Koopman Analysis of Chua's Circuit with EDMD

Nonlinear Dynamics and Chaos Project

Abstract

The Koopman operator provides a way to study nonlinear dynamical systems through a linear, but generally infinite-dimensional, perspective [2]. Instead of evolving the state directly, the Koopman framework evolves observables of the state. This is appealing in nonlinear dynamics because spectral ideas such as eigenvalues, eigenfunctions, and modes can then be used even when the original system is nonlinear. We first introduce the Koopman operator and show a simple example in which a nonlinear system can be made closed by augmenting the state with one nonlinear observable, then describe Extended Dynamic Mode Decomposition (EDMD), which approximates the Koopman operator from data [5]. We then apply EDMD to experimental data from Chua's circuit measured using an Arduino R4 Minima, analyzing four regimes — fixed point, limit cycle, period-doubled oscillation, and double-scroll chaos — and use Koopman eigenvalues, reconstructions, dictionary sweeps, and eigenfunction coloring to show what the Koopman viewpoint reveals about the route to chaos.

1 Introduction

One of the recurring themes in nonlinear dynamics is that even simple-looking differential equations can produce behavior that is difficult to understand directly in state space. Linear systems are much easier to analyze: eigenvalues describe growth, decay, and oscillation, and solutions can be decomposed into modal pieces. Nonlinear systems do not generally permit such a direct spectral analysis in their original coordinates.

The Koopman operator offers a different viewpoint [2]. Rather than seeking a linear description of the state variables themselves, it studies how *observables* evolve under the dynamics. Although the underlying system may be nonlinear, the Koopman operator acting on observables is linear. This move transfers some of the tools of linear algebra into nonlinear dynamics, provided one is willing to work in a larger space of functions.

This paper has two goals. First, we present the basic Koopman idea in a simple setting where the lifted system closes exactly. Second, we apply a data-driven approximation of Koopman theory to Chua's circuit, a canonical nonlinear electronic system that exhibits a progression from fixed-point behavior to periodic motion, period doubling, and chaos.

2 Koopman Background

2.1 The Koopman Operator

Consider a discrete-time dynamical system

$$x_{n+1} = F(x_n), \tag{1}$$

where $x_n \in \mathbb{R}^d$. Instead of following the state directly, we consider an observable $g : \mathbb{R}^d \to \mathbb{C}$. The Koopman operator $\mathcal{K}$ acts on observables by composition with the dynamics:

$$(\mathcal{K}g)(x) = g(F(x)). \tag{2}$$

This operator is linear, since for observables g_1, g_2 and scalars a, b,

$$\mathcal{K}(ag_1 + bg_2) = a\mathcal{K}g_1 + b\mathcal{K}g_2. \tag{3}$$

The key point is that the nonlinearity of F has not disappeared; rather, it has been moved into the function space on which $\mathcal{K}$ acts. In general this space is infinite-dimensional. The benefit is that if we can identify useful Koopman eigenfunctions ξ_k, satisfying

$$\mathcal{K}\xi_k = \lambda_k \xi_k, \tag{4}$$

then each eigenfunction evolves by a simple scalar multiplication:

$$\xi_k(x_{n+1}) = \lambda_k \xi_k(x_n). \tag{5}$$

This gives a linear evolution law for quantities derived from the nonlinear state.

2.2 A Simple Closed Koopman Representation

To make the Koopman idea concrete, consider the continuous-time nonlinear system

$$\dot{x}_1 = \mu x_1, \qquad \dot{x}_2 = \lambda(x_2 - x_1^2), \tag{6}$$

for which the augmented observable vector

$$\psi(x) = \begin{bmatrix} x_1 \\ x_2 \\ x_1^2 \end{bmatrix} \tag{7}$$

obeys a closed linear system. Indeed,

$$\frac{d}{dt}x_1 = \mu x_1, \tag{8}$$

$$\frac{d}{dt}x_2 = \lambda x_2 - \lambda x_1^2, \tag{9}$$

$$\frac{d}{dt}(x_1^2) = 2x_1\dot{x}_1 = 2\mu x_1^2. \tag{10}$$

Therefore

$$\frac{d}{dt}\begin{bmatrix} x_1 \\ x_2 \\ x_1^2 \end{bmatrix} = \begin{bmatrix} \mu & 0 & 0 \\ 0 & \lambda & -\lambda \\ 0 & 0 & 2\mu \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_1^2 \end{bmatrix}. \tag{11}$$

This is an exact finite-dimensional Koopman-invariant subspace. Although the original system is nonlinear in the state variables because of the x_1^2 term, the lifted observable vector evolves linearly. The example is useful because it shows what one hopes to achieve more generally: by choosing the right observables, nonlinear evolution in state space can become linear evolution in a lifted space.

This example also clarifies an important point. Koopman theory does not say that every nonlinear system has a low-dimensional exact linear representation. Rather, it says that nonlinear systems can be studied through linear evolution of observables, and in favorable cases a finite dictionary of observables may capture a useful approximation.

2.3 Koopman Modes, Eigenvalues, and Interpretation

If a vector of observables $\psi(x)$ can be expanded in terms of Koopman eigenfunctions, then the dynamics take the form of a modal decomposition:

$$\psi(x_n) \approx \sum_k a_k \, \lambda_k^n \, \phi_k, \qquad a_k = \xi_k(x_0), \tag{12}$$

where λ_k are Koopman eigenvalues, ξ_k are eigenfunctions evaluated at the initial state, and ϕ_k are Koopman *modes*. To recover physical state coordinates, one applies a selector matrix C that picks out the original states from the lifted vector, giving $x_n \approx \sum_k a_k \, \lambda_k^n \, (C\phi_k)$; the vectors $C\phi_k$ then record how strongly each measured state participates in the k-th eigenfunction.

The eigenvalues indicate decay, growth, or oscillation: complex conjugate pairs correspond to oscillatory behavior, and for periodic dynamics one expects eigenvalues near the unit circle. This spectral language is especially attractive in nonlinear dynamics because it lets us discuss fixed points, limit cycles, and more complicated attractors using a common linear vocabulary, even when the original equations are nonlinear.

