Reconstructing Chaotic Dynamics with Delay Embedding and Gaussian Process Regression: Applications to Ecological Biomass Data

Joey Park[1], Advisor: Prof. Sarah Day[2]
[1] Thomas Jefferson High School for Science and Technology, Fairfax, VA, USA
[2] Department of Mathematics, William & Mary, Williamsburg, VA, USA

*Corresponding author: joyful.joeypark@gmail.com

Keywords: state-space reconstruction, delay embedding, Gaussian process regression, false nearest neighbors, Wasserstein distance, ecology.

Abstract. We combine delay embedding and Gaussian Process Regression (GPR) to model chaotic dynamics and then apply the method to ecological time series. As a first step, we tested the approach on synthetic systems. For the logistic map, we ran GPR with $Q = 2, 3$ using only five data points (evenly spaced and randomly chosen) under different noise levels. For the Hénon map, we tried multiple measurement functions, identified three that kept the structure well, and used false nearest neighbors (FNN) to find the minimal embedding dimension Q. With the chosen functions and added noise, we trained GPR models and checked their accuracy using Wasserstein distance. Building on these trials, we then applied the same method to real ecological data, where we used delay embedding and GPR to test predictive power and robustness.

1 Introduction & Background

Chaotic systems show up in ecology, climate, and many other natural processes, where even small changes in starting conditions make prediction difficult. Understanding these patterns is important for ecology, since nonlinear interactions can cause sudden shifts and reliable early-warning signs are hard to find. A classical tool for working with this kind of data is Takens' embedding theorem, which shows how to rebuild the underlying state space from just one time series using delay embedding. Once the state space is rebuilt, we can use machine learning methods to model and forecast the system.

Gaussian Process Regression (GPR) is a flexible, nonparametric method that can learn nonlinear functions from small and noisy datasets. When combined with delay embedding, GPR can be trained on the embedded time series to approximate the system and make short-term forecasts. Recent work, such as the SSR–GPR framework, has shown that this approach can capture important features of chaotic maps, confirming the value of linking reconstruction with GPR.

We build on this work and test the framework in more practical settings. First, we study what happens with very limited and noisy data. For the logistic map, we ran GPR with embedding dimensions $Q = 2, 3$ using only five points (evenly spaced or randomly chosen) at different noise levels. For the Hénon map, we tried a range of measurement functions, picked the ones that best kept the structure, and used false nearest neighbors (FNN) to find the minimal embedding dimension. We then added noise, trained GPR models, and measured performance with Wasserstein distance against baseline data. These steps pushed the SSR–GPR method further than earlier synthetic demonstrations.

Finally, we applied the same approach to real ecological data: the Salvinia–weevil time series. This let us see if delay embedding plus GPR can provide useful predictions in an applied setting. In this way, we connect theory on chaotic dynamics with practical ecological forecasting.

The rest of the paper is organized as follows. Section 2 covers the logistic map experiments. Section 3 extends the method to the Hénon map, focusing on measurement functions and delay embedding. Section 4 applies the method to ecological data. Section 5 presents results and discussion, and Section 6 concludes with implications and directions for future work.

2 Logistic Map

The logistic map is a classic starting point for studying complex dynamics, showing how a simple one-dimensional quadratic recurrence can create many different behaviors. In this section, we look closely at its dynamics and test how well Gaussian Process Regression (GPR) can rebuild and approximate such nonlinear systems. We start by examining the bifurcation cascade and the onset of chaos as the growth parameter r increases, using cobweb plots and bifurcation diagrams for visualization. Next, we apply delay embedding to chaotic trajectories and show that GPR can recover the underlying one-step map even when we only have limited time series data. To see the effect of embedding more clearly, we also train GPR models directly on function samples across r without reconstruction, showing that the results reproduce the full bifurcation structure with high accuracy. Finally, we push the method further by testing it with very little data and added noise, showing how the GP noise hyperparameter α helps balance exact interpolation and smoothing. Together, these experiments establish the logistic map as a solid benchmark for testing state space reconstruction and nonparametric learning under both ideal and challenging settings.

2.1 Exploration of Logistic Map Dynamics

The logistic map

$$x_{t+1} = rx_t(1 - x_t)$$

is a classical example of how simple rules can produce complex dynamics. For certain values of r, trajectories converge to a fixed point or a short cycle, while larger r values generate aperiodic and chaotic behavior.

Figure 1 illustrates the chaotic case with $r = 3.8$. On the left, a time series starting from $x_0 = 0.5$ is iterated for 50 steps, showing the irregular fluctuations characteristic of chaos. On the right, a cobweb plot visualizes the same iterations against $f(x) = rx(1 - x)$ and the line $y = x$, making clear how successive points jump unpredictably rather than settling into a stable orbit.

