Lyapunov exponents

The orbits of a dynamical system may converge to an invariant set called *attractor*[1] as time tends to infinity. For example, a stable node is an attractor of dimension zero (a single point), while a stable periodic orbit (i.e., a stable limit cycle) is an attractor of dimension one. Attractors can have intricate geometric structure with possibly a non-integer (fractal) dimension. Such geometric sets are called *strange attractors*. Dynamics on strange attractors is non-periodic and sensitive to initial conditions, meaning that two trajectories corresponding to initial conditions that are very close tend to diverge in time, while remaining bounded.

Lyapunov exponents are quantities designed to quantify the exponential rates of divergence (or convergence) or nearby orbits in phase space, and can be used to provide qualitative properties such as the type of attractor, the attractor dimension, etc. A system exhibiting bounded dynamics with at least one positive Lyapunov exponent is defined to be *chaotic*, with the magnitude of the largest exponent reflecting the time scale on which system dynamics become unpredictable (predictability horizon). Systems with attractors are usually dissipative[2], i.e., volume-contracting. For dynamical systems whose equations of motion are explicitly known there is a straightforward technique for computing a complete Lyapunov spectrum, i.e., the whole set of Lyapunov exponents.

Lyapunov spectrum

Consider the n-dimensional nonlinear dynamical system

$$\begin{cases} \dfrac{d\boldsymbol{x}}{dt} = \boldsymbol{f}(\boldsymbol{x}) \\ \boldsymbol{x}(0) = \boldsymbol{x}_0 \end{cases} \tag{1}$$

We are interested in studying the long-term evolution of an infinitesimal n-sphere of initial conditions

$$S_\epsilon(\boldsymbol{x}_0) = \{\boldsymbol{x} \in \mathbb{R}^n : \ \|\boldsymbol{x} - \boldsymbol{x}_0\| \leq \epsilon\} \tag{2}$$

as its center is advected by the flow.

For infinitesimal ϵ the sphere becomes a n-ellipsoid due to the locally deforming nature of the flow (see Appendix A). Such local deformation can be fully characterized by the Jacobian of the vector field $\boldsymbol{J_f}(\boldsymbol{X}(t, \boldsymbol{x}_0))$ evaluated on a "fiducial trajectory" solving (1). Such fiducial trajectory is the trajectory of the center $\boldsymbol{x}_0$ of the infinitesimal sphere $S_\epsilon(\boldsymbol{x}_0)$. The i-th Lyapunov exponent for a trajectory starting at $\boldsymbol{x}_0$ is defined in terms of the length of the ellipsoidal principal axis $e_i(t)$ relative to the initial axes

$$\ell_i(\boldsymbol{x}_0) = \lim_{t\to\infty} \frac{1}{t} \log\left(\frac{e_i(t)}{e_i(0)}\right) \qquad i = 1, \ldots, n. \tag{3}$$

Equivalently,

$$\frac{e_i(t)}{e_i(0)} \simeq e^{\ell_i(\boldsymbol{x}_0)t} \tag{4}$$

for sufficiently large t. Hence, the Lyapunov exponents characterize the exponential rate of expansion or contraction of different directions in phase space.

[1] Attractors are never reached exactly in finite time unless the initial condition lies precisely on them. Trajectories starting off the attractor only approach it asymptotically as time tends to infinity. Attractors typically have associated *basins of attraction*, i.e., subsets of the phase space whose trajectories converge to the attractor over time.

[2] Non-dissipative systems such as Hamiltonian systems can exhibit chaotic behavior (e.g., the double pendulum). However, such chaotic behavior is not due to attractors but rather to level sets of invariant energy surfaces being filled.

Clearly, we are free to choose the principal axes of the initial sphere as we like; that is, we can pick any orthogonal basis with vectors of length ϵ, i.e., sitting on the infinitesimal sphere $S_\epsilon(\boldsymbol{x}_0)$. We can also normalize such basis vectors as needed. In fact, the Lyapunov exponents are invariant under simultaneous rescaling of both $e_i(0)$ and $e_i(t)$ by the same factor. For example, let $\{\boldsymbol{e}_1, \ldots, \boldsymbol{e}_n\}$ be vectors on the sphere $S_\epsilon(\boldsymbol{x}_0)$. These vectors, as well as their normalized versions $\boldsymbol{e}_j / \|\boldsymbol{e}_j\|$, evolve according to the same ODEs

$$\begin{cases} \dfrac{d\boldsymbol{\eta}_j}{dt} = \boldsymbol{J}_{\boldsymbol{f}}(\boldsymbol{X}(t,\boldsymbol{x}_0))\boldsymbol{\eta}_j \\ \boldsymbol{\eta}_j(0,\boldsymbol{x}_0) = \boldsymbol{e}_j \end{cases} \quad \Leftrightarrow \quad \begin{cases} \dfrac{d\widehat{\boldsymbol{\eta}}_j}{dt} = \boldsymbol{J}_{\boldsymbol{f}}(\boldsymbol{X}(t,\boldsymbol{x}_0))\widehat{\boldsymbol{\eta}}_j \\ \widehat{\boldsymbol{\eta}}_j(0,\boldsymbol{x}_0) = \dfrac{\boldsymbol{e}_j}{\|\boldsymbol{e}_j\|} \end{cases} \tag{5}$$

with

$$\boldsymbol{\eta}(t,\boldsymbol{x}_0) = \|\boldsymbol{e}_j\|\, \widehat{\boldsymbol{\eta}}_j(t,\boldsymbol{x}_0). \tag{6}$$

The orientation of the infinitesimal ellipsoid centered at $\boldsymbol{X}(t,\boldsymbol{x}_0)$ changes continuously as time evolves. Therefore, we cannot speak of a well-defined direction associated with a given exponent.

Properties of the Lyapunov spectrum

Let us begin with a theorem characterizing Lyapunov exponents at hyperbolic fixed points.

Theorem 1. Let $\boldsymbol{x}^*$ be a hyperbolic fixed point of (1). The Lyapunov exponents at the fixed point coincide with the real part of the eigenvalues of the Jacobian $\boldsymbol{J}_{\boldsymbol{f}}(\boldsymbol{x}^*)$.

Proof. Consider the infinitesimal sphere $S_\epsilon(\boldsymbol{x}^*)$ centered at the fixed point. The linearized flow in a neighborhood of $\boldsymbol{x}^*$ can be written as

$$X(t,\boldsymbol{x}_0) = \boldsymbol{x}^* + \boldsymbol{\eta}(t,\boldsymbol{x}_0) \tag{7}$$

where

$$\boldsymbol{\eta}(t,\boldsymbol{x}_0) = \boldsymbol{P}_R e^{t\boldsymbol{J}_R}\boldsymbol{P}_R^{-1}\boldsymbol{\eta}_0, \tag{8}$$

for all $\boldsymbol{\eta}_0$ in a neighborhood of $\boldsymbol{0}$. Here $\boldsymbol{J}_R$ is the real Jordan form of $\boldsymbol{J}_{\boldsymbol{f}}(\boldsymbol{x}^*)$. Of course, $\boldsymbol{\eta}(t,\boldsymbol{x}_0)$ is a linear combination of eigenvectors/generalized eigenvectors $\boldsymbol{v}_j$ (columns of $\boldsymbol{P}_R$) multiplied by the entries of $e^{t\boldsymbol{J}_R}$. The coefficients of the linear combination are the entries of $\boldsymbol{P}_R^{-1}\boldsymbol{\eta}_0$.