2.4 Brief Description of EDMD

Because the Koopman operator acts on an infinite-dimensional function space, practical computation requires a finite-dimensional approximation. EDMD [5] works in a *dictionary* of N_ψ observables

$$\boldsymbol{\psi}(\mathbf{x}) = \begin{bmatrix} \psi_1(\mathbf{x}) & \psi_2(\mathbf{x}) & \cdots & \psi_{N_\psi}(\mathbf{x}) \end{bmatrix}^\top \in \mathbb{C}^{N_\psi}. \tag{13}$$

The lifted snapshot matrices are

$$\Psi_X = \begin{bmatrix} \boldsymbol{\psi}(\mathbf{x}_1) & \cdots & \boldsymbol{\psi}(\mathbf{x}_M) \end{bmatrix} \in \mathbb{C}^{N_\psi \times M}, \qquad \Psi_Y = \begin{bmatrix} \boldsymbol{\psi}(\mathbf{x}_2) & \cdots & \boldsymbol{\psi}(\mathbf{x}_{M+1}) \end{bmatrix} \in \mathbb{C}^{N_\psi \times M}, \tag{14}$$

where $(\mathbf{x}_k, \mathbf{x}_{k+1})$ are consecutive snapshot pairs from one or more trajectories. EDMD then computes a matrix K such that

$$\Psi_Y \approx K \Psi_X. \tag{15}$$

In our implementation this approximation is obtained by least squares with truncated SVD regularization. The eigenvalues and eigenvectors of K approximate Koopman spectral data: its right eigenvectors, collected as columns of Φ, are the modes ϕ_k in (12), while the rows of $W = \Phi^+$ form the dual basis of left eigenvectors. The k-th Koopman eigenfunction can then be evaluated along any trajectory as

$$\xi_k(x_n) = w_k^\top \psi(x_n), \tag{16}$$

where w_k is the k-th row of W. The quality of the approximation depends heavily on the chosen dictionary: if the observables reflect the structure of the nonlinear system, the lifted linear model can be much more informative.

3 Experimental System: Chua's Circuit

3.1 Why Chua's Circuit

Chua's circuit is a standard example in nonlinear dynamics because it exhibits several qualitatively different behaviors as parameters are varied. In particular, it provides a physically accessible route from simple behavior to chaos. This makes it a natural test case for the Koopman viewpoint:

we can ask how the spectral picture changes as the system passes through a fixed point, a stable oscillation, a period-doubled oscillation, and finally the chaotic double scroll.

At the center of the circuit is its nonlinear element, commonly called the Chua diode. Despite the name, this is not an ordinary semiconductor diode, but an active nonlinear resistor with a piecewise-linear current-voltage characteristic [1, 3]. That nonlinearity is what gives the circuit its rich behavior. The slope changes in the Chua diode produce the switching structure that underlies the double-scroll attractor, and they also motivate the piecewise-linear observables used later in the EDMD dictionaries.

3.2 Data Collection

The experimental data were collected from Chua's circuit using an Arduino R4 Minima. In this implementation, the inductor was realized as a gyrator built from two op amps, and the Chua diode was also implemented using two op amps to produce the required piecewise-linear response. This allowed the full circuit to be constructed from active analog building blocks while preserving the dynamics of the standard Chua design.

To interface the analog circuit with the microcontroller, additional signal-conditioning stages were required. Summing amplifiers and diode clamps were used to shift and limit the measured voltages so that they lay within the readable range of the Arduino ADC, since the native circuit voltages are bipolar while the Arduino input expects a bounded unipolar range. After conditioning, the Arduino sampled three measured signals, denoted here by x, y, and z, corresponding to the state coordinates used throughout the analysis. With easy-to-find component values, Chua's circuit produces very fast dynamics; to keep up, we bypassed the Arduino `analogRead()` overhead and read directly from the R4's ADC registers. For each of the four regimes

- fixed point,
- limit cycle,
- period-doubled oscillation,
- double-scroll chaotic regime,

5000 samples were recorded at a sampling frequency of roughly 30 kHz, then organized into snapshot pairs for EDMD.

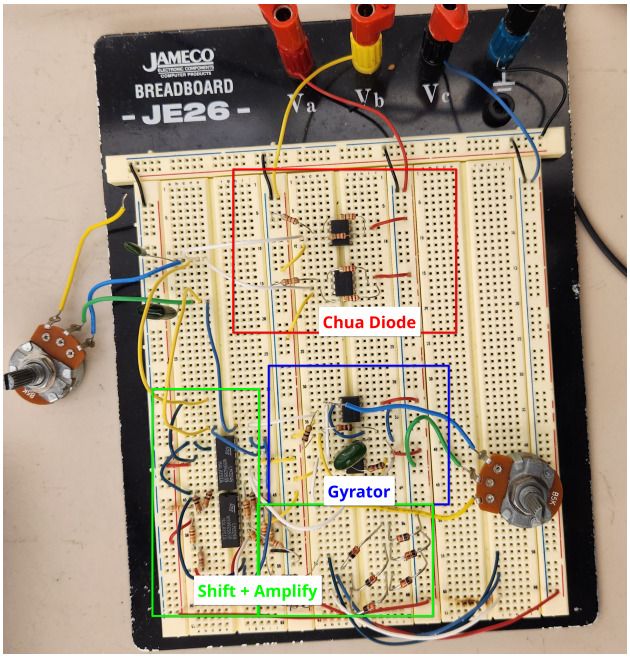

Figure 1: Implementation of Chua Circuit using 4 op-amps + circuit to shift voltage to 0-5V range for Arduino. Potentiometers used as the bifurcation parameter.

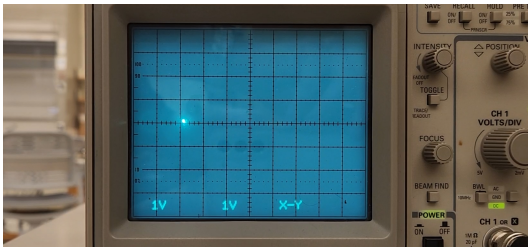

(a) Fixed point

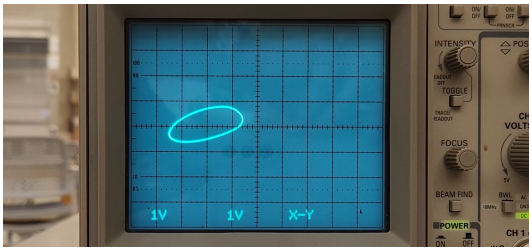

(b) Limit cycle

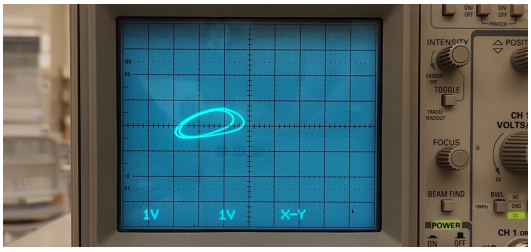

(c) Period-doubled oscillation

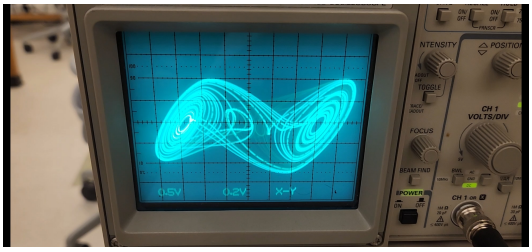

(d) Double-scroll chaos

Figure 2: Oscilloscope captures of the four regimes taken from the analog circuit during data collection, with the bifurcation potentiometer adjusted to walk through the route to chaos. These traces are the experimental view directly off the scope, before digitization; they should be compared with the sampled phase portraits in Figure 3, which show the same regimes after the Arduino ADC.