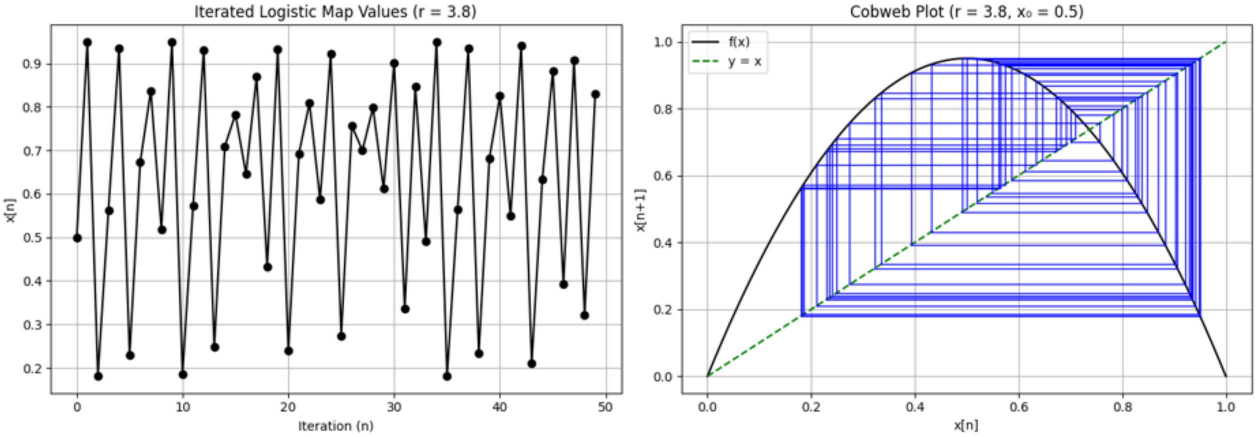

Figure 1: **Exploring the Logistic Map.** Iterations of the logistic map with $r = 3.8$ and $x_0 = 0.5$. Left: 50 iterations plotted as a time series, revealing irregular dynamics. Right: Cobweb plot of the same iterations, showing step-by-step progression along the parabola $f(x)$ and the diagonal $y = x$.

2.2 Delay Embedding + GPR Reconstruction at $r = 3.8$

To go beyond direct iteration, we used delay embedding to rebuild the state space from a univariate time series generated at $r = 3.8$, where the system shows strong chaos. Delay coordinates reveal the shape of the attractor without needing to know the original equation.

We then trained Gaussian Process Regression (GPR) models on the embedded points. Even with only a small number of samples, the GP was able to approximate the one-step dynamics and capture the curve of the logistic parabola. This shows that delay embedding with GPR can recover chaotic behavior from sparse data.

To make comparisons fair across different parameter values, we kept the sampling rule consistent: at each r, we used $N = 10$, the first ten points. This let us systematically check how reconstruction quality changed from periodic to chaotic cases.

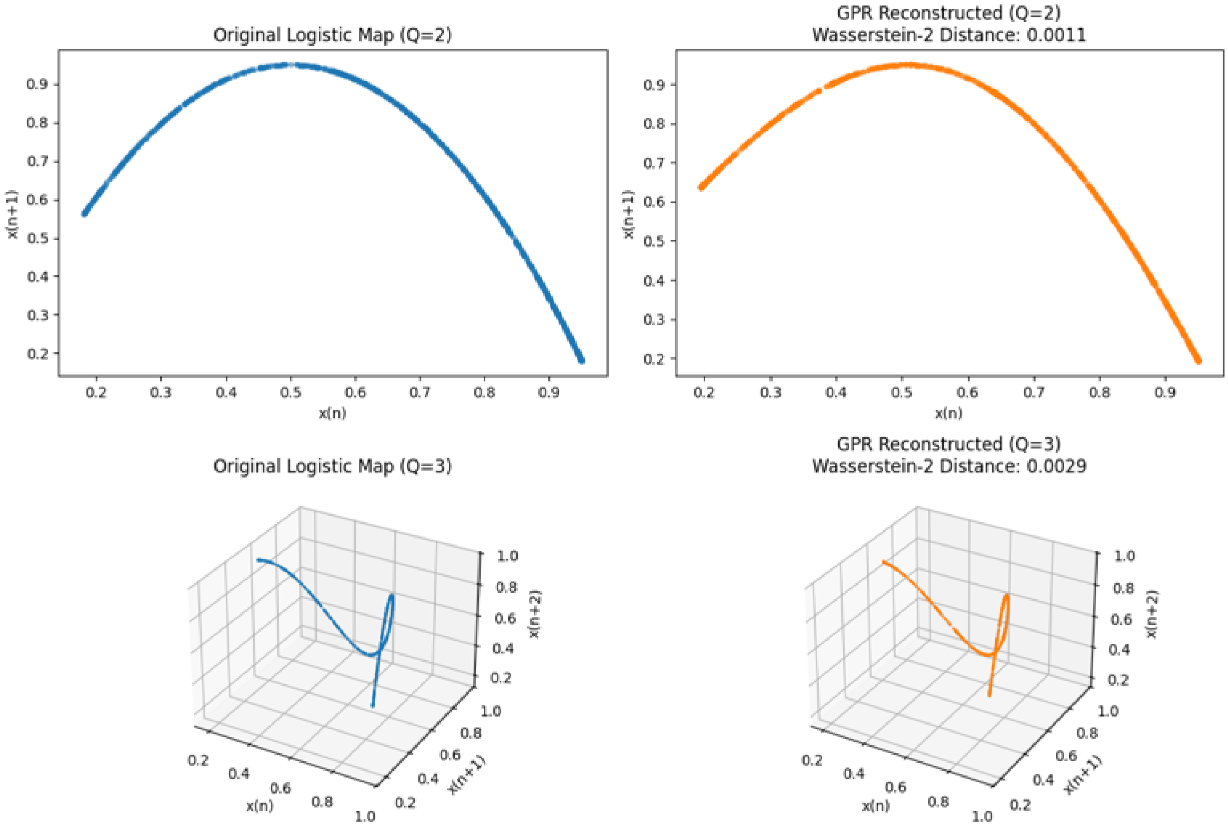
Original Logistic Map (Q=2)

GPR Reconstructed (Q=2)
Wasserstein-2 Distance: 0.0011

Original Logistic Map (Q=3)

GPR Reconstructed (Q=3)
Wasserstein-2 Distance: 0.0029

Figure 2: **Delay Embedding with GPR ($r = 3.8$).** Delay embedding rebuilds the state space from a univariate time series, and Gaussian Process Regression (GPR) approximates the one-step map under chaotic dynamics.