Let us now show that the Lyapunov exponent associated with the dynamics in a generalized eigenspace coincides with the real part of the eigenvalue corresponding to such generalized eigenspace. First let $\boldsymbol{v}_j(0)$ be a normalized eigenvector of $\boldsymbol{J}_R$ corresponding to a real eigenvalue λ_j. We know that $\boldsymbol{v}_j(t) = e^{\lambda_j}\boldsymbol{v}_j(0)$. In this case we have

$$\ell_j(\boldsymbol{x}_0) = \lim_{t\to\infty}\frac{1}{t}\log\frac{\|e^{\lambda_j t}\boldsymbol{v}_j(0)\|}{\|\boldsymbol{v}_j(0)\|} = \lim_{t\to\infty}\frac{\log(e^{t\lambda_j})}{t} = \lambda_j.$$

If λ_j is a complex eigenvalue (with its conjugate one) then

$$\boldsymbol{v}_j(t) = e^{t\,\mathrm{Re}(\lambda_j)}\left[\cos(\mathrm{Im}(\lambda_j)t)\boldsymbol{u}_j + \sin(\mathrm{Im}(\lambda_j)t)\boldsymbol{w}_j\right].$$

Here $\boldsymbol{u}_j$ and $\boldsymbol{w}_j$ represent the real and the imaginary parts of the complex eigenvector corresponding λ_j. In this case we have

$$\begin{aligned}\ell_j(\boldsymbol{x}_0) &= \lim_{t\to\infty}\frac{1}{t}\log\left(\frac{\left\|e^{t\,\mathrm{Re}(\lambda_j)}\left[\cos(\mathrm{Im}(\lambda_j)t)\boldsymbol{u}_j + \sin(\mathrm{Im}(\lambda_j)t)\boldsymbol{w}_j\right]\right\|}{\|\boldsymbol{v}_j(0)\|}\right)\\ &= \lim_{t\to\infty}\frac{1}{t}\log\left(e^{t\,\mathrm{Re}(\lambda_j)}\right) + \lim_{t\to\infty}\frac{1}{t}\log\left(\frac{\|\cos(\mathrm{Im}(\lambda_j)t)\boldsymbol{u}_j + \sin(\mathrm{Im}(\lambda_j)t)\boldsymbol{w}_j\|}{\|\boldsymbol{v}_j(0)\|}\right)\\ &= \mathrm{Re}(\lambda_j).\end{aligned} \tag{9}$$

If λ_j is real and degenerate, then

$$\boldsymbol{v}_j(t) = e^{t\lambda_j}\boldsymbol{u}_j + te^{t\lambda_j}\boldsymbol{w}_j.$$

Here $\boldsymbol{u}_j$ is an eigenvector while $\boldsymbol{w}_j$ is a generalized eigenvector. We have,

$$\ell_j(\boldsymbol{x}_0) = \lim_{t\to\infty}\frac{1}{t}\log\left(\frac{\left\|e^{t\lambda_j}\boldsymbol{u}_j + te^{t\lambda_j}\boldsymbol{w}_j\right\|}{\|\boldsymbol{v}_j(0)\|}\right) = \lim_{t\to\infty}\frac{\log(e^{\lambda_j t})}{t} + \lim_{t\to\infty}\frac{1}{t}\log\left(\frac{\|\boldsymbol{u}_j + t\boldsymbol{w}_j\|)}{\|\boldsymbol{v}_j(0)\|}\right) = \lambda_j.$$

The last equality holds because the numerator in the second term grows logarithmically (like $\log(t)$) and thus vanishes as $t\to\infty$. Cases involving higher-order eigenvalue degeneracies can be treated similarly.

□

Next, we discuss a theorem characterizing Lyapunov exponents for initial conditions $\boldsymbol{x}_0$ in an invariant set that does not include any fixed points.

Theorem 2. Consider the dynamical system (1) and suppose that there exists a bounded invariant set[3] $\mathcal{M}\subset\mathbb{R}^n$ with no fixed points in it. Then for any initial condition $\boldsymbol{x}_0\in\mathcal{M}$, the system has at least one zero Lyapunov exponent.

Proof. Let $\boldsymbol{X}(t,\boldsymbol{x}_0)$ be a trajectory the flow of (1) corresponding to $\boldsymbol{x}_0\in\mathcal{M}$. Consider the evolution of a vector $\boldsymbol{v}(t)$ centered at $\boldsymbol{X}(t,\boldsymbol{x}_0)$ that is tangent to the flow, i.e.,

$$\boldsymbol{v}(t) = \frac{\partial}{\partial t}\boldsymbol{X}(t,\boldsymbol{x}_0) = \boldsymbol{f}(\boldsymbol{X}(t,\boldsymbol{x}_0)).$$

Of course, we can rescale the vector $\boldsymbol{v}(t)$ so that $\|\boldsymbol{v}(0)\| = 1$ (or rescale it to be on the infinitesimal n-sphere $S_\epsilon(\boldsymbol{x}_0)$). By differentiating both sides with respect to time we obtain

$$\frac{d\boldsymbol{v}(t)}{dt} = J_{\boldsymbol{f}}(\boldsymbol{X}(t,\boldsymbol{x}_0))\boldsymbol{f}(\boldsymbol{X}(t,\boldsymbol{x}_0)) = \boldsymbol{J}_{\boldsymbol{f}}(\boldsymbol{X}(t,\boldsymbol{x}_0))\boldsymbol{v}(t),$$

by the chain rule. Thus, $\boldsymbol{v}(t) = \boldsymbol{f}(\boldsymbol{X}(t,\boldsymbol{x}_0))$ solves the linearized system. Let us now compute the Lyapunov exponent on the fiducial trajectory $\boldsymbol{X}(t,\boldsymbol{x}_0)$ corresponding to a vector that is tangent to the flow

$$\ell(\boldsymbol{x}_0) = \lim_{t\to\infty}\frac{1}{t}\log\left(\frac{\|\boldsymbol{v}(t)\|}{\|\boldsymbol{v}(0)\|}\right) = \lim_{t\to\infty}\frac{1}{t}\log\left(\frac{\|\boldsymbol{f}(\boldsymbol{X}(t,\boldsymbol{x}_0))\|}{\|\boldsymbol{f}(\boldsymbol{x}_0)\|}\right).$$

Since $\boldsymbol{f}$ is smooth and $\boldsymbol{X}(t,\boldsymbol{x}_0)\in\mathcal{M}$, i.e., $\boldsymbol{X}(t,\boldsymbol{x}_0)$ is bounded, we have that

$$\ell(\boldsymbol{x}_0) = 0.$$

For dissipative systems it can be shown that one of the principal axes of the infinitesimal ellipsoid locks in the direction of the flow, i.e., at least of the singular values of the matrix $\boldsymbol{F}$ in Appendix A converges to one. For energy-preserving systems there are are usually multiple zero Lyapunov exponents and the ellipsoid semi-axies may not necessarily align with the direction ofthe flow.

□

Theorem 3. Let $\boldsymbol{X}(t,\boldsymbol{x}_0)$ be the flow generated by a smooth vector field $\boldsymbol{f}(\boldsymbol{x})$, and let

$$J(t,\boldsymbol{x}_0) = \det\left(\frac{\partial \boldsymbol{X}(t,\boldsymbol{x}_0)}{\partial \boldsymbol{x}_0}\right) \tag{11}$$

denote the Jacobian determinant of the flow with respect to the initial condition. Then the sum of the Lyapunov exponents at $\boldsymbol{x}_0$ satisfies

$$\sum_{i=1}^{n} \ell_i(\boldsymbol{x}_0) = \lim_{t\to\infty} \frac{1}{t}\log(J(t,\boldsymbol{x}_0)). \tag{12}$$

Proof. Let $\boldsymbol{X}(t,\boldsymbol{x}_0)$ be the flow generated by the vector field $\boldsymbol{f}(\boldsymbol{x})$, and let

$$\boldsymbol{F}(t,\boldsymbol{x}_0) = \frac{\partial \boldsymbol{X}(t,\boldsymbol{x}_0)}{\partial \boldsymbol{x}_0} \tag{13}$$

be the Jacobian matrix of the flow with respect to the initial condition. We have seen that

$$\frac{d}{dt}\boldsymbol{F}(t,\boldsymbol{x}_0) = \boldsymbol{J}_{\boldsymbol{f}}(\boldsymbol{X}(t,\boldsymbol{x}_0))\boldsymbol{F}(t,\boldsymbol{x}_0), \qquad \boldsymbol{F}(0,\boldsymbol{x}_0) = \boldsymbol{I},$$

where $\boldsymbol{J}_{\boldsymbol{f}}(\boldsymbol{X}(t,\boldsymbol{x}_0))$ is the Jacobian matrix of the vector field $\boldsymbol{f}(\boldsymbol{x})$. Consider an infinitesimal sphere of initial conditions $S_\epsilon(\boldsymbol{x}_0)$ centered at $\boldsymbol{x}_0$. In Appendix A we show that this sphere is mapped by the flow to an ellipsoid whose principal semi-axes are determined by the singular values $\sigma_1(t) \geq \sigma_2(t) \geq \cdots \geq \sigma_n(t) > 0$ of $\boldsymbol{F}(t,\boldsymbol{x}_0)$. Using the change of variable theorem in multivariate integration we have

$$V(t) = V(0)J(t,\boldsymbol{x}_0) \tag{14}$$

where $V(0)$ is volume of the infinitesimal n-sphere. and $V(t)$ is the volume of the ellipsoid. The Jacobian determinant $J(t,\boldsymbol{x}_0)$ is the product of the singular values of $\boldsymbol{F}(t,\boldsymbol{x}_0)$. This yields

$$V(t) = V(0)\prod_{i=1}^{n}\sigma_i(t).$$

Moreover, $\sigma_j(t)$ represents the length of the j-th semi-axis of the infinitesimal ellipsoid (Corollary 1 in Appendix A). Therefore,

$$\ell_j(\boldsymbol{x}_0) = \lim_{t\to\infty}\frac{1}{t}\log\sigma_j(t), \qquad j = 1,\ldots,n$$

Taking the logarithm of $V(t)/V(0)$ yields

$$\log(J(t,\boldsymbol{x}_0)) = \sum_{i=1}^{n}\log\sigma_i(t).$$

Dividing by t and taking the limit yields

$$\lim_{t\to\infty}\frac{1}{t}\log(J(t,\boldsymbol{x}_0)) = \sum_{j=1}^{n}\lim_{t\to\infty}\frac{1}{t}\log\sigma_j(t),$$

where the interchange of limit and summation is justified because each limit exists individually. Therefore,

$$\sum_{j=1}^{n}\ell_j(\boldsymbol{x}_0) = \lim_{t\to\infty}\frac{1}{t}\log(J(t,\boldsymbol{x}_0)),$$

which completes the proof.