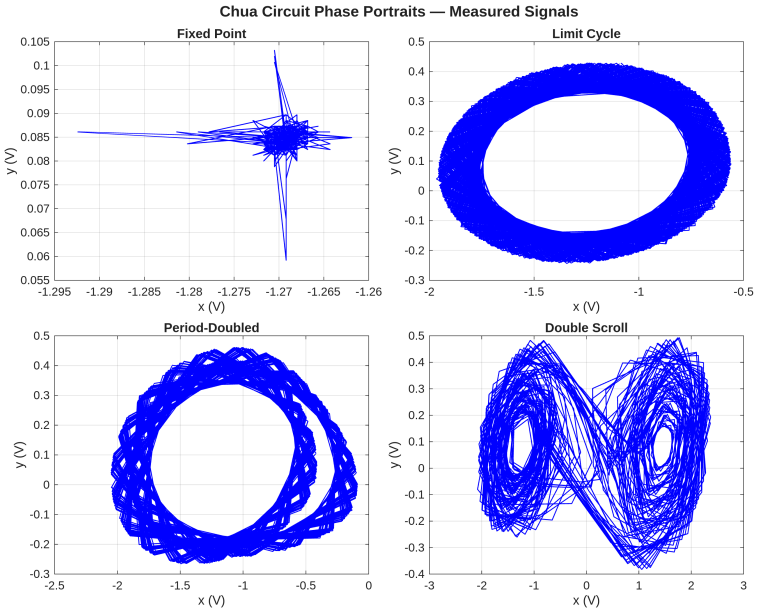

Figure 3: Phase portraits (x versus y) of the four measured regimes. The fixed-point regime (top left) shows scatter around a single point, indicating that the trajectory has settled to an equilibrium with measurement noise. The limit cycle (top right) traces a single closed loop, the geometric signature of sustained periodic oscillation. The period-doubled regime (bottom left) traces a doubled loop — the trajectory must complete two passes before returning to its starting point. The double-scroll regime (bottom right) shows the iconic two-lobe chaotic attractor that distinguishes Chua's circuit. These four panels illustrate the qualitative progression through the route to chaos in raw state space, before any Koopman analysis is applied.

4 Classical Nonlinear Analysis of the Four Regimes

Before applying the Koopman framework, we characterize each regime using two classical measures from nonlinear dynamics: the box-counting fractal dimension and the largest Lyapunov exponent. These provide a quantitative ground truth about whether each regime is truly chaotic and serve as a baseline for interpreting the Koopman spectral picture.

4.1 Box-Counting Fractal Dimension

A strange attractor is characterized geometrically by a non-integer dimension. The box-counting (capacity) dimension D_0 counts how many ε-boxes are needed to cover the attractor:

$$D_0 = -\lim_{\varepsilon \to 0} \frac{\log N(\varepsilon)}{\log \varepsilon}. \tag{17}$$

In practice the limit is estimated by a linear regression of $\log N$ against $\log(1/\varepsilon)$ over a scaling region that avoids both the data-limited small-scale floor and the coarse-scale regime where the power law has not yet established. The three-dimensional state (x, y, z) is normalized axis-by-axis to $[0, 1]^3$ before covering to preserve local geometry.

The results are shown in Figure 4 and summarized in Table 1. The fixed-point regime yields a near-zero dimension ($D_0 = 0.181$), confirming that the trajectory has collapsed to a compact neighborhood of an equilibrium; the non-integer value reflects measurement noise smearing the

point into a small cluster. The limit-cycle and period-doubled regimes give dimensions close to 1 ($D_0 = 1.101$ and 1.052), consistent with a one-dimensional curve in state space. The double-scroll regime yields $D_0 = 1.519$, clearly above 1 and providing direct geometric evidence of a strange attractor. The underestimate of the theoretical value $D_F \approx 2.05$–2.3 [1] is a finite-sample artifact — reliable box-counting requires $\sim 10^{D_F}$ points to resolve the scaling region — but the qualitative ordering across regimes is unambiguous.

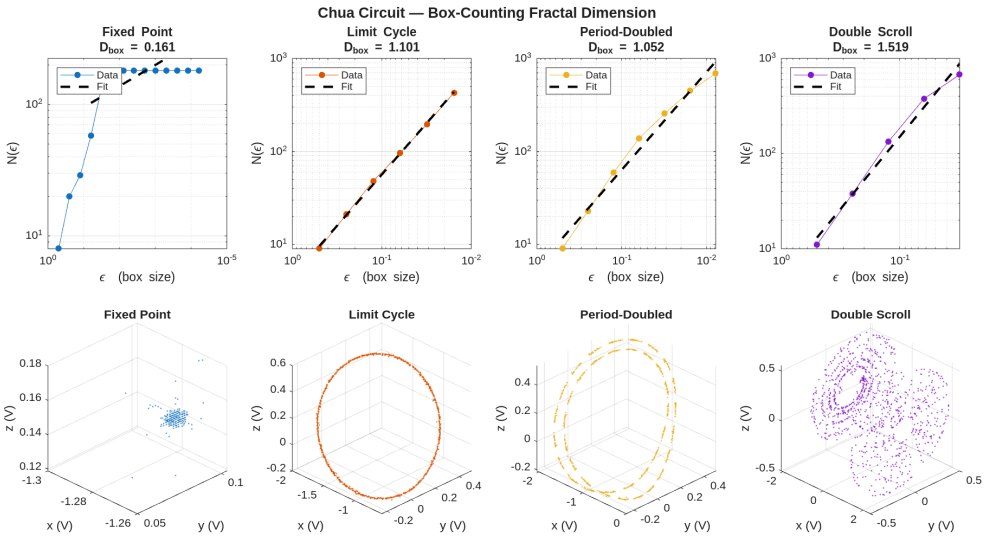

Figure 4: Box-counting fractal dimension of the four Chua regimes. Top row: $\log N(\varepsilon)$ versus $\log \varepsilon$ with linear fit (dashed); the slope gives D_0. Bottom row: 3-D phase portraits. The dimension increases from near zero (fixed point) through approximately 1 (periodic regimes) to 1.519 (double scroll), confirming the fractal geometry of the strange attractor.

4.2 Largest Lyapunov Exponent

Lyapunov exponents measure the mean exponential rate of separation of nearby trajectories; a positive largest exponent λ_1 is the defining hallmark of chaos. We use the algorithm of Rosenstein et al. [4], which tracks the mean log-divergence of nearest-neighbor pairs. For each reference point i, the nearest spatial neighbor j is found outside a Theiler exclusion window W (set to twice the first zero-crossing lag of the autocorrelation of the x channel) to suppress temporally correlated neighbors. The average over all pairs,

$$\left\langle \ln \frac{d_k}{d_0} \right\rangle \approx \lambda_1 k + \text{const}, \tag{18}$$

grows linearly in the initial phase at rate λ_1.

