on 2D cubic polynomial systems (Refs. [46], Fig. S17(a); [74]). The cartographer is still able to distinguish different attractors [Fig. S17(b)] and embeds such chaotic systems “away” from in-distribution systems [Fig. S17(b)]. Despite this indication that the systems lie out of the training distribution, the linear classifier trained on 2D polynomial systems is still able to make meaningful identifications; for example, classifying the two-lobed Lorenz systems as bistable and recognizing the Aizawa toroidal attractor as an oscillator [Fig. S17(c)].

D. Cartographer detects nonlinear transients

We show above that the cartographer segregates systems by topology. However, by construction, h is not fully diffeomorphism-invariant, and it is thus expected to be sensitive to geometric information as well. This suggests that the

cartographer should distinguish more subtle forms of global structure beyond topologically invariant features.

While there exist statistical tests and canonical observables to identify common behaviors [75], understanding, implementing, and running all possible tests without prior knowledge and assumptions demands considerable researcher and computer time. Our method provides general capabilities that enable quick exploration of dynamical data. We illustrate such a potential use case for the cartographer.

The cartographer h provides the ability to compare inputs for which there are no obvious discriminating observables. Consider the problem of distinguishing between a linear system and a nonlinear excitable system with the same linear dynamics but which has not spiked. The presence of nonlinearities leads to subtle differences in the trajectory statistics, which reflect the existence of a threshold in phase space: Beyond this threshold, trajectories escape the stable domain before coming back from a different direction [Fig. 4(a)]. We find that our system can distinguish between the dynamics of the two systems [Fig. 4(c)]. The map thus provides a way to turn nonlinear analytical problems into standard statistical




FIG. 4. The cartographer identifies local nonlinearities. (a) A noisy excitable system irregularly exhibits spiking trajectories. (b) Trajectories of an excitable system that has not spiked (has not crossed the threshold indicated by the dashed line) and its linear approximation (without the terms in red) which cannot spike are visually similar. (c) Embedding trajectories shows that nonlinear systems are discernible from linear systems in the latent space.

discrimination tests in latent space, without having to design specific statistical tests, which, for this example, would be based on higher-order moments of the trajectories (Ref. [46], Sec. VIII).

Another application discussed in Ref. [46], Sec. IX is the detection of nonpotential dynamics, by comparing data generated by systems deriving from a potential $\dot{\mathbf{x}} = -\nabla V(\mathbf{x})$ against observed trajectories (Fig. S15). Additionally, inspired by the notion of canonical systems used in theoretical neuroscience [24], we can also use geometric structure in the trajectory data to introduce a finer categorization of latent space beyond topologically invariant features (Ref. [46], Sec. X, Fig. S16).

E. Contrastive learning assists experimental analysis

The contrastive learning strategy allows us to identify meaningful features summarizing complex dynamics. However, many dynamics may not be fully reducible to a low-dimensional manifold, and experimental measurements of them are likely to be only partial characterizations of the system state. Can we still use our self-supervised learning strategy to characterize such data?

As an example, we consider a publicly available dataset of run-and-tumble motion of *Escherichia coli* bacteria [76] [Fig. 5(a)]. These dynamics are commonly modeled by a stochastic process with a complex high-dimensional control system, and considerable effort has been devoted to designing statistics to characterize the trajectories. We ask whether the contrastive approach can identify meaningful behavioral variability, without requiring prior knowledge. We expect that important features of the velocity dynamics are invariant under orthogonal transformations of the coordinate system (x, y) and that a contrastive approach may be useful for identifying these invariant features.

The data include measurements tracking wild type *E. coli* and mutants which vary with respect to their probability to tumble (tumble bias), which affects chemotaxis [76]. Here we train a neural network on the experimental velocity trajec-

tries $(v_x(t), v_y(t))$ of the wild type, with the same architecture (Ref. [46], Sec. II A) and InfoNCE loss as the previous cartographer, using data augmentations in the orthogonal group $O(2)$ (Ref. [46], Sec. XII). We then use the trained cartographer to embed the velocity trajectories of the mutants. Evaluating the first principal component of the embedding, we find that it explains a large part of the variance, in contrast to the principal components obtained from a direct analysis of the velocity data [Fig. 5(b)]. This suggests that enforcing rotational invariance improves the ability to identify the key low-dimensional feature of these dynamics. This is also the case when using nonlinear dimensional-reduction methods (Ref. [46], Sec. XII, Fig. S18).