□

Recalling the Liouville theorem

$$J(t, \boldsymbol{x}_0) = \exp\left[\int_0^t \nabla \cdot \boldsymbol{f}(\boldsymbol{X}(\tau, \boldsymbol{x}_0))d\tau\right] \tag{15}$$

we can write (12) as

$$\sum_{i=1}^{n} \ell_i(\boldsymbol{x}_0) = \lim_{t\to\infty} \frac{1}{t}\int_0^t \nabla \cdot \boldsymbol{f}(\boldsymbol{X}(\tau, \boldsymbol{x}_0))d\tau. \tag{16}$$

Hence, for any dissipative (volume contracting) dynamical system, we have that *at least one Lyapunov exponent must be negative*, and the sum of all Lyapunov exponents is strictly negative.

The exponential expansion associated with a positive Lyapunov exponent is incompatible with motion on a bounded attractor unless some sort of folding process merges widely separated trajectories. Each positive exponent reflects a "direction" in which the system experiences the repeated stretching and folding that decorrelates nearby states on the attractor.

Computation of Lyapunov exponents

The Lyapunov exponents are defined by the long-term evolution of the axes of an infinitesimal sphere of states. To compute them we proceed as follows:

1. We first compute a fiducial trajectory $\boldsymbol{X}(t, \boldsymbol{x}_0)$ by integrating the nonlinear equations of motion (1) for some "post-transient" initial condition $\boldsymbol{x}_0$. This allows us to make sure that we pick an initial condition that is reasonably close to the attractor.

2. Simultaneously, we integrate the linearized equations of motion (5), hereafter rewritten for convenience,

$$\frac{d\boldsymbol{\eta}_j}{dt} = \boldsymbol{J}_{\boldsymbol{f}}(\boldsymbol{X}(t, \boldsymbol{x}_0))\boldsymbol{\eta}_j \tag{17}$$

   for initial conditions defining an arbitrarily oriented frame of n-orthonormal vectors $\{\boldsymbol{\eta}_1(0), \ldots, \boldsymbol{\eta}_n(0)\}$. Each vector $\boldsymbol{\eta}_j(t)$ may diverge in magnitude and will tends to fall along the local direction of most rapid growth (see Figure 2 and Figure 6). The collapse toward a common direction causes the orientation of all axis vectors $\boldsymbol{\eta}_j(t)$ to eventually become indistinguishable. These two problems can be overcome by the repeated use of the Gram- Schmidt (or QR) orthonormalization procedure on the vector frame. To this end, we integrate (17) forward in time for some time and then we orthonormalize the vectors $\boldsymbol{\eta}_j(t)$ relative an arbitrary vector[4] $\boldsymbol{\eta}_1$

$$\widehat{\boldsymbol{\eta}}_1 = \boldsymbol{\eta}_1, \qquad \boldsymbol{w}_1 = \frac{\widehat{\boldsymbol{\eta}}_1}{\|\widehat{\boldsymbol{\eta}}_1\|}, \tag{18}$$

$$\widehat{\boldsymbol{\eta}}_i = \boldsymbol{\eta}_i - \sum_{j=1}^{i-1}\langle \boldsymbol{\eta}_i, \boldsymbol{w}_j\rangle \boldsymbol{w}_j \qquad \boldsymbol{w}_i = \frac{\widehat{\boldsymbol{\eta}}_i}{\|\widehat{\boldsymbol{\eta}}_i\|}, \qquad \text{for } i = 2, \ldots, n. \tag{19}$$

   The frequency of re-orthonormalization is not critical, so long as neither the magnitude nor the orientation divergences have exceeded computer limitations.

3. The projection of the vectors $\boldsymbol{\eta}_j$ onto the new orthonormal frame correctly updates the *rates of growth* of each of the ellipsoid principal axes. Such projections are given by $\widehat{\boldsymbol{\eta}}_j$ (before normalization). Each time we perform Gram-Schmidt we record the stretching factors

$$S_i \leftarrow S_i + \log\|\widehat{\boldsymbol{\eta}}_i\|_2, \quad \text{for each } i = 1, \ldots, n \tag{20}$$

[4] The Gram Schmidt orthogonalization procedure never affects the direction of the first vector in the system. Hence this vector tends to seek out the direction which is most rapidly growing.

and restart integration of both the fiducial trajectory and the linearized equation (17) from the new orthonormal initial condition $\{\boldsymbol{w}_1, \ldots, \boldsymbol{w}_n\}$.

4. After evolving the system up to total time T, the Lyapunov exponents are approximated by

$$\ell_i \approx \frac{S_i}{T}, \quad \text{for each } i = 1, \ldots, n. \tag{21}$$

Each exponent ℓ_i corresponds to the average logarithmic growth rate along an orthogonal direction.

Remark: This procedure ensures that the first vector $\widehat{\boldsymbol{\eta}}_1$ (or $\boldsymbol{w}_1$) aligns with the direction of maximal expansion, corresponding to the largest Lyapunov exponent. In dissipative systems, the dynamics naturally cause one of the other orthonormal vectors to lock onto the flow direction asymptotically in time (see Figure 2). Such direction is characterized by neither expansion nor contraction (zero Lyapunov exponent).

Remark: To better understand step 3 of the algorithm, let $\boldsymbol{Q}_0$ denote the initial orthonormal frame (for example, $\boldsymbol{Q}_0 = \boldsymbol{I}$). Let $\boldsymbol{A}_1$ be the linear operator (propagator) that evolves the frame $\boldsymbol{Q}_0$ according to the linear equation (17), producing the matrix $\boldsymbol{Y}_1$ whose columns are the solutions at time t_1 obtained by integrating the initial condition $\boldsymbol{Q}_0$. We have

$$\boldsymbol{Y}_1 = \boldsymbol{A}_1 \boldsymbol{Q}_0. \tag{22}$$

The columns of $\boldsymbol{Y}_1$ are the vectors $\boldsymbol{\eta}_j(t_1)$ defined in equation (17). Let us do a QR of $\boldsymbol{Y}_1$

$$\boldsymbol{Y}_1 = \boldsymbol{Q}_1 \boldsymbol{R}_1,$$

and use $\boldsymbol{Q}_1$ as the new orthonormal initial frame to integrate the system forward from t_1 to t_2. The columns of $\boldsymbol{Q}_1$ are $\boldsymbol{w}_j(t_1)$ (orthonormal) in previous notation. The diagonal entries of $\boldsymbol{R}_1$ are the norms of $\widehat{\boldsymbol{\eta}}(t_1)$ (prior to normalization), i.e.,

$$[\boldsymbol{R}_1]_{jj} = \left\| \widehat{\boldsymbol{\eta}}_j(t_1) \right\|_2. \tag{23}$$

Let $\boldsymbol{A}_2$ denote the propagator from t_1 to t_2. Then the resulting matrix $\boldsymbol{Y}_2$ is given by

$$\boldsymbol{Y}_2 = \boldsymbol{A}_2 \boldsymbol{Q}_1 \boldsymbol{R}_1.$$

We now compute a second QR decomposition of $\boldsymbol{A}_2 \boldsymbol{Q}_1$, yielding

$$\boldsymbol{Y}_2 = \boldsymbol{Q}_2 \boldsymbol{R}_2 \boldsymbol{R}_1. \tag{24}$$