Results are shown in Figure 5 and collected in Table 1. The fixed-point regime is the only one with a clearly negative slope ($\lambda_1 = -0.0042$/sample), reflecting contraction toward the equilibrium. The limit-cycle and period-doubled regimes yield small positive values ($+0.0172$ and $+0.0182$/sample); these arise because the Rosenstein estimator fits the initial transient growth before the divergence curve saturates as neighbors leave the compact attractor, and measurement noise introduces a slight upward bias on periodic orbits. The qualitatively decisive result is the

double-scroll regime, which gives $\lambda_1 = +0.0726$/sample — more than four times the periodic values and with a clean linear growth region — confirming sensitive dependence on initial conditions.

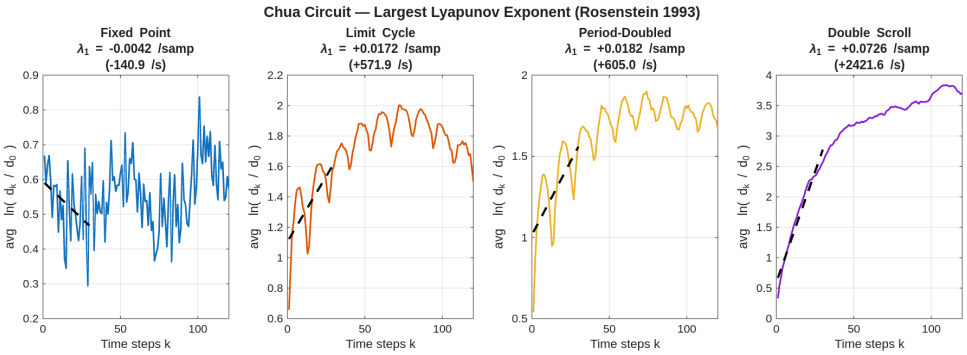

Figure 5: Largest Lyapunov exponent via the Rosenstein method for all four Chua regimes. Solid curves show the mean log-divergence $\langle \ln(d_k/d_0) \rangle$; dashed lines are the linear fits to the initial growth region. The fixed-point regime is the only one with a clearly negative slope. The double-scroll regime shows a clean linear initial growth at $\lambda_1 = +0.0726$/sample, more than four times the small positive bias seen in the periodic regimes.

Table 1: Classical nonlinear measures for the four Chua regimes. D_0: box-counting fractal dimension. λ_1: largest Lyapunov exponent (Rosenstein method, per sample).

Regime	D_0	λ_1 (/sample)
Fixed Point	0.181	-0.0042
Limit Cycle	1.101	$+0.0172$
Period-Doubled	1.052	$+0.0182$
Double Scroll	1.519	$+0.0726$

5 EDMD Analysis of the Four Regimes

5.1 Choice of Dictionary

Because EDMD depends strongly on the observable dictionary, part of the analysis focused on comparing different families of observables. The implementation includes five families: linear (the raw state), polynomial (monomials of total degree at most d), radial basis function (Gaussian bumps with k-means centers), Chua-matched piecewise-linear (the seven observables built from the kink terms $|x_1 \pm 1|$ and $x_1|x_1 \pm 1|$), and a "full" family that combines polynomial, RBF, and piecewise-linear terms. The Chua-specific observables are motivated by the piecewise-linear nonlinearity in the circuit, so they are expected to capture the dynamics more efficiently than a completely generic basis. This follows the broader EDMD philosophy that approximation quality depends not only on the data, but on whether the chosen observables are rich enough to represent the important Koopman-invariant structures [5].

This comparison is important conceptually. Koopman theory is not a black-box replacement for nonlinear dynamics. Rather, it is a framework in which physical insight enters through the observables one chooses to evolve.

5.2 Dictionary Sweep: Accuracy Versus Cost

To quantify the effect of dictionary choice, we performed a sweep over the five families above and over dictionary sizes, measuring the relative one-step error in lifted space together with computation time. For each regime the data were min-max normalized to $[-1, 1]$ before lifting, so that high-degree monomials remain $O(1)$ and the Gram matrix stays well conditioned.

The results are shown in Figure 6. The expected trend, that richer dictionaries reduce one-step prediction error, is visible across all four regimes: the polynomial, RBF, and full families all improve substantially as N_ψ grows, with the chaotic double-scroll regime showing the largest absolute improvement because it has the most complex spectral content to capture. The linear dictionary saturates at relatively high error in every regime, confirming that purely linear observables cannot describe the Chua dynamics. Note that adding more polynomial terms does not help describe the singular fixed point.

The right column shows that this accuracy comes at a cost. The dominant scaling is the cubic complexity of the EDMD solve in N_ψ, so timing rises monotonically for the polynomial and RBF families. The full dictionary timing is non-monotonic because different "full" configurations include or exclude the RBF block (and the associated k-means call); when the configurations are sorted by N_ψ for plotting, RBF-free configurations sit between RBF-containing ones of similar size, producing visible dips. The conceptual point is that dictionary choice is genuinely a tradeoff: enrichment improves accuracy but does so at increasing cost and with the risk of numerical conditioning issues at high degree if normalization is not used.