What is this feature? Examining correlations of experimental measurements with the first principal component, we recover the tumble bias as the most salient feature of the observed dynamics [Fig. 5(c)]. Consistent with the analysis in Ref. [76], the first principal component also correlates strongly with the ratio of the logarithms of the concentrations of the proteins [CheB] and [CheR], demonstrating a direct link between behavioral features and biochemical parameters. Determining the tumble bias from experimental bacterial tracks required the authors of Ref. [76] to construct a probabilistic model to identify tumbling and running periods. The fact that the symmetry-informed contrastive approach introduced here yields the known relationships without the need for such a system-specific model shows the usefulness of the approach for interpreting experimental data.

III. DISCUSSION

Here we show how the minimization of a contrastive learning objective, paired with a data-augmentation strategy enforcing coordinate invariance, allows learning an interpretable representation of multidimensional dynamical trajectories even with a relatively simple neural network. Our work provides a physics-informed approach for using machine learning with dynamical data distinct from the common generative model, in which machine learning-based




FIG. 5. Contrastive cartography for bacterial behavior characterization. (a) Single bacterial tracks alternate between running phases (higher velocity, more persistent orientation) and tumbling phases (slower velocity, rapid orientation change). (b) As a representation tool, contrastive cartography produces embeddings with better linear separability than the original data, with more variance captured along the first few principal components. (c) The first principal component correlates with tumble bias, as well as the ratio of the proteins $\log[\text{CheB}]/\log[\text{CheR}]$. See Ref. [46], Sec. XII, Fig. S18 for t-SNE representation and additional correlates.

regressions are used to find optimal dynamical update rules to extrapolate dynamics from data [5]. Indeed, having a generative model, whether symbolic or data driven, does not directly provide a qualitative characterization because analytical methods can rapidly become limited in the face of nonlinearities and growing numbers of parameters.

We focused on 2D dynamical systems and showed that the cartographer trained on 2D systems can still meaningfully separate projections of 3D systems. However, a full analysis of higher-dimensional systems would require consideration of the much wider range of behaviors possible in these systems, including chaotic dynamics as well as additional nonchaotic attractors. To explore the potential utility of the cartographer for higher dimensions, we trained a cartographer on 3D cubic polynomial systems (Ref. [46], Sec. XI B) and embedded the set of example chaotic systems. The 3D cartographer is able to better distinguish the different chaotic attractors, indicating that extension to 3D may be feasible and useful [Fig. S17(d)]. However, we leave further characterization of the very broad set of dynamics possible in 3D to future work. Beyond 3D, the training ensemble of cubic-order models becomes harder to sample and more refined forms of sampling (e.g., Ref. [77]) might become necessary to provide a sufficiently rich training ensemble.

An additional possible application of our method is to leverage the “phase diagram” encoded by the latent space to identify tipping points in complex systems, a problem which has attracted both classical and data-driven work [17,78–80]. As dynamics in the vicinity of a bifurcation (or generalizations thereof in the presence of noise) are expected to be included in the polynomial class of systems used to train the network, one could use low-dimensional representations of the complex system dynamics to detect sudden changes in embedding without assumptions on the specific kind of bifurcations present in the system.

One could also hope to leverage cartography as a heuristic to help infer interpretable candidate models from data, similarly to sparsity [6,81]. A heuristic to pick a simple model amenable to theory work—e.g., a normal form that is close in

some sense to the observed dynamics—could help provide the elements of a Landau-like phenomenological description. By providing otherwise inaccessible observables characteristic of canonical nonlinear models, such a construction could be relevant even in nonequilibrium and weakly symmetric systems where global phase-space features are often relevant, such as animal behavior [39], biological development, or immunity [82].