2.3 Direct GP Learning Across the Parameter r (No Embedding)

To test how well a nonparametric regressor can learn the logistic map directly, we chose *not* to use delay embedding here. For each parameter value $r \in [2.5, 4.0]$, we built a training set of 50 noiseless samples on $[0, 1]$: $\{(x_i,\ f_r(x_i))\}_{i=1}^{50}$ with $f_r(x) = r\,x(1 - x)$ and x_i evenly spaced. A Gaussian Process (GP) with an RBF kernel was then trained on these pairs to approximate the one-step map $x_{t+1} = f_r(x_t)$. We iterated the learned GP map (starting from a fixed x_0) and plotted only the later part of each trajectory, just as we did for the true logistic map, to form a bifurcation diagram.

The resulting "Figure 3 – GP bifurcation diagram" matches the true diagram very closely across fixed-point, periodic, and chaotic regions. Even near the start of chaos and in fully chaotic ranges, the GP iterates reproduce the correct branching and spread of the attractor. This shows that with only a modest set of samples at each r, GP regression can learn and iterate a highly nonlinear map with high accuracy-capturing both the qualitative and much of the quantitative geometry of the chaotic attractor-without relying on delay embedding.

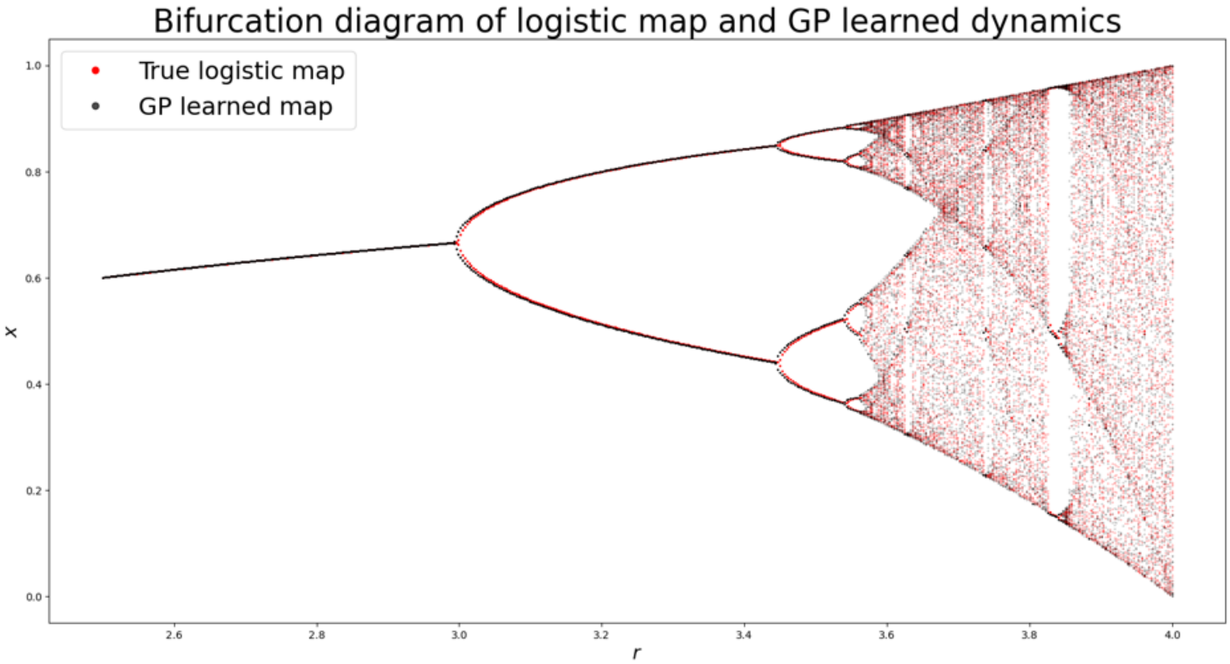

Figure 3: **GP bifurcation diagram.** At each parameter value r, $N = 50$ evenly spaced samples are used for training and evaluation, keeping comparisons consistent across periodic and chaotic dynamics.

2.4 GPR with Limited Data and Noise

We tested GP learning under extreme data scarcity by training on only five samples from the logistic map while varying the observation noise level $\sigma \in \{0.00, 0.02, 0.05, 0.10\}$. In Figure 4, the dotted red curve shows the true logistic map $f_r(x) = r\,x(1 - x)$ and the solid blue curve is the GP posterior mean fit.

We compared two settings for the GP noise parameter:

- **(a) Noise matched**: α was set to the square of the noise standard deviation used to generate the data, that is $\alpha = \sigma^2$ (so $\alpha = 0$, 4×10^{-4}, 2.5×10^{-3}, 10^{-2}). This encourages the GP to balance data fit and smoothness according to the assumed noise. With only five training points, however, the implied α is still too small when noise is low, and the fit shows near interpolation or overfitting for $\sigma \in \{0.00, 0.02, 0.05\}$. By $\sigma = 0.10$, the larger α produces visibly smoother and more stable estimates.