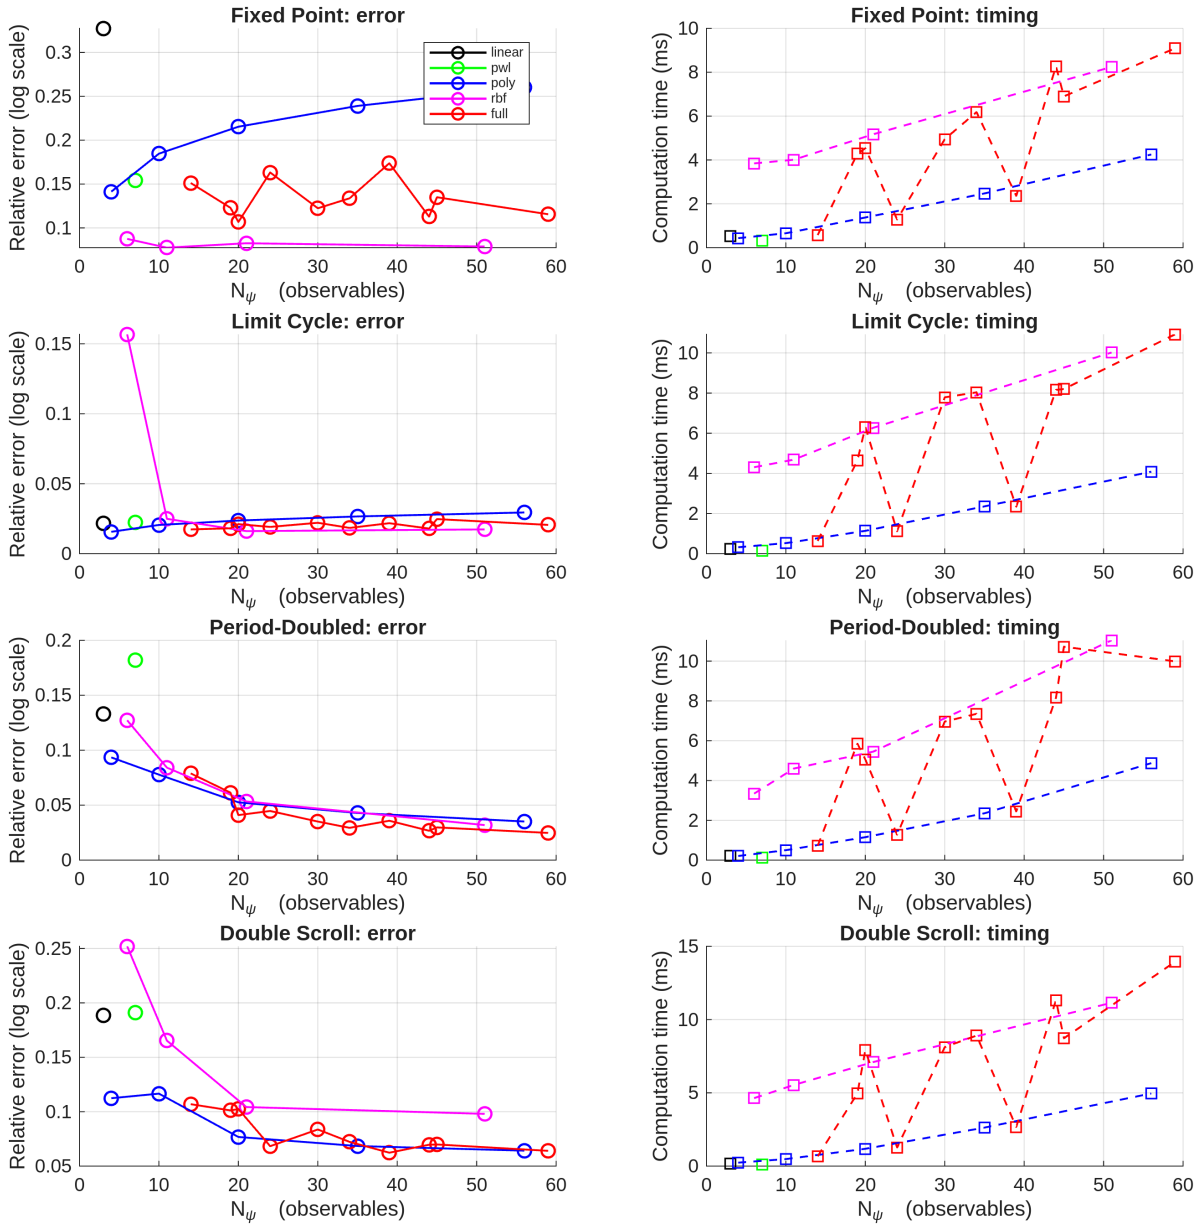

Figure 6: Dictionary sweep across the four regimes. Each row corresponds to a regime, with the relative one-step error in lifted space (left, log scale) and the EDMD wall-clock time in milliseconds (right) plotted against the number of observables N_ψ. Marker color encodes the dictionary family: black (linear), green (PWL), blue (poly), magenta (RBF), red (full). Solid and dashed lines connect points within the same family. Error decreases substantially with richer dictionaries, while computation time rises with N_ψ as expected from the cubic-in-N_ψ EDMD solve.

5.3 Koopman Eigenvalues Across Regimes

For each regime, EDMD produces a finite-dimensional Koopman matrix whose eigenvalues approximate the spectral behavior of the system. These eigenvalues are especially informative when compared across regimes, because the qualitative differences between fixed-point, periodic, period-doubled, and chaotic dynamics translate into qualitative differences in the spectrum.

A common dictionary — monomials up to total degree three, giving $N_\psi = 20$ observables — was used for all four regimes so that the panels in Figure 7 are directly comparable. To make conjugate pairs visually identifiable, the plotting routine identifies each λ together with its conjugate $\bar{\lambda}$ and assigns them a single color from a qualitative palette; real eigenvalues are drawn in black. The unit circle is drawn for reference, since for a periodic attractor the physical Koopman eigenvalues lie exactly on it (any deviation inward is an EDMD truncation artifact).

The expected pattern, regime by regime, is the following. For a stable fixed point, all dynamics decay exponentially without oscillation, so the spectrum should consist of real eigenvalues clustered inside the unit circle, with no nontrivial complex pairs. For a stable limit cycle of period T samples, the true Koopman eigenvalues are the T-th roots of unity $e^{i2\pi k/T}$, all on the unit circle, with conjugate pairs at the fundamental frequency and its harmonics. For a period-doubled orbit, the period in samples doubles, so additional eigenvalues appear at the new $2T$-th roots of unity; the spectrum splits, and the half-frequency subharmonic shows up as a conjugate pair on the unit circle that was not present in the limit-cycle spectrum. The most direct fingerprint of period doubling is therefore the appearance of new conjugate pairs at angles between those of the original limit cycle. For chaos, no finite spectrum on the unit circle can capture the dynamics exactly; the EDMD spectrum spreads across a broader region of the complex plane, with magnitudes both on and inside the unit circle.

These predictions are borne out in Figure 7. The fixed-point panel shows only real eigenvalues, all black, with one near $\lambda = 1$ (the constant observable, which is always a Koopman eigenfunction with eigenvalue 1) and the rest clustered well inside the unit circle, consistent with non-oscillatory decay to equilibrium. The limit-cycle panel shows the constant mode at $\lambda = 1$, a real mode near $\lambda = -1$, and conjugate pairs near $\pm i$, all very close to the unit circle. These are precisely the fourth roots of unity, indicating that the limit cycle has approximately four samples per period. The period-doubled panel retains the structure of the limit cycle but adds several new conjugate pairs at intermediate angles around the unit circle (visible as additional colored pairs near $e^{\pm i\pi/4}$ and $e^{\pm i3\pi/4}$), which is the spectral signature of the doubled orbit: the new pairs sit at the subharmonic frequency and its odd harmonics. The double-scroll panel shows the broadest distribution: many conjugate pairs are scattered both on and inside the unit circle, consistent with chaotic dynamics that are not captured by any finite set of unit-circle frequencies.