Recent methods for characterizing dynamics for such complex systems involve fitting effective linear time-evolution operators that can then be used as fingerprints of the underlying dynamics [30–32,34]. Our trajectory-based work is complementary to those operator-theoretic approaches, trading the inference of linear, often higher-dimensional operators for an interpretable latent space that captures nonlinear effects in low-dimensional systems. Combining the strengths of both approaches could help provide widely applicable characterizations that capture nonlinear phenomena.

A limitation of our method comes from the difficulty of interpreting latent spaces without external knowledge. For example, the knowledge of the topological structure of polynomial systems allows us to use a linear classifier to formalize the connection between latent space coordinates and topological structure in Fig. 3. This classifier can serve as a set of feature labels for the latent space; even in cases where the data are clearly outside of the training distribution, as in the chaotic systems, the cartographer identifies the “closest” trained 2D topology. This training and classification strategy could be extended, for example, to different types of stochastic dynamics. When the cartographer is trained directly on data, as is done here in Fig. 5, correlating external information with the found latent space is essential for interpretability. As for the polynomial systems, supervised machine learning can then help formalize mappings between relevant parameters or properties and latent space position.

As shown in Fig. 3, another fundamental limitation is that sufficient amounts of noise obscure features of the underlying fixed-point topology, making classification less reliable. An additional limitation arises from the challenges

inherent to characterizing multiscale phenomena, whether in time or in amplitude. In particular, detecting oscillations requires sampling faster than all relevant frequencies. Observing and successfully identifying rare events could also require prohibitively long data inputs which our current architecture might have trouble supporting.

While it is now possible to embed practically any data into a vector space [83], making such embeddings interpretable and useful for further theory work is an outstanding challenge. Using symmetries and invariances that are common in physics, contrastive learning can identify forms of order that are otherwise difficult to characterize [50]. With time-resolved high-dimensional data—whether from biological imaging, “-omics” datasets, or large-scale sensor batteries—becoming increasingly common, we hope that this strategy makes identifying distinct behaviors a simpler task on the path to model and control complex systems. Code and trajectory datasets are available [84,85].

ACKNOWLEDGMENTS

We gratefully acknowledge helpful discussions and code from Arvind Murugan, Matthew S. Schmitt, Peter Y. Lu, Hermann Riecke, and Daniel S. Seara. N.R. acknowledges support from the University of Chicago Biological Physics and Center for Living Systems Fellowships. E.R.J. acknowledges support from the Burroughs Wellcome Fund through a Career Award at the Scientific Interface and the Henry Luce Foundation through a Clare Boothe Luce Assistant Professorship. This research was supported by the Physics Frontier Center for Living Systems funded by the National Science Foundation (PHY-2317138). Computing resources were provided by the University of Chicago Research Computing Center.

DATA AVAILABILITY

The data and code that support the findings of this article are openly available [84,85].