- **(b) Fixed regularization**: α was fixed at 0.005 across all noise levels, independent of σ. This works as a constant regularizer, giving consistently smoother fits that preserve the global curve of the logistic parabola while avoiding the high-variance fluctuations seen in (a) under low noise.

Overall, matching α to σ^2 is reasonable for Gaussian noise, but with very few training points it can provide too little regularization when σ is small, leading to overfit. A modest fixed α (here 0.005) offers reliable performance across noise settings when the true variance is unknown or the sample size is very limited.

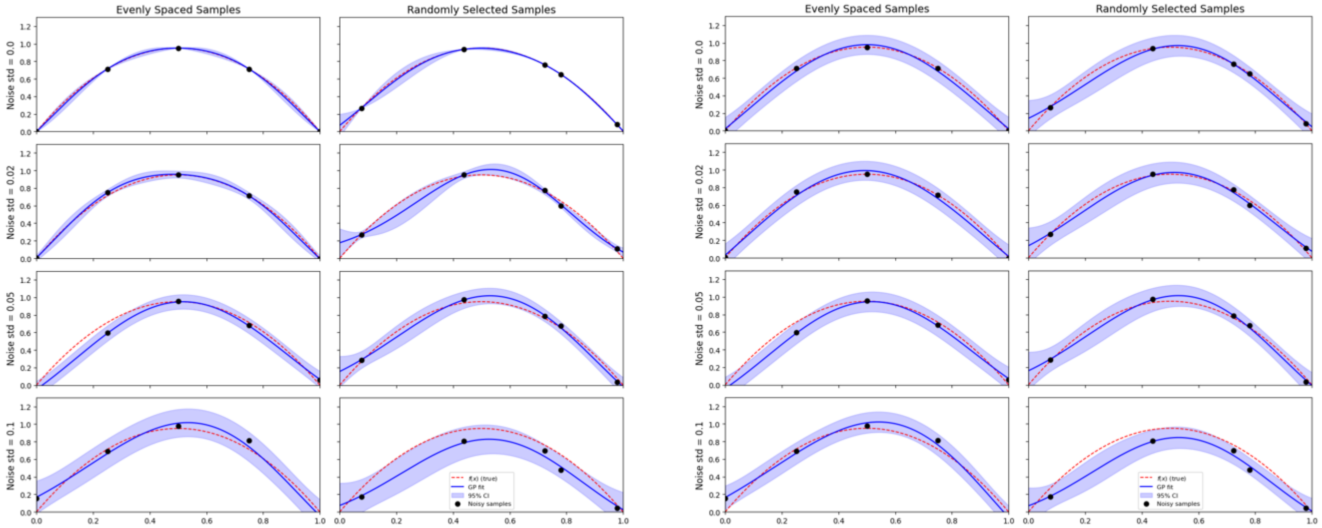

(a) **Noise matched** $\alpha = \sigma^2$ (b) **Fixed** $\alpha = 0.005$ (independent of σ)

Figure 4: **GPR on the logistic map with only 5 training points.** Dotted red shows the true logistic map; solid blue shows the GP fit. **(a)** Using $\alpha = \sigma^2$ balances data fit by the assumed noise, but with five samples it still overfits for $\sigma \in \{0, 0.02, 0.05\}$. By $\sigma = 0.10$ the fit smooths noticeably. **(b)** A fixed $\alpha = 0.005$ produces smoother and more consistent fits across all noise levels, trading exact interpolation for stability and faithful global shape.

3 Hénon Map

In this section we use the two dimensional Hénon map as a standard example for state space reconstruction and learning from partial observations. We start by recalling its quadratic form and the well known boomerang shaped strange attractor at $(a, b) = (1.4, 0.3)$, which gives a low dimensional but strongly dissipative setting for testing methods. We then study how the choice of measurement function $h(x, y)$ changes delay reconstructions: among several options, smooth low frequency functions (for example $h(x, y) = \sin(x/2) + \cos(y/2)$) produce the clearest 2D and 3D embeddings, while highly oscillatory or multiplicative functions cause tangling and artificial overlaps. Next, after fixing a good h, we vary the delay τ and show the trade off between redundancy (when τ is too small) and spreading (when τ is too large) in low dimensional reconstructions. Moving to 4D embeddings helps separate trajectories further, although very high frequency measurements remain unstable. To check embedding sufficiency, we perform a False Nearest Neighbors (FNN) analysis across Q and h, finding that smooth measurement functions need smaller Q while oscillatory ones require larger values. Finally, we train Gaussian Process Regression (GPR) models on the reconstructed coordinates and compare the learned dynamics to the original attractor. With more training points the Wasserstein distance decreases, and even under added noise the method stays qualitatively robust, though extra data improves both accuracy and stability. Altogether, these experiments give practical guidelines for selecting measurements, delays, and embedding dimensions before applying nonparametric learning to chaotic planar dynamics.

3.1 Hénon Map: General Information

The Hénon map is a two dimensional quadratic map defined as

$$x_{t+1} = 1 - a x_t^2 + y_t, \quad y_{t+1} = b x_t,$$

where (a, b) are parameters. At the classical chaotic values $a = 1.4, b = 0.3$, the system produces the well known boomerang shaped strange attractor, which has fractal dimension ≈ 1.26.