The columns of $\boldsymbol{Y}_2$ are the vectors $\boldsymbol{\eta}_j(t_2)$. The diagonal entries of $\boldsymbol{R}_2$ represent the scaling factors by which the corresponding directions grow in that step, i.e.,

$$[\boldsymbol{R}_2]_{jj} = \|\widehat{\boldsymbol{\eta}}(t_2)\|_2$$

Furthermore, the product $\boldsymbol{R}_2 \boldsymbol{R}_1$ is upper triangular and has as diagonal entries the product of the diagonal entries of $\boldsymbol{R}_1$ and $\boldsymbol{R}_1$. Hence,

$$\log\left([\boldsymbol{R}_2]_{jj}[\boldsymbol{R}_1]_{jj}\right) = \log\left([\boldsymbol{R}_2]_{jj}\right) + \log\left([\boldsymbol{R}_1]_{jj}\right). \tag{25}$$

For the first vector (the one pointing towards the direction of largest stretching) (24) reduces to

$$\boldsymbol{\eta}_1(t_2) = \boldsymbol{w}_1(t_2) \left\|\widehat{\boldsymbol{\eta}}_1(t_2)\right\| \left\|\widehat{\boldsymbol{\eta}}_1(t_1)\right\|, \tag{26}$$

which implies

$$\log(\|\boldsymbol{\eta}_1(t_2)\|) = \log(\|\widehat{\boldsymbol{\eta}}_1(t_2)\|) + \log(\|\widehat{\boldsymbol{\eta}}_1(t_1)\|). \tag{27}$$

This justifies why summing the logarithms of the norms of the vectors in the orthonormal frame (prior to normalization), as in (20), provides an approximation to the logarithm of the norm of $\boldsymbol{\eta}_j(t)$ for arbitrary t. In particular, for long times

$$\sum_{k=1}^{N} \log\left([\boldsymbol{R}_k]_{jj}\right)$$

approximates $\log \|\boldsymbol{\eta}_j(t)\|_2$, where $t = t_N$.

Remark: Lyapunov exponents are not local quantities in either the spatial or temporal sense. Each exponent arises from the average of the local deformation of various phase space directions. Each is determined by the long-time evolution of a single volume element.

Lyapunov time and predictability horizon

The Lyapunov exponents, in particular the largest postitive one ℓ_1, measure the rate at which nonlinear systems destroy information. In fact, if an initial point is specified with an accuracy δ_0, then the future behavior of the system could not be accurately predicted if t exceeds a certain value know as *predictability horizon*. If two trajectories in phase space are initially separated by a small perturbation δ_0 and the system has a positive Lyapunov exponent, then their separation tends to grow exponentially in time

$$\delta(t) \approx \delta_0 e^{\ell_1 t}.$$

- The *Lyapunov time* T_L is defined to be the time after which the separation grows by a factor $e = 2.718281828$, i.e., $\delta(T_L) \approx \delta_0 e$. This yields

$$T_L = \frac{1}{\ell_1}.$$

- The *predictability horizon* T_p is defined as the maximum time up to which predictions remain accurate, given a maximum acceptable error $\delta_{\max}$. Using the same steps as above, we obtain

$$\delta_{\max} \sim \delta_0 e^{\ell_1 T_p},$$

i.e.,

$$T_p \approx \frac{1}{\ell_1} \log\left(\frac{\delta_{\max}}{\delta_0}\right).$$

Remark: Larger (positive) Lyapunov exponents correspond to shorter predictability horizons. Note that the units of T_L and T_p depend on the units of time used in the system. If ℓ_1 is measured in inverse seconds then T_L and T_p will be in seconds.

Attractor dimension

The geometry of an attractor in a dissipative system, and in particular a strange attractor, can be characterized using various notions of *fractal dimension*, such as the box-counting dimension, the Kaplan–Yorke dimension, the Housdorff dimension, and the information dimension. Under certain technical assumptions, these different notions of fractal dimension coincide.

Box dimension. Let $A \subset \mathbb{R}^n$ be and invariant set that is attracting trajectories of a dynamical system. The *box dimension* (also called the box-counting dimension) of $A \subset \mathbb{R}^n$ is defined by the following limiting

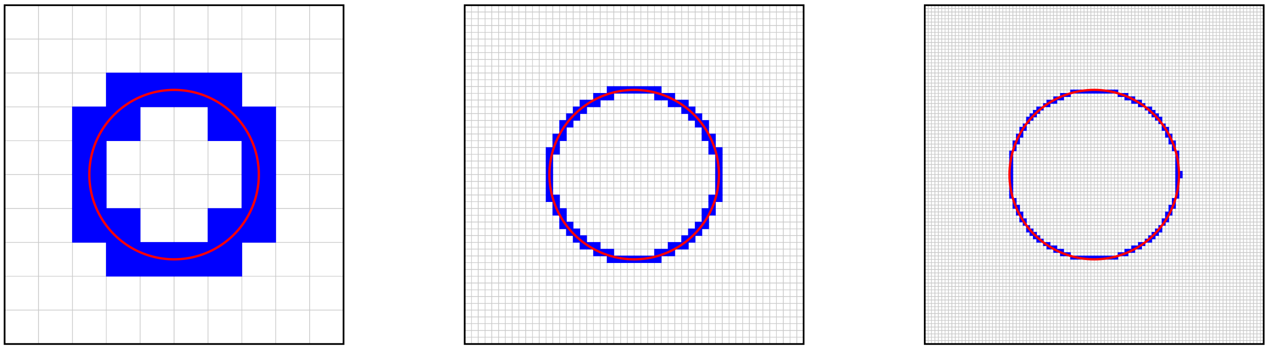

Figure 1: Example calculation of box dimension for a circle. The smallest number of boxes to cover the circle is illustrated in blue. The ratio $\log(N(\varepsilon))/\log(1/\varepsilon)$ is $\{1.3, 1.18, 1.13\}$ from left to right .

procedure: Cover the set A with a grid of boxes (cubes) $B_i(\varepsilon)$ of side length $\varepsilon > 0$; Let $N(\varepsilon)$ be the smallest number of boxes needed to cover A. The *box dimension* of A is defined as

$$D_{\text{box}}(A) = \lim_{\varepsilon\to 0} \frac{\log N(\varepsilon)}{\log(1/\varepsilon)},$$

if the limit exists (otherwise, use the lim sup or lim inf). For a smooth curve (e.g., a periodic orbit), $N(\varepsilon) \sim 1/\varepsilon$ (see Figure 1). For a smooth surface, $N(\varepsilon) \sim 1/\varepsilon^2$, so $D_{\text{box}} = 2$. For a fractal, $N(\varepsilon) \sim 1/\varepsilon^d$ with non-integer d, leading to a non-integer box dimension.

Kaplan-Yorke (Lyapunov) dimension. The Kaplan-Yorke (KY) dimension of an attractor is defined in terms of the Lyapunov spectrum as follows

$$D_{KY} = j + \frac{1}{|\ell_{j+1}|} \sum_{i=1}^{j} \ell_i, \tag{28}$$

where j is defined by the condition that

$$\sum_{i=1}^{j} \ell_i \geq 0 \qquad \sum_{i=1}^{j+1} \ell_i < 0. \tag{29}$$

The calculation of dimension from this equation requires knowledge of all but the most negative Lyapunov exponents.

It has been shown that for smooth flows on *uniformly hyperbolic attractors*, the *Kaplan–Yorke (KY) dimension* coincides with several other notions of fractal dimension, such as the box-counting dimension, the Hausdorff dimension, and the information dimension[5]. A compact invariant set $\mathcal{M} \subset \mathbb{R}^n$ is called

[5]The *information dimension* of an attractor, denoted by D_{info}, quantifies how the probability mass is distributed over the attractor with respect to scale. Let μ be an invariant probability measure supported on the attractor, and consider a cover of the attractor by boxes $\{B_i(\varepsilon)\}$ of side length ε. Define

$$\mu_i = \mu(B_i(\varepsilon)),$$

the measure of the i-th box. The information dimension is given by

$$D_{\text{info}} = \lim_{\varepsilon\to 0} \frac{H(\varepsilon)}{\log(1/\varepsilon)}, \qquad \text{where} \qquad H(\varepsilon) = -\sum_i \mu_i \log \mu_i$$

is the Shannon entropy of the coarse-grained measure at scale ε.

uniformly hyperbolic if for each $\boldsymbol{x} \in \mathcal{M}$, the tangent space admits the splitting

$$T_{\boldsymbol{x}}\mathcal{M} = E^s_{\boldsymbol{x}} \oplus E^u_{\boldsymbol{x}} \oplus E^0_{\boldsymbol{x}},$$

where E^s and E^u are the stable and unstable subspaces, and E^0 is the flow direction. Vectors in E^s are contracted under the flow, and vectors in E^u are expanded. Uniformly hyperbolic attractors admit a *Sinai–Ruelle–Bowen (SRB) measure*, that is, an invariant probability measure μ on $\mathcal{M}$ such that for almost every initial condition $\boldsymbol{x}_0 \in \mathcal{M}$ (except for sets with zero Lebesgue measure), the time averages of continuous observables φ converge to space averages with respect to μ. Formally, for every continuous function $\varphi : \mathcal{M} \to \mathbb{R}$

$$\lim_{T\to\infty} \frac{1}{T}\int_0^T \varphi(\boldsymbol{X}(t,\boldsymbol{x}_0))\, dt = \int_M \varphi(\boldsymbol{x})\, d\mu(\boldsymbol{x}),$$

for almost every initial condition $\boldsymbol{x}_0 \in \mathcal{M}$, where $\boldsymbol{X}(t,\boldsymbol{x}_0)$ denotes the flow starting at $\boldsymbol{x}_0$.