- [1] S. H. Strogatz, *Nonlinear Dynamics and Chaos*, 3rd ed. (Chapman & Hall/CRC, Philadelphia, PA, 2024)
- [2] E. M. Izhikevich, *Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting* (The MIT Press, Cambridge, MA, 2006).
- [3] Y. A. Kuznetsov, *Elements of Applied Bifurcation Theory* (Springer International Publishing, Cham, 2023).
- [4] J. D. Meiss, *Differential Dynamical Systems* (Society for Industrial and Applied Mathematics, Philadelphia, PA, 2007).
- [5] W. Gilpin, Generative learning for nonlinear dynamics, *Nat. Rev. Phys.* **6**, 194 (2024).
- [6] S. L. Brunton, J. L. Proctor, and J. N. Kutz, Discovering governing equations from data by sparse identification of nonlinear dynamical systems, *Proc. Natl. Acad. Sci. USA* **113**, 3932 (2016).
- [7] R. Supekar, B. Song, A. Hastewell, G. P. T. Choi, A. Mietke, and J. Dunkel, Learning hydrodynamic equations for active matter from particle simulations and experiments, *Proc. Natl. Acad. Sci. USA* **120**, e2206994120 (2023).
- [8] M. Schmidt and H. Lipson, Distilling free-form natural laws from experimental data, *Science* **324**, 81 (2009).
- [9] P. Y. Lu, S. Kim, and M. Soljačić, Extracting interpretable physical parameters from spatiotemporal systems using unsupervised learning, *Phys. Rev. X* **10**, 031056 (2020).
- [10] A. Frishman and P. Ronceray, Learning force fields from stochastic trajectories, *Phys. Rev. X* **10**, 021009 (2020).
- [11] T. Kadeethum, D. O’Malley, J. N. Fuhg, Y. Choi, J. Lee, H. S. Viswanathan, and N. Bouklas, A framework for data-driven solution and parameter estimation of PDEs using conditional generative adversarial networks, *Nat. Comput. Sci.* **1**, 819 (2021).
- [12] M. J. Falk, F. D. Roach, W. Gilpin, and A. Murugan, Curiosity-driven search for novel nonequilibrium behaviors, *Phys. Rev. Res.* **6**, 033052 (2024).
- [13] M. Fruchart, R. Hanai, P. B. Littlewood, and V. Vitelli, Non-reciprocal phase transitions, *Nature (London)* **592**, 363 (2021).
- [14] U. C. Täuber, *Critical Dynamics: A Field Theory Approach to Equilibrium and Non-equilibrium Scaling Behavior* (Cambridge University Press, Cambridge, UK, 2014).
- [15] M. B. Elowitz and S. Leibler, A synthetic oscillatory network of transcriptional regulators, *Nature (London)* **403**, 335 (2000).
- [16] U. Alon, *An Introduction to Systems Biology: Design Principles of Biological Circuits*, 2nd ed. (CRC Press, Boca Raton, FL, 2020).
- [17] L. Dai, D. Vorselen, K. S. Korolev, and J. Gore, Generic indicators for loss of resilience before a tipping point leading to population collapse, *Science* **336**, 1175 (2012).
- [18] V. Dakos, B. Matthews, A. P. Hendry, J. Levine, N. Loeuille, J. Norberg, P. Nosil, M. Scheffer, and L. De Meester, Ecosystem tipping points in an evolving world, *Nat. Ecol. Evol.* **3**, 355 (2019).
- [19] S. Brunton and J. Kutz, *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control* (Cambridge University Press, Cambridge, UK, 2022).
- [20] P. Welch, The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms, *IEEE Trans. Audio Electroacoust.* **15**, 70 (1967).
- [21] D. A. Dickey and W. A. Fuller, Distribution of the estimators for autoregressive time series with a unit root, *J. Am. Stat. Assoc.* **74**, 427 (1979).
- [22] A. L. Hutchison, M. Maienschein-Cline, A. H. Chiang, S. M. A. Tabei, H. Gudjonson, N. Bahroos, R. Allada, and A. R. Dinner, Improved statistical methods enable greater sensitivity in rhythm detection for genome-wide data, *PLoS Comput. Biol.* **11**, e1004094 (2015).
- [23] F. C. Hoppensteadt and E. M. Izhikevich, in *Weakly Connected Neural Networks*, edited by J. E. Marsden, L. Sirovich, and F. John, Applied Mathematical Sciences Vol. 126 (Springer, New York, 1997).
- [24] E. M. Izhikevich, Neural excitability, spiking and bursting, *Int. J. Bifurcat. Chaos* **10**, 1171 (2000).
- [25] A. Hastings, K. C. Abbott, K. Cuddington, T. Francis, G. Gellner, Y.-C. Lai, A. Morozov, S. Petrovskii, K. Scranton, and M. L. Zeeman, Transient phenomena in ecology, *Science* **361**, eaat6412 (2018).
- [26] D. Koch, A. Nandan, G. Ramesan, I. Tyukin, A. Gorban, and A. Koseska, Ghost channels and ghost cycles guiding long