This attractor is a standard test case in nonlinear dynamics. Even though the equations are simple polynomials, the system produces rich chaotic behavior with strong sensitivity to initial conditions. Beyond its mathematical appeal, the Hénon map is historically important as a simplified model for turbulence and weather like unpredictability. Because it is low dimensional and strongly contracting, it gives a setting that is simple enough to handle but still complex enough to be a real test for state space reconstruction and learning methods.

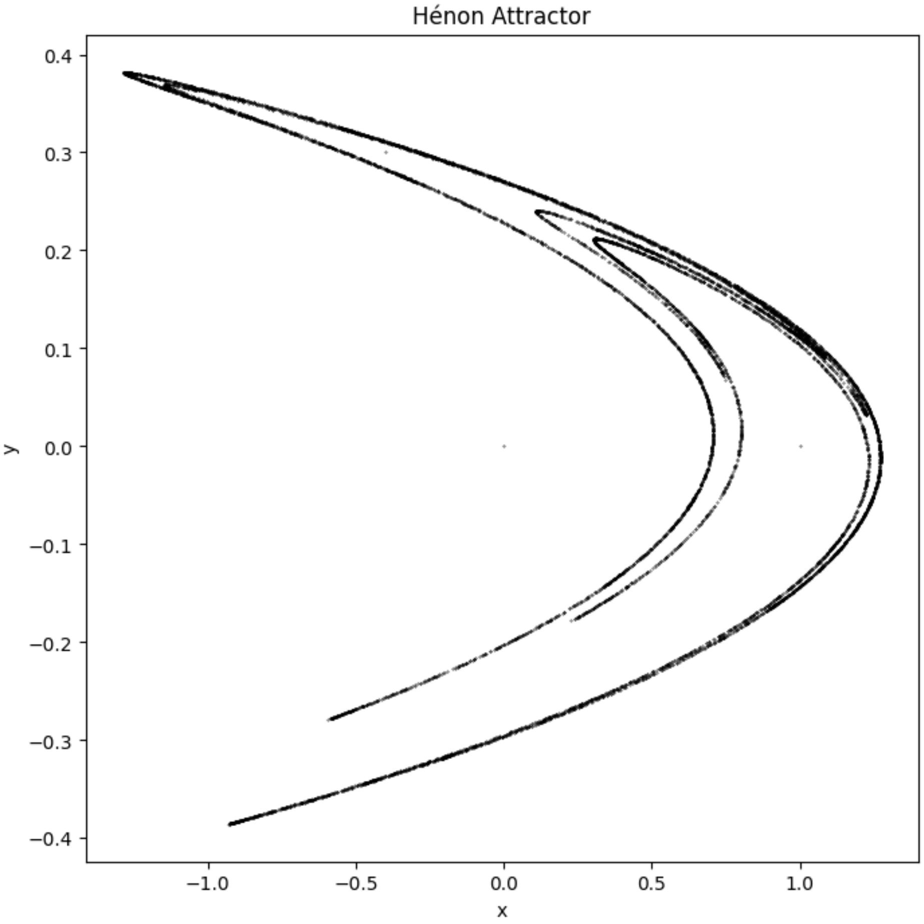

Figure 5: **Hénon attractor at** $(a, b) = (1.4, 0.3)$**.**

3.2 2D and 3D Delay Reconstructions with Measurement Functions

To study how different measurement functions affect reconstruction, we applied delay embedding with $\tau = 1$ to univariate series $z_t = h(x_t, y_t)$ generated from the Hénon attractor at $(a, b) = (1.4, 0.3)$. We tested four representative measurement functions (Figure 6):

$$h_1(x, y) = \sin(x/2) + \cos(y/2), \quad h_2(x, y) = \sin(6x), \quad h_3(x, y) = \sin(xy), \quad h_4(x, y) = x - y.$$

- $h_1(x, y) = \sin(x/2) + \cos(y/2)$: Produced the cleanest reconstruction. Both 2D and 3D embeddings preserve the boomerang like geometry of the original Hénon attractor without visible intersections. Although chaotic, the structure remains well organized, making this the best candidate.

- $h_2(x, y) = \sin(6x)$: The high frequency oscillations caused very tangled trajectories. Reconstructions appear extremely chaotic and highly sensitive to small changes, making this function poor for predictive modeling.

- $h_3(x, y) = \sin(xy)$: Created embeddings with noticeable overlaps, showing that different dynamical states can collapse to the same observed values. Such overlap limits faithful state space reconstruction.
- $h_4(x, y) = x - y$: Produced reasonably unfolded embeddings without severe distortion. While it resembles the original attractor, its geometry is less stable than h_1 and still shows some flattening.

Among the tested measurement functions, $h_1(x, y) = \sin(x/2) + \cos(y/2)$ preserves the Hénon attractor most clearly in both 2D and 3D embeddings. Its clean geometry without intersections suggests strong potential for later Gaussian Process Regression (GPR) modeling.