Qualitative description of attractors using Lyapunov exponents. Lyapunov exponents can be used to describe qualitatively attractors. For instance, in a three-dimensional dissipative dynamical system the only possible cases Lyapunov spectra (hereafter summarized by the signs of the exponents), and the attractors they describe, are as follows:

- strange attractor $(+, 0, -)$
- stable limit cycle $(0, -, -)$
- stable node $(-, -, -)$

In four dimensions there are, in principle, three possible "attractors" $(+, 0, 0, -)$, $(+, +, 0, -)$, and $(+, 0, -, -)$. The first one may be either a dissipative system or a Hamiltonian system.

Numerical examples

In this section, we present three numerical examples illustrating the computation of Lyapunov exponents.

Van der Pol oscillator. Consider the Van der Pol Oscillator

$$\ddot{x} - \mu(1 - x^2)\dot{x} + x = 0 \tag{30}$$

As is well known, this system can be written as a system of two first-order ODEs as

$$\begin{cases} \dot{x}_1 = x_2, \\ \dot{x}_2 = \mu(1 - x_1^2)x_2 - x_1, \end{cases} \tag{31}$$

In Figure 2 we plot the phase portrait associated with this system, together with the temporal dynamics of a circle volume of initial conditions close to the globally attracting *stable limit cycle* that attracts all orbits. The circle evolves according to equation (17). As easily seen, as time increases, the circle becomes an ellipse which then shrinks to zero, and two vectors in black (initially orthogonal) tend to align with the direction of the largest Lyapunov exponent, which in this case is zero (see Figure 3. In other words, the vectors tend to align to the flow as time increases, as clearly demonstrated in Figure 2. In Figure 3 we plot the Lyapunov exponents computed with the Gram-Schmidt algorithm described above. We perform re- orthogonalization every 0.5 time units. The Lyapunov exponents we obtained with this procedure are

$$\ell_1 = 0.000345 \qquad \ell_2 = -1.060644. \tag{32}$$

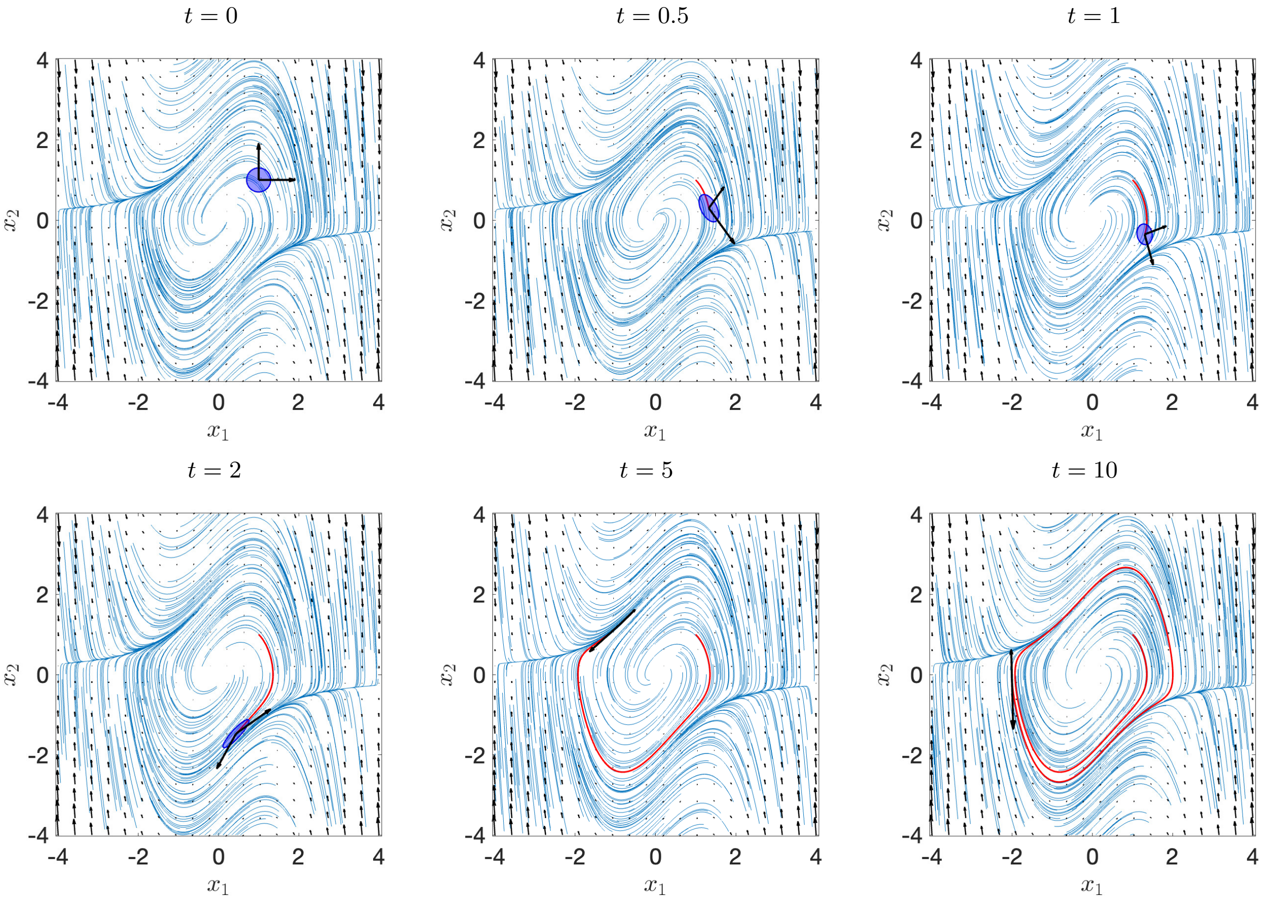

Figure 2: Phase portrait of the Van der Pol oscillator (31) for $\mu = 2$. It is seen that as time increases the initial circle of initial conditions (governed by (17)) becomes an ellipse and shrinks to zero. The two vectors in black (initially orthogonal) tend to align with the direction of the largest Lyapunov exponent, which in this case is zero (see Figure 3). This implies that the vectors tend to align to the flow as time increases. This motivates the Gram-Schmidt orthogonalization procedure.

The dimension of the attractor can be estimated using the Kaplan–Yorke formula (28), which in this case yields

$$D_{KY} \simeq 1.$$

This means that the attractor is one-dimensional, i.e., the curve in $\mathbb{R}^2$ representing the limit cycle shown in Figure 2.

Lorenz-63 system. Consider the Lorenz-63 system

$$\begin{cases} \dot{X} = -\sigma X + \sigma Y \\ \dot{Y} = rX - Y - XZ \\ \dot{Z} = XY - bZ \end{cases} \tag{33}$$

with parameters $\sigma = 10$ and $b = 8/3$ and $r = 28$. We set the initial condition[6] as $(X_0, Y_0, Z_0) = \begin{bmatrix} 5 & 5 & 5 \end{bmatrix}$. With these parameters the Lorenz system (33) exhibits an aperiodic behavior that is very sensitive to initial condition (see Figure 4).

[6]The Lyapunov exponents are independent of the set of initial condition

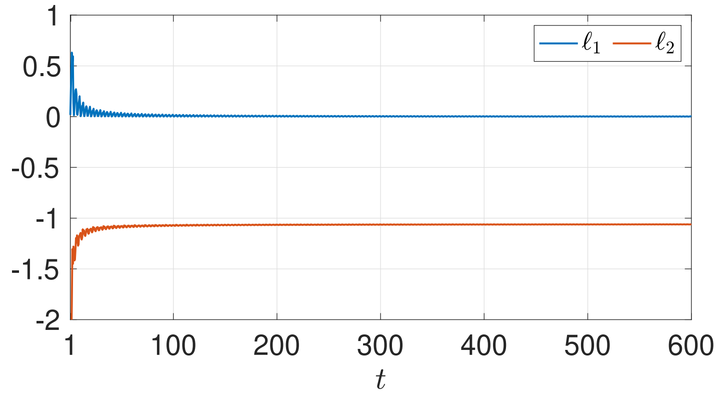

Figure 3: Lyapunov exponents for the Van der Pol system (31). The Kaplan–Yorke dimension of the attractor, computed using (28), is $D_{KY} \simeq 1$.