- transients in dynamical systems, *Phys. Rev. Lett.* **133**, 047202 (2024).
- [27] C. Kwon, P. Ao, and D. J. Thouless, Structure of stochastic dynamics near fixed points, *Proc. Natl. Acad. Sci. USA* **102**, 13029 (2005).
 - [28] L. Billings and I. B. Schwartz, Identifying almost invariant sets in stochastic dynamical systems, *Chaos* **18**, 023122 (2008).
 - [29] J. E. Chung, J. F. Magland, A. H. Barnett, V. M. Tolosa, A. C. Tooker, K. Y. Lee, K. G. Shah, S. H. Felix, L. M. Frank, and L. F. Greengard, A fully automated approach to spike sorting, *Neuron* **95**, 1381 (2017).
 - [30] A. C. Costa, T. Ahamed, and G. J. Stephens, Adaptive, locally linear models of complex dynamics, *Proc. Natl. Acad. Sci. USA* **116**, 1501 (2019).
 - [31] N. Romeo, A. Hastewell, A. Mietke, and J. Dunkel, Learning developmental mode dynamics from single-cell trajectories, *eLife* **10**, e68679 (2021).
 - [32] A. E. Cohen, A. D. Hastewell, S. Pradhan, S. W. Flavell, and J. Dunkel, Schrödinger dynamics and Berry phase of undulatory locomotion, *Phys. Rev. Lett.* **130**, 258402 (2023).
 - [33] A. B. Wiltshko, M. J. Johnson, G. Iurilli, R. E. Peterson, J. M. Katon, S. L. Pashkovski, V. E. Abaira, R. P. Adams, and S. R. Datta, Mapping sub-second structure in mouse behavior, *Neuron* **88**, 1121 (2015).
 - [34] G. Sridhar, M. Vergassola, J. C. Marques, M. B. Orger, A. C. Costa, and C. Wyart, Uncovering multiscale structure in the variability of larval zebrafish navigation, *Proc. Natl. Acad. Sci. USA* **121**, e2410254121 (2024).
 - [35] M. S. Schmitt, M. Koch-Janusz, M. Fruchart, D. S. Seara, M. Rust, and V. Vitelli, Information theory for data-driven model reduction in physics and biology, *arXiv:2312.06608*.
 - [36] P. Conti, J. Kneifl, A. Manzoni, A. Frangi, J. Fehr, S. L. Brunton, and J. N. Kutz, VENI, VINDy, VICI: A variational reduced-order modeling framework with uncertainty quantification, *arXiv:2405.20905*.
 - [37] K. Zeng, C. E. P. De Jesús, A. J. Fox, and M. D. Graham, Autoencoders for discovering manifold dimension and coordinates in data from complex dynamical systems, *Mach. Learn.: Sci. Technol.* **5**, 025053 (2024).
 - [38] S. R. Achar, F. X. P. Bourassa, T. J. Rademaker, A. Lee, T. Kondo, E. Salazar-Cavazos, J. S. Davies, N. Taylor, P. François, and G. Altan-Bonnet, Universal antigen encoding of T cell activation from high-dimensional cytokine dynamics, *Science* **376**, 880 (2022).
 - [39] G. J. Berman, D. M. Choi, W. Bialek, and J. W. Shaevitz, Mapping the stereotyped behaviour of freely moving fruit flies, *J. R. Soc. Interface* **11**, 20140672 (2014).
 - [40] O. Yair, R. Talmon, R. R. Coifman, and I. G. Kevrekidis, Reconstruction of normal forms by learning informed observation geometries from data, *Proc. Natl. Acad. Sci. USA* **114**, E7865 (2017).
 - [41] A. van den Oord, Y. Li, and O. Vinyals, Representation learning with contrastive predictive coding, *arXiv:1807.03748*.
 - [42] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, A simple framework for contrastive learning of visual representations, *arXiv:2002.05709*.
 - [43] Y.-S. Chuang, R. Dangovski, H. Luo, Y. Zhang, S. Chang, M. Soljačić, S.-W. Li, W. tau Yih, Y. Kim, and J. Glass, DiffCSE: Difference-based contrastive learning for sentence embeddings, *arXiv:2204.10298*.