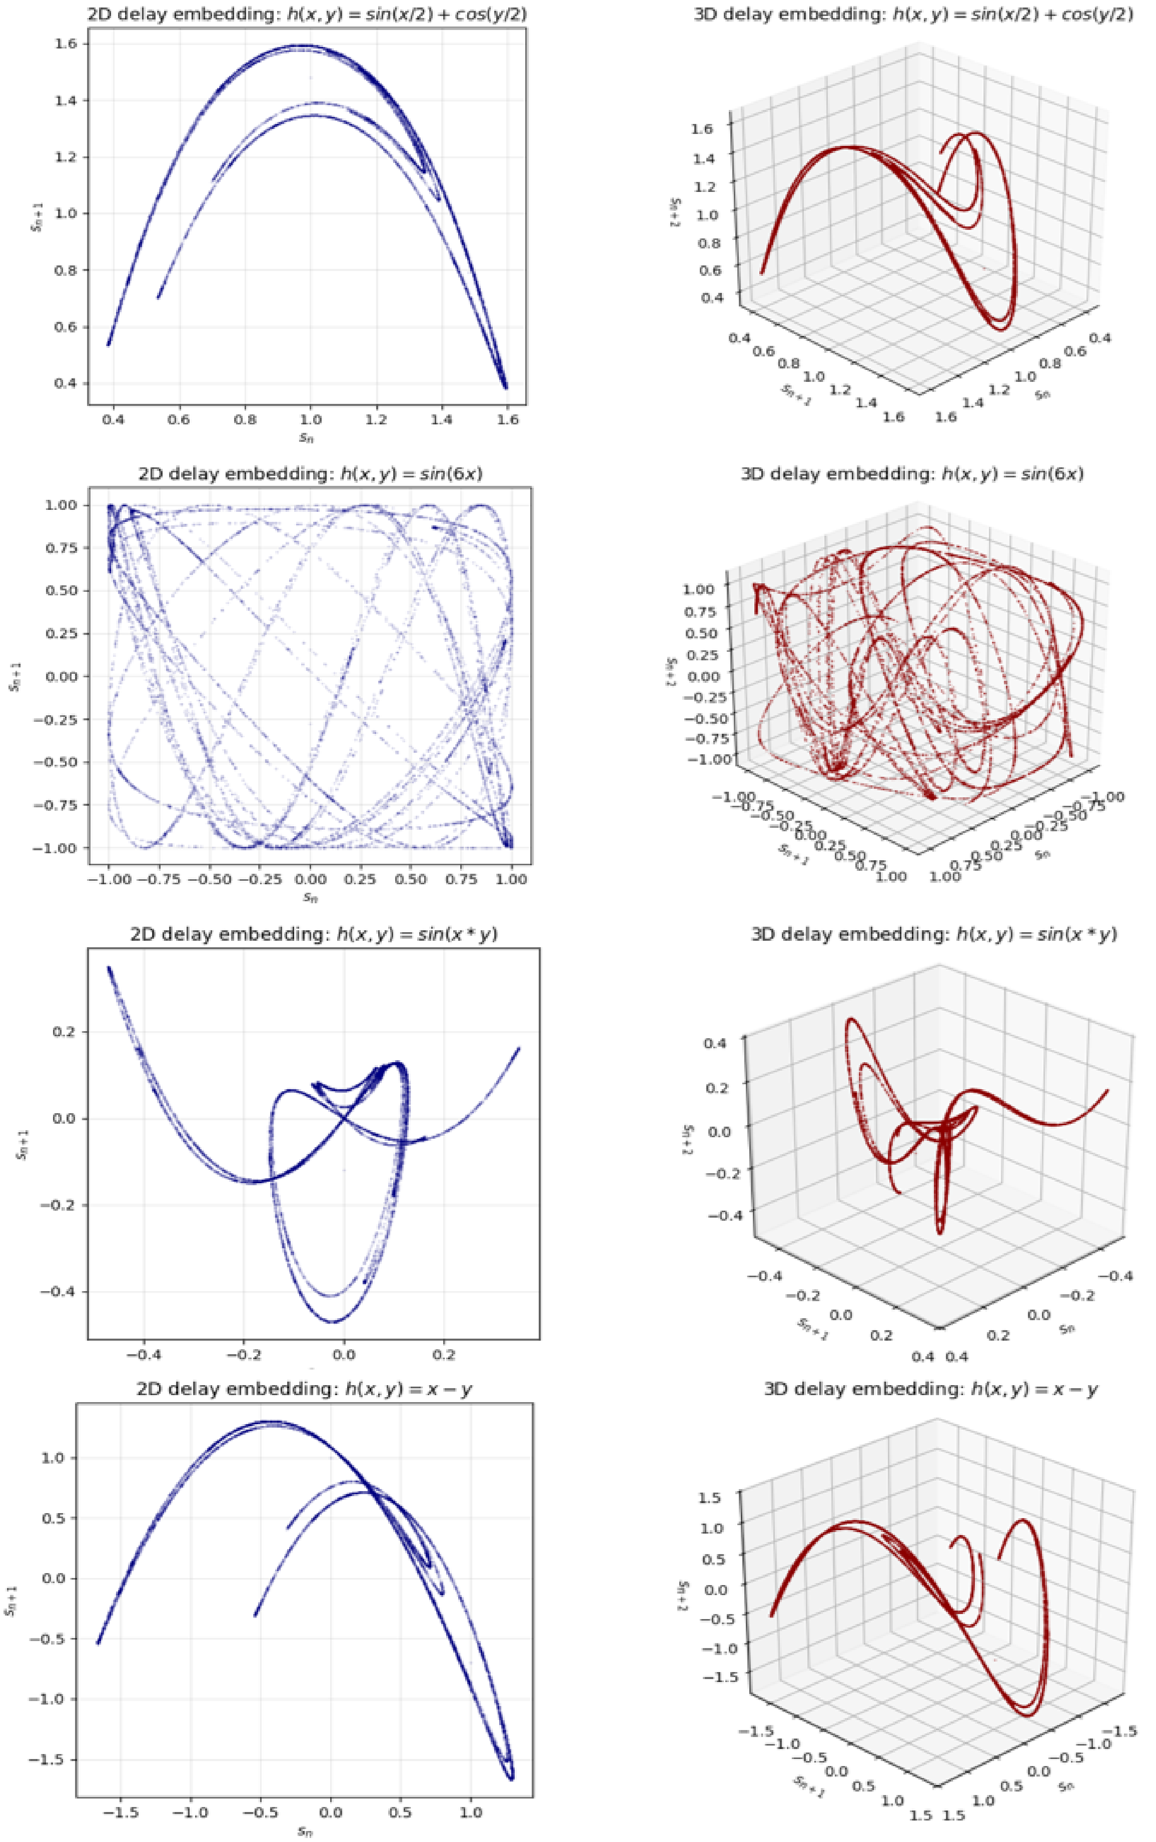

Figure 6: **2D and 3D delay embeddings of the Hénon map using four measurement functions ($\tau = 1$).** Top: $h(x, y) = \sin(x/2) + \cos(y/2)$; second row: $h(x, y) = \sin(6x)$; third row: $h(x, y) = \sin(xy)$; bottom: $h(x, y) = x - y$. The first column

3.3 3D Delay Reconstructions with Varying Delays

We next fixed the measurement function as

$$h(x, y) = \sin(x/2) + \cos(y/2),$$

which gave the clearest structure in Section 3.2. We then varied the delay parameter $\tau = 1, 2, \ldots, 6$ while keeping the embedding dimension fixed at $Q = 3$ (Figure 7).

- For $\tau = 1$ and $\tau = 2$, the reconstructed attractor kept a clear boomerang shape with only minor distortion.
- For larger delays such as $\tau \geq 3$, the embeddings became tangled and visually harder to interpret. The attractor's shape looked flattened or folded in ways that hid its original organization.