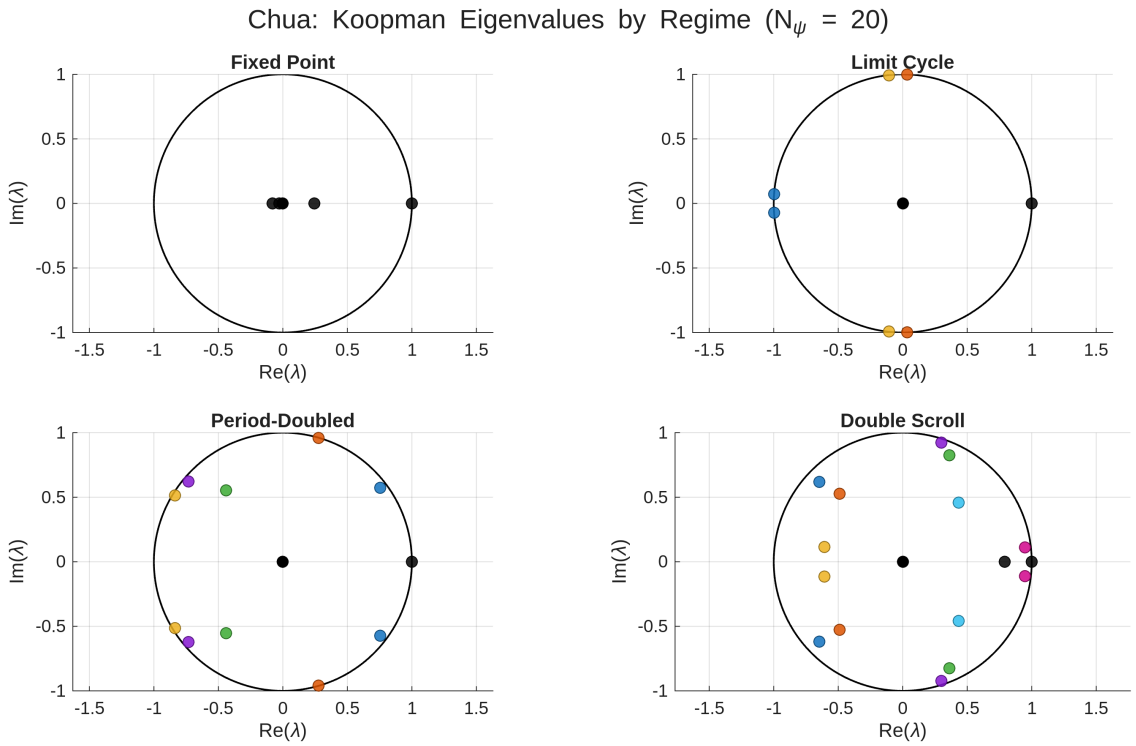

Figure 7: Koopman eigenvalues for the four Chua regimes, computed from EDMD with a degree-3 polynomial dictionary ($N_\psi = 20$). Each conjugate pair is plotted in a single color; real eigenvalues are shown in black. The fixed-point spectrum is purely real and inside the unit circle. The limit-cycle spectrum sits near the fourth roots of unity. The period-doubled spectrum adds new conjugate pairs at intermediate angles (the subharmonic and its odd harmonics). The double-scroll spectrum is broadly distributed both on and inside the unit circle, reflecting the absence of a finite set of dominant frequencies in chaos.

5.4 Period-Doubled Reconstruction from Koopman Modes

Beyond reading off eigenvalues, the modal expansion (12) can be used to reconstruct the measured state directly. With the EDMD matrix K, its right eigenvectors ϕ_k, and the amplitudes $a_k = (\Phi^+\psi(x_0))_k$ in hand, projecting back through the selector C gives $x_n \approx \sum_k a_k\,\lambda_k^n\,(C\phi_k)$.

The period-doubled regime is particularly instructive because its spectral picture reflects the bifurcation directly. The dominant oscillatory content is no longer just the original fundamental frequency f_0 inherited from the limit cycle; instead, the leading conjugate pair sits near $f_0/2$, the subharmonic that is the hallmark of period doubling. The script ranks bounded modes by their contribution to the x-state energy, taking conjugate pairs together so that real-valued reconstructions are preserved, and selects the top ten such modes for comparison with the measured state.

One subtlety is that EDMD systematically pulls the eigenvalues of a sustained periodic orbit slightly inside the unit circle: a finite dictionary cannot represent the global Koopman operator exactly, so the least-squares fit absorbs some genuine oscillation as a small numerical decay. Over thousands of samples this artifact accumulates and the modal reconstruction $a_k\lambda_k^n \to 0$ even though the physical signal does not. To remove this, the reconstruction projects each oscillatory eigenvalue back to the unit circle, $\lambda_k \mapsto \lambda_k/|\lambda_k|$, while leaving its identified frequency and phase unchanged. Real (DC) eigenvalues are not modified. This gives a cleaner physical interpretation of what the spectrum is asserting — that these modes oscillate at fixed frequencies forever — without altering

any of the dynamical content that EDMD actually identified.

The clearest geometric test of the reconstruction is the phase plane. A period-doubled orbit traces two nearby loops in state space: the trajectory must complete two passes before closing, so the attractor topologically consists of two interleaved loops rather than one. Figure 8 compares the measured phase portrait with the top-ten Koopman reconstruction in three projections (x-y, x-z, y-z). To make the doubled-loop structure visually unambiguous, each sample is colored by the sign of the $f_0/2$ Koopman eigenfunction: blue when $\mathrm{Re}(\xi_{f_0/2}) \geq 0$ (Loop A), red otherwise (Loop B). This is a satisfying internal consistency check, since the $f_0/2$ eigenfunction is by definition the observable that changes sign exactly once per doubled period — its own zero crossings are what define the two loops — so using it as a loop label closes the loop between the bifurcation's spectral signature (a new λ at frequency $f_0/2$) and its geometric signature (a doubled orbit). The bottom row shows that the reconstruction from ten Koopman modes recovers both loops in all three projections, confirming that the spectral content identified by EDMD is sufficient to rebuild the doubled-loop topology, not just the dominant frequency.