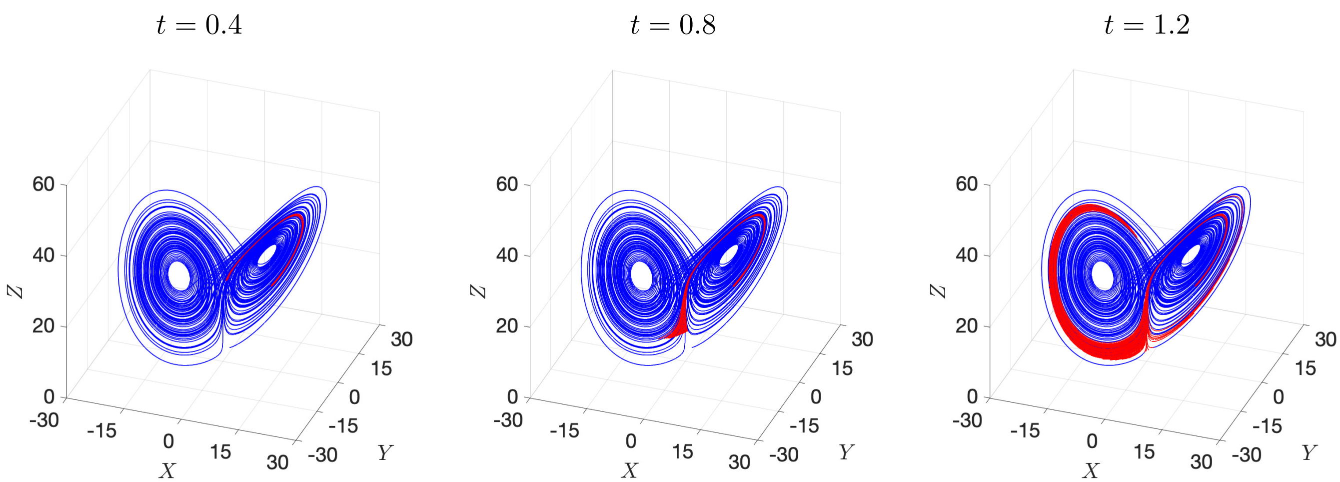

Figure 4: Lorenz's attractor for $r = 28$, $\sigma = 10$ and $b = 8/3$, sensitivity to initial conditions. Shown are trajectories of (33) corresponding to an infinitesimal ball of initial conditions placed nearby the attractor. As time increases such small ball of red initial conditions paints the entire attractor.

In Figure 5 we plot the Lyapunov exponents computed with the Gram-Schmidt algorithm described above. We perform re-orthogonalization every 0.5 time units. The Lyapunov exponents we obtained with this procedure are

$$\ell_1 = 0.901 \qquad \ell_2 = 0.001 \qquad \ell_3 = -14.566. \tag{34}$$

This indicates that the system has a strange attractor characterized by a strongly contracting direction associated with ℓ_3, a zero Lyapunov exponent corresponding to the vector tangent to the flow, and an expanding direction. The sum of all Lyapunov exponents equals the time average of the divergence of the vector field (see Eq. (16)). Specifically, we obtain

$$\sum_{j=1}^{3} \ell_j = -13.6631 \tag{35}$$

which is in excellent agreement with the analytical value[7] -13.666. The dimension of the strange attractor

[7]The divergence of the vector field for the Lorenz system (33) is given by

$$\nabla \cdot \boldsymbol{f}(\boldsymbol{x}) = -(\sigma + b + 1). \tag{36}$$

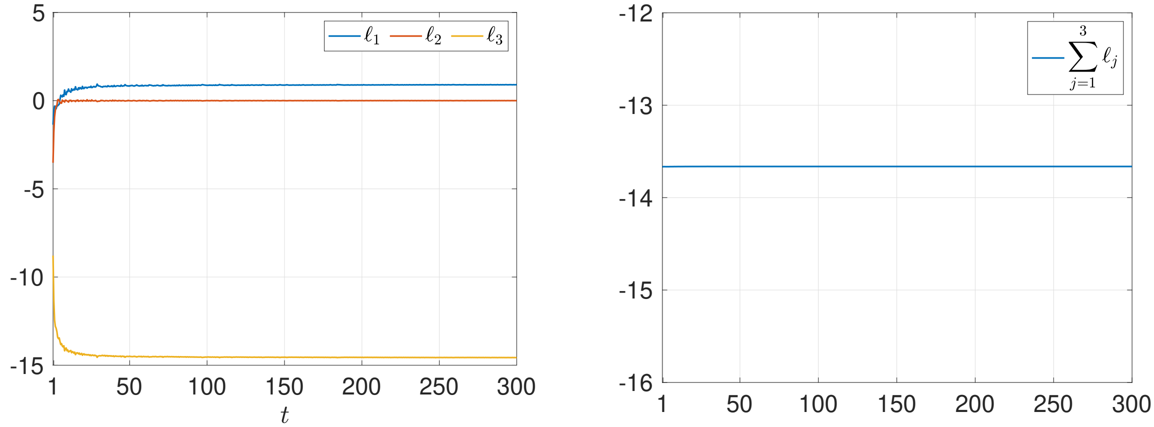

Figure 5: Lyapunov exponents for the Lorenz system (33). The Kaplan–Yorke dimension of the Lorenz attractor, computed using (28), is $D_{KY} = 2.057$. The plot also shows the sum of all Lyapunov exponents, which corresponds to the temporal average of the divergence of the vector field $\boldsymbol{f}(\boldsymbol{x})$ (see Eq. (16)).

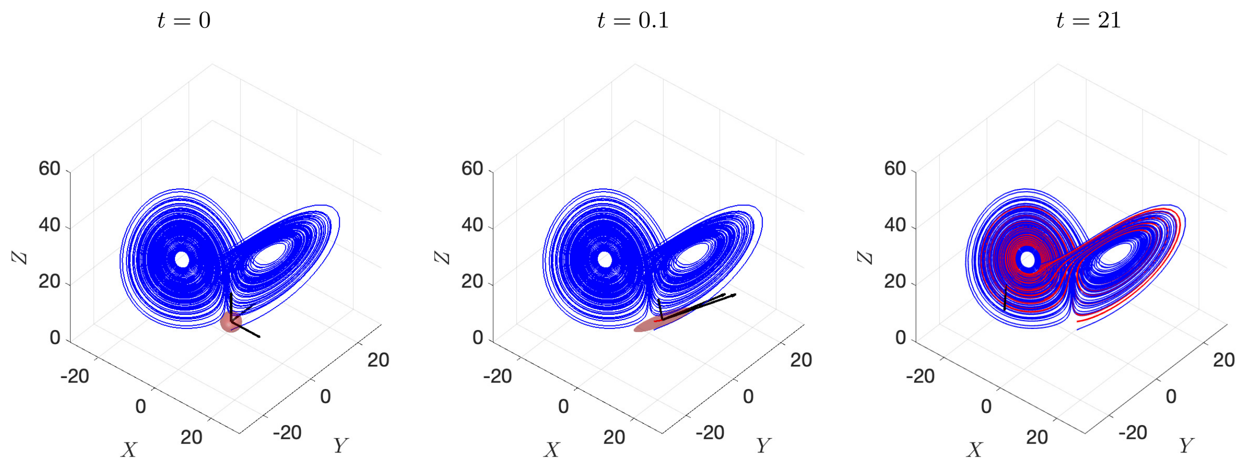

Figure 6: Lorenz system (33). Temporal evolution of an orthonormal basis of vectors (using (17)) along a trajectory of (33). It is seen that as time evolves all vectors collapse to the direction of largest expansion which, in this case is transverse to the flow.

can be estimated using the Kaplan–Yorke formula (28), yielding

$$D_{KY} = 2.062.$$

Thus, the attractor of the Lorenz 63 system is "almost a surface". If we take a section of the attractor that is transverse to the flow we see indeed that there is a tiny bit of roughness (see Figure 7), i.e., we do not have a smooth curve. Such roughness is responsible for the fractal dimension of the attractor.