It might seem natural to assume that smaller delays always work better, but this is not always true. The best value of τ depends on how it interacts with the chosen embedding dimension Q. If the delay is too small, the coordinates become redundant and fail to open up the attractor. If the delay is too large, the points spread out too much and local details are lost. Takens' theorem guarantees that reconstruction is possible when Q is large enough, but in practice choosing both τ and Q requires balancing overlap and independence.

In our 3D visualizations, we could see that geometric clarity dropped off as delays grew larger. Higher-dimensional embeddings could help separate the trajectories, but then the results become difficult to visualize. This experiment shows a practical tradeoff: in low-dimensional reconstructions, small delays ($\tau = 1$ or $\tau = 2$) give the cleanest results, while larger delays reduce interpretability.

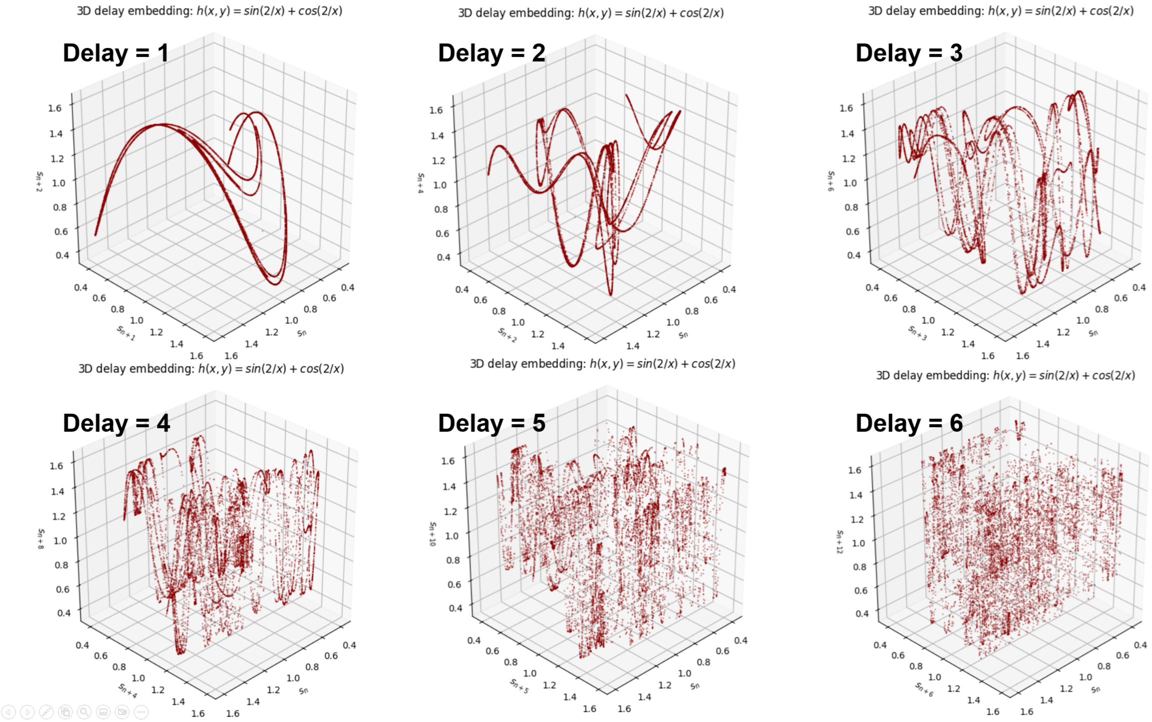

Figure 7: **3D delay embeddings of the Hénon map with delays $\tau = 1$ to 6.** Using $h(x, y) = \sin(x/2) + \cos(y/2)$ and embedding dimension $Q = 3$, small delays ($\tau = 1, 2$) preserve the attractor's shape, while larger delays distort the geometry.