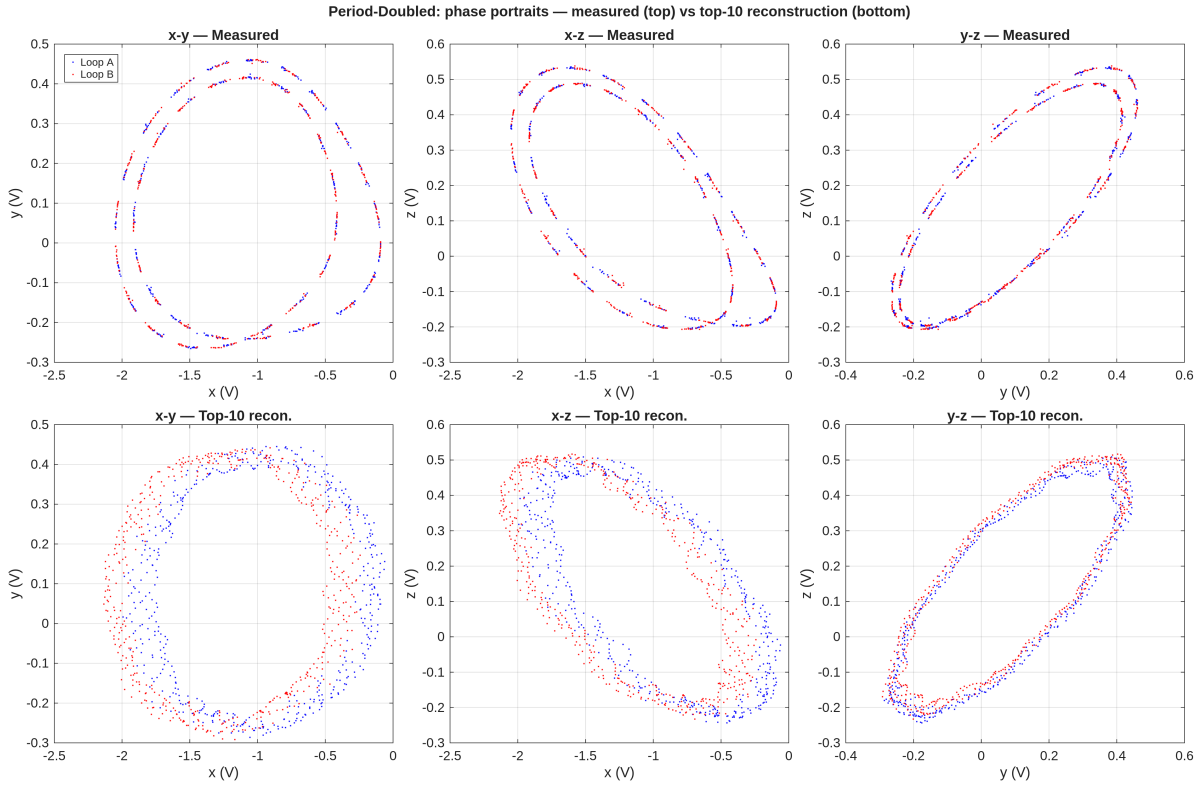

Figure 8: Phase portraits of the period-doubled regime in three projections (x-y, x-z, y-z). *Top row:* measured trajectory. *Bottom row:* top-ten Koopman mode reconstruction with each oscillatory eigenvalue projected to the unit circle. Samples are colored by the sign of the $f_0/2$ Koopman eigenfunction — blue for Loop A ($\mathrm{Re}(\xi_{f_0/2}) \geq 0$) and red for Loop B — which makes the two interleaved loops of the doubled orbit clearly visible. The reconstruction recovers the same two-loop topology as the measurement, demonstrating that the Koopman modes selected by x-energy ranking are sufficient to reproduce both the dominant subharmonic and the higher harmonics that shape each loop.

5.5 Double Scroll and Eigenfunction Coloring

The chaotic double-scroll regime is the most visually striking part of Chua's circuit, but it is also where the Koopman viewpoint becomes especially insightful. In raw state space the trajectory wanders unpredictably between two lobes, and no individual voltage—x, y, or z—cleanly distinguishes them. The two lobes overlap in every linear projection: a threshold on x alone, for example, will incorrectly assign points near the center of the attractor because both lobes pass through that region. This is a fundamental limitation of working directly with state coordinates in a chaotic system.

The Koopman framework offers a different approach. Rather than looking for a separator in the original state space, we look for an eigenfunction of the Koopman operator, a nonlinear observable that evolves by simple multiplication at each timestep. By definition $\xi_k(x_{n+1}) = \lambda_k \, \xi_k(x_n)$, so that $\xi_k(x_n) = \lambda_k^n \, \xi_k(x_0)$: the state x_n itself evolves chaotically, but the scalar $\xi_k(x_n)$ executes a rigid rotation in the complex plane at angular frequency $\arg(\lambda_k)$ per timestep. This is precisely the Koopman linearity guarantee: nonlinear dynamics in state space become linear dynamics in observable space.

We compute ξ_k along the measured trajectory using (16) and select for visualization the bounded oscillatory mode with the highest variance in $\mathrm{Re}(\xi_k)$, which is the mode that most strongly discriminates different regions of the attractor. For this regime a degree-3 polynomial dictionary augmented with the four piecewise-linear Chua kink terms was used, since a small dictionary does not provide enough expressive power to resolve the eigenfunction structure of the double scroll.

Interpreting the four panels. The visualization in Figure 9 presents four complementary views of the same eigenfunction.

The top two panels color every attractor point by $\arg(\xi_k)$ using a cyclic (HSV) colormap, shown in the x-y and x-z projections respectively. Because ξ_k advances by a fixed angle each step, points visited at similar phases of the rotation receive similar colors. If the eigenfunction genuinely reflects the attractor's geometry, the color should sweep smoothly around each lobe rather than appearing randomly speckled. Smooth, coherent color bands confirm that ξ_k is acting as an intrinsic phase coordinate on the attractor—not a noise artifact—and that the chosen mode captures persistent oscillatory structure even within the chaos.

The bottom-left panel reduces the complex eigenfunction to a binary label: points where $\mathrm{Re}(\xi_k) \geq 0$ are shown in blue, and those where $\mathrm{Re}(\xi_k) < 0$ in red. This is the most direct test of the eigenfunction's utility as a lobe separator. If it succeeds, blue and red should cleanly partition the two scrolls in the full three-dimensional state space, achieving in a single learned observable what no voltage threshold can accomplish. The key is that ξ_k is a nonlinear function of all three state variables: the EDMD procedure has found, from data alone, the combination of polynomial and piecewise-linear terms that most cleanly partitions the attractor.

The bottom-right panel shows $\mathrm{Re}(\xi_k)$ and $\mathrm{Im}(\xi_k)$ as time series. Despite the irregular switching visible in the raw voltages, the eigenfunction trace should show approximately sinusoidal oscillation whose real part changes sign each time the trajectory transitions between lobes. This is the time-domain signature of the linear evolution $\xi_k(x_{n+1}) = \lambda_k \xi_k(x_n)$: the chaotic switching in physical coordinates maps to a regular phase rotation in the eigenfunction coordinate.

Why this demonstrates the advantage of Koopman. The central challenge for any analysis method applied to chaos is that the dynamics are sensitive to initial conditions and superficially irregular. State-space methods that rely on linear combinations of voltages are limited because nonlinear geometric features—such as lobe membership—have no simple linear signature. The

Koopman approach circumvents this by explicitly searching the space of nonlinear observables for one that evolves linearly. The eigenfunction coloring demonstrates this advantage concretely: it transforms the question "which lobe is the trajectory in?" from an intractable problem in state space into a trivial threshold on a single learned coordinate. Moreover, the eigenvalue λ_k associated with that coordinate tells us the frequency at which the trajectory winds around the attractor, providing quantitative dynamical information that is completely invisible in the raw time series.