The Lyapunov time is $T_L = 1.11$ (in Lorenz-63 time units). The predictability horizon, defined as the time required for an initial error δ_0 to grow to $\delta_{\max} = 100\delta_0$ (a one-hundred-fold amplification), is given

Hence,

$$\sum_{j=1}^{3} \ell_j = \lim_{t\to\infty} \frac{1}{t} \int_0^t \nabla \cdot \boldsymbol{f}(\boldsymbol{X}(\tau, \boldsymbol{x}_0))d\tau = -(\sigma + b + 1) = -\left(10 + \frac{8}{3} + 1\right) = -13.666 \tag{37}$$

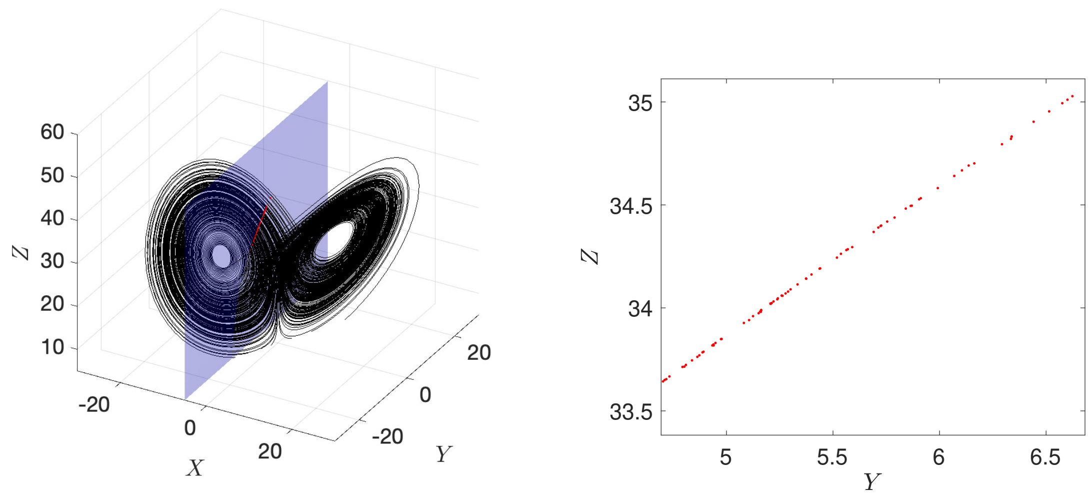

Figure 7: Poincaré section (plane $X = -5$) of the Lorenz attractor . It is seen that points on the section (right) have a tiny bit of roughness, i.e., they do not exactly sit on a one-dimensional manifold but rather on a very thin set.

by

$$T_p = \frac{\log(100)}{\ell_1} = 4.19$$

This means that two trajectories corresponding to initial conditions that are 10^{-3} apart (length of the vector that connects the two initial conditions) will be approximately 10^{-1} apart after $T_p = 4.19$ time units.

Kuramoto-Sivashinsky equation. Consider the initial-boundary value problem for the Kuramoto-Sivashinsky equation in the periodic spatial domain $[-L, L]$

$$\begin{cases} \dfrac{\partial u}{\partial t} + u\dfrac{\partial u}{\partial y} + \dfrac{\partial^2 u}{\partial y^2} + \dfrac{\partial^4 u}{\partial y^4} = 0 \qquad t \geq 0 \qquad y \in [-L, L] \\ u(y, 0) = \sin(y)e^{-(y-10)^2/2} \\ \text{Periodic boundary conditions} \end{cases} \tag{38}$$

We approximate the solution to this PDE using finite differences on the spatial grid

$$y_k = -L + (k-1)\Delta y \qquad k = 1, \ldots, n, \qquad \frac{2L}{n}. \tag{39}$$

This yields the n-dimensional nonlinear dynamical system

$$\frac{dx_j}{dt} = -x_j\frac{x_{j+1} - x_{j-1}}{2\Delta y} - \frac{x_{j-1} - 2x_j + x_{j+1}}{\Delta y^2} - \frac{x_{j-2} - 4x_{j-1} - 6x_j + 4x_{j+1} + x_{j-2}}{\Delta y^4} \qquad j = 1, \ldots, n \tag{40}$$

In Figure 8(a) we show one trajectory of this system for $L = 25$ and $n = 200$ (i.e., a 200-dimensional system) corresponding to the initial condition

$$x_{0k} = u(y_k, 0). \tag{41}$$

In Figure 8(b) we show all Lyapunov exponents for this system. The system is clearly contracting (sum of Lyapunov exponents is negative), and it has three positive Lyapunov exponents

$$\ell_1 = 0.1022, \qquad \ell_2 = 0.0972, \qquad \ell_3 = 0.0595. \tag{42}$$

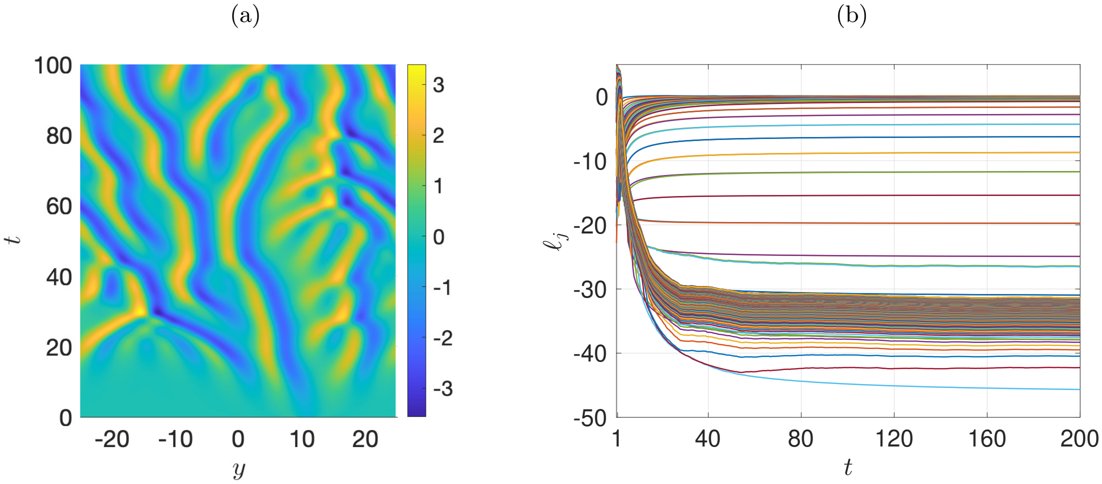

Figure 8: Kuramoto-Sivashinsky equation (38). In (a) we show one trajectory of the dynamical system (40) approximating the solution of (38) for $L = 25$ and $n = 200$. In (b) we show all Lyapunov exponents for this system. The system has a strange attractor with Kaplan–Yorke dimension $D_{KY} = 8.67$, computed using (28).

Hence, the system is approaching a strange attractor with Kaplan–Yorke dimension $D_{KY} = 8.67$, computed using (28).

Lyapunov spectra for Hamiltonian systems. In general, the Lyapunov spectrum of Hamiltonian systems is constrained by time-reversal invariance (reversibility) to have all of the exponents in positive/negative pairs with the same modulus. Moreover, because of time-translation invariance (vector tangent to the flow) and conservation of energy, there are at least two exact zero exponents[8]. The number of zero exponents can be higher than two, and equal to the number of integrals of motion, i.e., the number of conserved quantities. Hamiltonian systems preserve in phase space volume and therefore they do not have an attractor in the dissipative sense. The chaotic dynamics is instead due to invariant energy surfaces being filled by trajectories as time evolves. The dimension of the "attractor" is essentially the dimension of energy level set being filled. For instance, the level sets of the double pendulum Hamiltonian are three-dimensional manifolds (level set of a scalar function in 4 variables). Trajectories of the double span a three dimensional energy manifold, sometimes densely, sometimes on complicated lower-dimensional structures inside the manifold.

Appendix A: Dynamics of infinitesimal n-spheres

In this appendix we show that an infinitesimal sphere is mapped to an infinitesimal ellipsoid along a trajectory, and that the singular values of the Jacobian of the flow coincide with length of the ellipsoid semi-axes.