 - [44] C. Loh, T. Christensen, R. Dangovski, S. Kim, and M. Soljačić, Surrogate- and invariance-boosted contrastive learning for data-scarce applications in science, *Nat. Commun.* **13**, 4223 (2022).
 - [45] C. Rackauckas and Q. Nie, DifferentialEquations.jl—A performant and feature-rich ecosystem for solving differential equations in Julia, *J. Open Res. Softw.* **5**, 15 (2017).
 - [46] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/jc9p-m3m> for implementation details, benchmarks and additional use cases.
 - [47] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, PyTorch: An imperative style, high-performance deep learning library, *arXiv:1912.01703*.
 - [48] F. Mezzadri, How to generate random matrices from the classical compact groups, *NOTICES of the AMS* **54**, 592 (2007).
 - [49] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, *arXiv:1412.6980*.
 - [50] D. E. Gökmen, Z. Ringel, S. D. Huber, and M. Koch-Janusz, Symmetries and phase diagrams with real-space mutual information neural estimation, *Phys. Rev. E* **104**, 064106 (2021).
 - [51] K. A. Murphy and D. S. Bassett, Machine-learning optimized measurements of chaotic dynamical systems via the information bottleneck, *Phys. Rev. Lett.* **132**, 197201 (2024).
 - [52] R. Jiang, P. Y. Lu, E. Orlova, and R. Willett, Training neural operators to preserve invariant measures of chaotic attractors, in *Advances in Neural Information Processing Systems*, edited by A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Curran Associates, New York, 2023), Vol. 36, pp. 27645–27669.
 - [53] R. FitzHugh, Impulses and physiological states in theoretical models of nerve membrane, *Biophys. J.* **1**, 445 (1961).
 - [54] J. Nagumo, S. Arimoto, and S. Yoshizawa, An active pulse transmission line simulating nerve axon, *Proc. IRE* **50**, 2061 (1962).
 - [55] F. Takens, Detecting strange attractors in turbulence, in *Dynamical Systems and Turbulence, Warwick 1980*, edited by D. Rand and L.-S. Young (Springer, Berlin, 1981), Vol. 898, pp. 366–381.
 - [56] M. Hoffmann, M. Scherer, T. Hempel, A. Mardt, B. de Silva, B. E. Husic, S. Klus, H. Wu, N. Kutz, S. L. Brunton, and F. Noé, Deeptime: A Python library for machine learning dynamical models from time series data, *Mach. Learn.: Sci. Technol.* **3**, 015009 (2022).
 - [57] J. J. Tyson, Modeling the cell division cycle: cdc2 and cyclin interactions, *Proc. Natl. Acad. Sci. USA* **88**, 7328 (1991).
 - [58] D. Martin, D. Seara, Y. Avni, M. Fruchart, and V. Vitelli, The transition to collective motion in nonreciprocal active matter: Coarse graining agent-based models into fluctuating hydrodynamics, (2024) *arXiv:2307.08251*.
 - [59] J. M. G. Vilar, H. Y. Kueh, N. Barkai, and S. Leibler, Mechanisms of noise-resistance in genetic oscillators, *Proc. Natl. Acad. Sci. USA* **99**, 5988 (2002).
 - [60] P. Gaspard, G. Nicolis, A. Provata, and S. Tasaki, Spectral signature of the pitchfork bifurcation: Liouville equation approach, *Phys. Rev. E* **51**, 74 (1995).
 - [61] K. P. Murphy, *Probabilistic Machine Learning: An Introduction* (MIT Press, Cambridge, MA, 2022).