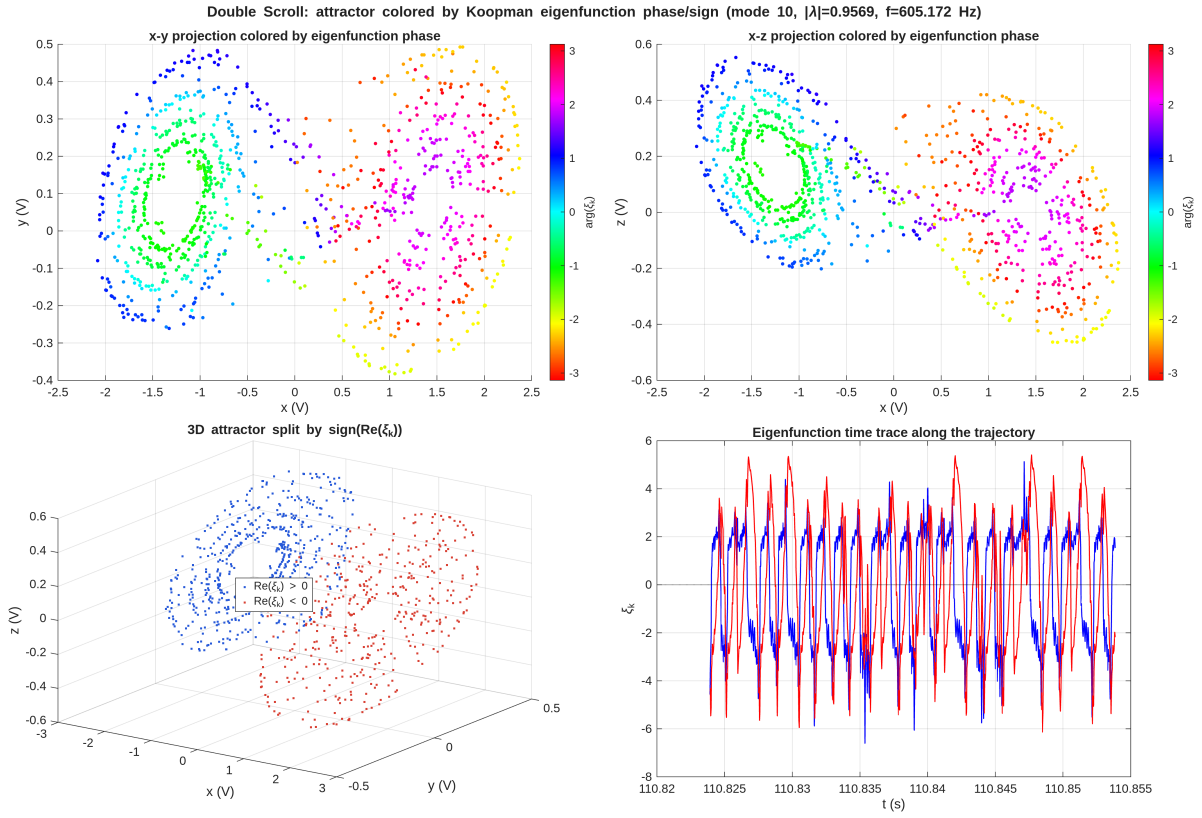

Figure 9: Double-scroll attractor colored by a Koopman eigenfunction ξ_k with $|\lambda_k| = 0.957$ and identified frequency $f_k \approx 605$ Hz. *Top row:* attractor points colored by $\arg(\xi_k)$ in the x-y and x-z projections; the smooth, coherent color bands sweeping around each lobe confirm that ξ_k is acting as a genuine phase coordinate on the attractor rather than a noise-driven artifact. *Bottom left:* three-dimensional attractor split by $\mathrm{sign}(\mathrm{Re}(\xi_k))$; blue and red cleanly separate the two scrolls, a partition that no single voltage threshold could achieve in raw state space. *Bottom right:* time trace of $\mathrm{Re}(\xi_k)$ (blue) and $\mathrm{Im}(\xi_k)$ (red); despite the irregular chaotic switching in the raw voltages, the eigenfunction itself executes a near-sinusoidal complex rotation, the time-domain signature of the linear evolution $\xi_k(x_{n+1}) = \lambda_k \xi_k(x_n)$.

6 Discussion

Several points emerge from this project. First, the Koopman framework provides a natural spectral language for nonlinear dynamics. Even though the systems under study are nonlinear, their behavior can be described through eigenvalues, eigenfunctions, and modes once appropriate observables are chosen.

Second, dictionary choice is central. The dictionary sweep shows that better observables can

substantially improve the EDMD approximation, but not for free: richer dictionaries increase computation time and can worsen numerical conditioning. This is both a practical issue and a conceptual one, since it emphasizes that Koopman analysis depends on choosing observables that reflect the underlying physics.

Third, the Chua circuit data illustrate a meaningful progression through the route to chaos. The fixed-point regime is spectrally simple, the limit cycle exhibits conjugate-pair structure associated with sustained oscillation, the period-doubled regime reveals subharmonic content, and the double-scroll regime shows that Koopman eigenfunctions can still uncover coherent structure even in chaos. This is all still done without any knowledge of the underlying governing dynamical equations.

Fourth, the classical and Koopman analyses are complementary rather than competing. The box-counting dimension ($D_0 = 1.519$) confirms that the double-scroll attractor is geometrically fractal, and the positive Lyapunov exponent ($\lambda_1 = +0.0726$/sample) certifies chaos. These classical tools answer the diagnostic question "is it chaotic?" definitively. The Koopman framework then adds a different kind of understanding: even within the chaotic attractor, there exist nonlinear observables that evolve approximately linearly and encode the attractor's geometry. The eigenfunction coloring of the double scroll, for instance, achieves a lobe partition that no voltage threshold can replicate — a result that is invisible to Lyapunov or fractal-dimension analysis but falls naturally out of the spectral structure of the Koopman operator. The two approaches together provide a more complete picture than either one alone.

7 Conclusion

This project used the Koopman operator as a bridge between linear spectral ideas and nonlinear dynamical behavior. A simple background example showed how augmenting the observable space can produce a closed linear description when the right nonlinear term is included. EDMD then provided a practical way to approximate Koopman spectral quantities from experimental data [2, 5].

Applied to Chua's circuit, this framework gave a unified way to study four regimes measured with an Arduino R4 Minima: fixed point, limit cycle, period doubling, and double-scroll chaos. The most compelling result is that the Koopman viewpoint does not merely fit the data; it organizes the behavior. In the periodic regimes it yields interpretable oscillatory modes, and in the chaotic regime it produces eigenfunctions that reveal hidden coordinates on the attractor. This makes the Koopman approach especially valuable for a nonlinear dynamics setting, where the central challenge is often not only to simulate nonlinear motion, but to understand its structure.

References

[1] Leon O. Chua, Mitsuhiro Komuro, and Takashi Matsumoto. The double scroll family. *IEEE Transactions on Circuits and Systems*, 33(11):1072–1118, 1986.

[2] B. O. Koopman. Hamiltonian systems and transformation in hilbert space. *Proceedings of the National Academy of Sciences*, 17(5):315–318, 1931.

[3] Takashi Matsumoto. A chaotic attractor from chua's circuit. *IEEE Transactions on Circuits and Systems*, CAS-31(12):1055–1058, 1984.

[4] Michael T. Rosenstein, James J. Collins, and Carlo J. De Luca. A practical method for calculating largest Lyapunov exponents from small data sets. *Physica D: Nonlinear Phenomena*, 65 (1–2):117–134, 1993.

[5] Matthew O. Williams, Ioannis G. Kevrekidis, and Clarence W. Rowley. A data-driven approximation of the koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*, 25(6):1307–1346, 2015.