[8]In practice, if the set of initial perturbation vectors $\boldsymbol{\eta}_j(0)$ does not include a component along $\boldsymbol{f}(\boldsymbol{x}_0)$, the numerical procedure for computing Lyapunov exponents may fail to detect the zero exponent corresponding to time-translation symmetry. To guarantee that this neutral direction is captured, it is important to explicitly include $\boldsymbol{f}(\boldsymbol{x}_0)$ as one of the initial perturbations when computing the full Lyapunov spectrum, i.e., set $\boldsymbol{\eta}_1(0) = \boldsymbol{f}(\boldsymbol{x}_0)$ as initial condition for the linearized equation (17). The reason is that the zero Lyapunov exponent in Hamiltonian systems is associated with a neutral subspace that is at least two-dimensional (flow direction + energy conservation). Unlike dissipative systems, where the neutral direction is unique and alignment along the flow is natural, in Hamiltonian systems perturbations do not preferentially align along the flow direction alone. The alignment tends to be arbitrary within the entire neutral subspace, reflecting the higher-dimensional symmetry and conservation properties of the system.

Theorem 4. Let $\boldsymbol{X}(t, \boldsymbol{x}_0)$ be the flow generated by a smooth vector field $\boldsymbol{f}(\boldsymbol{x})$, and let

$$\boldsymbol{F}(t, \boldsymbol{x}_0) = \frac{\partial \boldsymbol{X}(t, \boldsymbol{x}_0)}{\partial \boldsymbol{x}_0} \tag{43}$$

denote the Jacobian matrix of the flow with respect to the initial condition. Consider the SVD

$$\boldsymbol{F} = \boldsymbol{U}(t)\boldsymbol{\Sigma}\boldsymbol{V}^T, \tag{44}$$

and an infinitesimal sphere centered $S_\epsilon(\boldsymbol{x}_0)$ centered at $\boldsymbol{x}_0$. Then the image of the sphere under $\boldsymbol{F}(t, \boldsymbol{x}_0)$ is an ellipsoid centered at $\boldsymbol{x}_0$ with

1. Principal axes aligned with the columns of $\boldsymbol{U}$ (the left singular vectors of $\boldsymbol{F}$),
2. Semi-axis lengths equal to $r\sigma_1, \ldots, r\sigma_n$.

Proof. Consider the infinitesimal sphere $S_\epsilon(\boldsymbol{x}_0)$ defined in (2) and centered at $\boldsymbol{x}_0$ in the phase space. Under the flow, the image of $S_\epsilon(\boldsymbol{x}_0)$ at time t can be approximated (for infinitesimal ϵ) by the linearized map $\boldsymbol{F}(t, \boldsymbol{x}_0)$. To this end, let $\boldsymbol{\eta}_j(0)$ be such that $\|\boldsymbol{\eta}_j(0)\| \leq \epsilon$. The image of the n-sphere $S_\epsilon(\boldsymbol{x}_0)$ under the flow $\boldsymbol{X}(, \boldsymbol{x}_0)$ is represented by the vectors

$$\boldsymbol{\eta}(t) = \boldsymbol{X}(t, \boldsymbol{x}_0 + \boldsymbol{\eta}_j(0)) - \boldsymbol{X}(t, \boldsymbol{x}_0), \qquad \text{for all} \qquad \boldsymbol{\eta}_j(0) \in S_\epsilon(\boldsymbol{0}). \tag{45}$$

Using Taylor expansions we have

$$\boldsymbol{\eta}_j(t) = \boldsymbol{F}(t, \boldsymbol{x}_0)\boldsymbol{\eta}_j(0) + o(\epsilon),$$

Higher-order terms can be neglected as $\epsilon \to 0$. The matrix $\boldsymbol{F}(t, \boldsymbol{x}_0)$ can be always decomposed onto a symmetric part representing a *pure deformation* and a skew-symmetric part representing a *rotation* as follows

$$\boldsymbol{F}(t, \boldsymbol{x}_0) = \frac{1}{2}\left[\boldsymbol{F}(t, \boldsymbol{x}_0) + \boldsymbol{F}^T(t, \boldsymbol{x}_0)\right] + \frac{1}{2}\left[\boldsymbol{F}(t, \boldsymbol{x}_0) - \boldsymbol{F}^T(t, \boldsymbol{x}_0)\right]. \tag{46}$$

To first-order in ϵ, the infinitesimal sphere $S_\epsilon(\boldsymbol{x}_0)$ is mapped into an infinitesimal ellipsoid by $\boldsymbol{F}(t, \boldsymbol{x}_0)$ (rotation + stretching). This can be also seen by computing the SVD of $\boldsymbol{F} = \boldsymbol{F}(t, \boldsymbol{x}_0)$

$$\boldsymbol{F} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^T,$$

where $\boldsymbol{U}$, $\boldsymbol{V}$ are matrices with orthonormal columns, $\boldsymbol{\Sigma}$ is a diagonal matrix of nonzero singular values s

$$\boldsymbol{\Sigma} = \mathrm{diag}(\sigma_1\sigma_2, \ldots, \sigma_n),$$

For each fixed t and $\boldsymbol{x}_0$, the action of $\boldsymbol{F}(t, \boldsymbol{x}_0)$ on a vector $\boldsymbol{\eta} \in S_\epsilon(\boldsymbol{0})$ can thus be described as

i) Rotation (or reflection) by $\boldsymbol{V}^T$;

ii) Stretch along the coordinate axes by factors σ_i;

iii) Rotation (or reflection) by $\boldsymbol{U}$.

Let us express $\boldsymbol{\eta}$ relative to the orthonormal basis given by the columns of $\boldsymbol{V}$, i.e., $\{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n\}$.

$$\boldsymbol{\eta} = \sum_{i=1}^{n} \alpha_i \boldsymbol{v}_i.$$

We have,

$$\|\boldsymbol{\eta}\|_2^2 = \sum_{i=1}^{n} \alpha_i^2 \leq \epsilon^2. \tag{47}$$

If $\boldsymbol{\eta}$ is on the surface of the infinitesimal sphere then $\|\boldsymbol{\eta}\|_2^2 = \epsilon^2$ Applying $\boldsymbol{F}(t, \boldsymbol{x}_0)$ to $\boldsymbol{\eta}$ yields

$$\boldsymbol{F}\boldsymbol{\eta} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^T\boldsymbol{\eta}$$

with

$$\boldsymbol{\Sigma}\boldsymbol{V}^T\boldsymbol{\eta} = \begin{bmatrix} \sigma_1\alpha_1 \\ \vdots \\ \sigma_n\alpha_n \end{bmatrix}. \tag{48}$$

Thus, each component α_i is stretched by the corresponding singular value σ_i. In particular, if $\boldsymbol{\eta}$ is on the surface of the the sphere $S_\epsilon(\mathbf{0})$ then after rotation by $\boldsymbol{V}^T$, rescaling by $\boldsymbol{\Sigma}$, and rotation by $\boldsymbol{U}$ we obtain a vector on the surface of an ellipsoid. To show this, let us write $\boldsymbol{y} = \boldsymbol{F}\boldsymbol{\eta}$ as

$$\boldsymbol{y} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{\alpha} \qquad \Leftrightarrow \qquad \boldsymbol{\alpha} = \boldsymbol{\Sigma}^{-1}\boldsymbol{U}^T\boldsymbol{y}. \tag{49}$$

Let $\boldsymbol{z} = \boldsymbol{U}^T\boldsymbol{y}$ (canonical coordinates – projection of $\boldsymbol{y}$ onto the orthonormal basis $\boldsymbol{U}$) and let α_j satisfy

$$\sum_{j=1}^{n} \alpha_j^2 = \epsilon^2. \tag{50}$$

Substituting equation (49) into (50) yields

$$\epsilon^2 = \|\boldsymbol{\alpha}\|^2 = (\boldsymbol{\Sigma}^{-1}\boldsymbol{z})^T(\boldsymbol{\Sigma}^{-1}\boldsymbol{z}) = \boldsymbol{z}\boldsymbol{\Sigma}^{-2}\boldsymbol{z}. \tag{51}$$

Therefore, the image of any point $\boldsymbol{\eta}$ on the sphere of radius ϵ satisfies (in normal coordinates $\boldsymbol{z}$)

$$\boldsymbol{z}^T\boldsymbol{\Sigma}^{-2}\boldsymbol{z} = \epsilon^2, \tag{52}$$

which is an ellipsoid with semi-axes of length $\sigma_j\epsilon$. Such semi-axes are directed along the column vectors of $\boldsymbol{U}$ (left singular vectors of $\boldsymbol{F}$).

□

Corollary 1. Let $\sigma_1(t), \ldots, \sigma_n(t)$ be the singular values of the Jacobian matrix (43). Then the Lyapunov exponents can be expressed as

$$\ell_j(\boldsymbol{x}_0) = \lim_{t\to\infty} \frac{1}{t} \log(\sigma_j(t)). \tag{53}$$

The proof of this corollary follows immediately from the previous Theorem and definition (3).