- [62] V. John and G. Matthies, Higher-order finite element discretizations in a benchmark problem for incompressible flows, *Int. J. Numer. Methods Fluids* **37**, 885 (2001).
- [63] I. A. Baratta, J. P. Dean, J. S. Dokken, M. Habera, J. S. Hale, C. N. Richardson, M. E. Rognes, M. W. Scroggs, N. Sime, and G. N. Wells, DOLFINx: The next generation FEniCS problem solving environment (2023), doi:[10.5281/zenodo.10447666](https://doi.org/10.5281/zenodo.10447666).
- [64] M. S. Alnaes, A. Logg, K. B. Ølgaard, M. E. Rognes, and G. N. Wells, Unified form language: A domain-specific language for weak formulations of partial differential equations, *ACM Trans. Math. Softw.* **40**, 1 (2014).
- [65] M. W. Scroggs, I. A. Baratta, C. N. Richardson, and G. N. Wells, Basix: A runtime finite element basis evaluation library, *J. Open Source Softw.* **7**, 3982 (2022).
- [66] M. W. Scroggs, J. S. Dokken, C. N. Richardson, and G. N. Wells, Construction of arbitrary order finite element degree-of-freedom maps on polygonal and polyhedral cell meshes, *ACM Trans. Math. Softw.* **48**, 1 (2022).
- [67] P. J. Schmid, Dynamic mode decomposition of numerical and experimental data, *J. Fluid Mech.* **656**, 5 (2010).
- [68] P. Breiding and S. Timme, Homotopycontinuation.jl: A package for homotopy continuation in Julia, [arXiv:1711.10911](https://arxiv.org/abs/1711.10911).
- [69] L. van der Maaten and G. Hinton, Visualizing data using t-SNE, *J. Mach. Learn. Res.* **9**, 2579 (2008).
- [70] G. Datseris, K. Luiz Rossi, and A. Wagemakers, Framework for global stability analysis of dynamical systems, *Chaos* **33**, 073151 (2023).
- [71] L. McInnes, J. Healy, and J. Melville, UMAP: Uniform manifold approximation and projection for dimension reduction, [arXiv:1802.03426](https://arxiv.org/abs/1802.03426).
- [72] L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. A. Driscoll, Hydrodynamic stability without eigenvalues, *Science* **261**, 578 (1993).
- [73] R. Sciammas, Y. Li, A. Warmflash, Y. Song, A. R. Dinner, and H. Singh, An incoherent regulatory network architecture that orchestrates B cell diversification in response to antigen signaling, *Mol. Syst. Biol.* **7**, 495 (2011).
- [74] W. Gilpin, Chaos as an interpretable benchmark for forecasting and data-driven modelling, [arXiv:2110.05266](https://arxiv.org/abs/2110.05266).
- [75] M. Christ, N. Braun, J. Neuffer, and A. W. Kempa-Liehr, Time series feature extraction on basis of scalable hypothesis tests (tsfresh—a Python package), *Neurocomputing* **307**, 72 (2018).
- [76] Y. S. Dufour, S. Gillet, N. W. Frankel, D. B. Weibel, and T. Emonet, Direct correlation between motile behavior and protein abundance in single cells, *PLoS Comput. Biol.* **12**, e1005041 (2016).
- [77] C. Chi, J. Weare, and A. R. Dinner, Sampling parameters of ordinary differential equations with langevin dynamics that satisfy constraints, [arXiv:2408.15505](https://arxiv.org/abs/2408.15505).
- [78] M. Scheffer, J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, H. Held, E. H. Van Nes, M. Rietkerk, and G. Sugihara, Early-warning signals for critical transitions, *Nature (London)* **461**, 53 (2009).
- [79] T. M. Bury, R. I. Sujith, I. Pavithran, M. Scheffer, T. M. Lenton, M. Anand, and C. T. Bauch, Deep learning for early warning signals of tipping points, *Proc. Natl. Acad. Sci. USA* **118**, e2106140118 (2021).
- [80] Y. Huang, S. Bathiany, P. Ashwin, and N. Boers, Deep learning for predicting rate-induced tipping, *Nat. Mach. Intell.* **6**, 1556 (2024).
- [81] P. J. Baddoo, B. Herrmann, B. J. McKeon, J. Nathan Kutz, and S. L. Brunton, Physics-informed dynamic mode decomposition, *Proc. Roy. Soc. A* **479**, 20220576 (2023).
- [82] Y. Lebel, T. Milo, A. Bar, A. Mayo, and U. Alon, Excitable dynamics of flares and relapses in autoimmune diseases, *iScience* **26**, 108084 (2023).
- [83] P. Mehta, A twenty-first century statistical physics of life, [arXiv:2410.20506](https://arxiv.org/abs/2410.20506).
- [84] Code is available at <https://github.com/NicoRomeo/DynCarto>.
- [85] Trajectory datasets are available at Zenodo (2025), doi:[10.5281/zenodo.16763216](https://doi.org/10.5281/zenodo.16763216).