3.4 4D Delay Reconstructions of Hénon Map

3.4.1 With the measurement function $h(x, y) = \sin(2/x) + \cos(2/x)$

We built a 4D delay embedding from the univariate series $z_t = h(x_t, y_t)$ of the Hénon map at $(a, b) = (1.4, 0.3)$ with $\tau = 1$. In Figure 8, the embedding (s_n, s_{n+1}, s_{n+2}) is shown with colors based on the next coordinate s_{n+3}, which helps visualize how the system evolves.

From inspection, the 3D reconstruction already provides a reasonable picture of the attractor. The 4D embedding, however, separates trajectories more clearly and reduces overlaps. This matches the idea from Takens' theorem, which states that using a large enough embedding dimension makes it possible to recover the original dynamics. In practice, coloring by s_{n+3} shows smooth changes along the attractor and confirms that the 4D reconstruction works as intended.

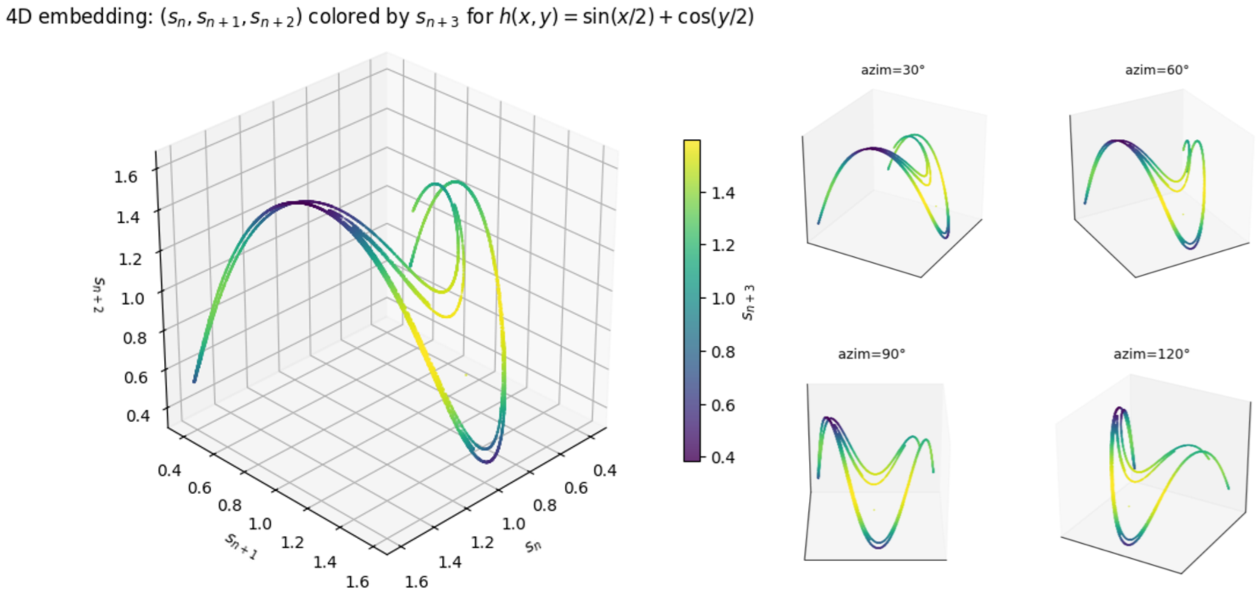

Figure 8: **4D delay reconstruction with** $h(x, y) = \sin(2/x) + \cos(2/x)$**,** $\tau = 1$**.** The embedding (s_n, s_{n+1}, s_{n+2}) is colored by s_{n+3}. While the 3D view already captures much of the attractor, the 4D embedding unfolds the dynamics more cleanly.

3.4.2 Testing High Frequency Measurement Functions

We also tested measurement functions of the form $h(x) = \sin(kx)$ with $k \in \{1, 2, 4, 6, 8, 10\}$, using 4D delay embeddings with lag $\tau = 1$. Figure 9 shows the resulting reconstructions.

- For low frequencies ($k = 1, 2$), the embeddings produced smooth and clear attractor shapes, fairly close to the original Hénon geometry.
- At moderate frequency ($k = 4$), the structure became more complex, with extra folding but still possible to interpret.
- For higher frequencies ($k \geq 6$), the trajectories broke into tangled curves. The embeddings looked noisy and harder to interpret, showing strong sensitivity to small changes in the time series.

Increasing k amplifies small variations in x, which makes the chaotic fluctuations already in the system appear stronger. As a result, high frequency measurement functions can create embeddings that look more detailed but are also more fragile to noise and limited data. This experiment suggests that moderate oscillatory functions (such as $k = 2$ or $k = 4$) can enrich the reconstructed geometry, but very high frequencies ($k \geq 6$) obscure the attractor's overall structure. For later modeling (such as GPR on ecological data), this instability makes high frequency choices less practical. In contrast, functions like $h(x,y) = \sin(x/2) + \cos(y/2)$ strike a better balance between detail and robustness.

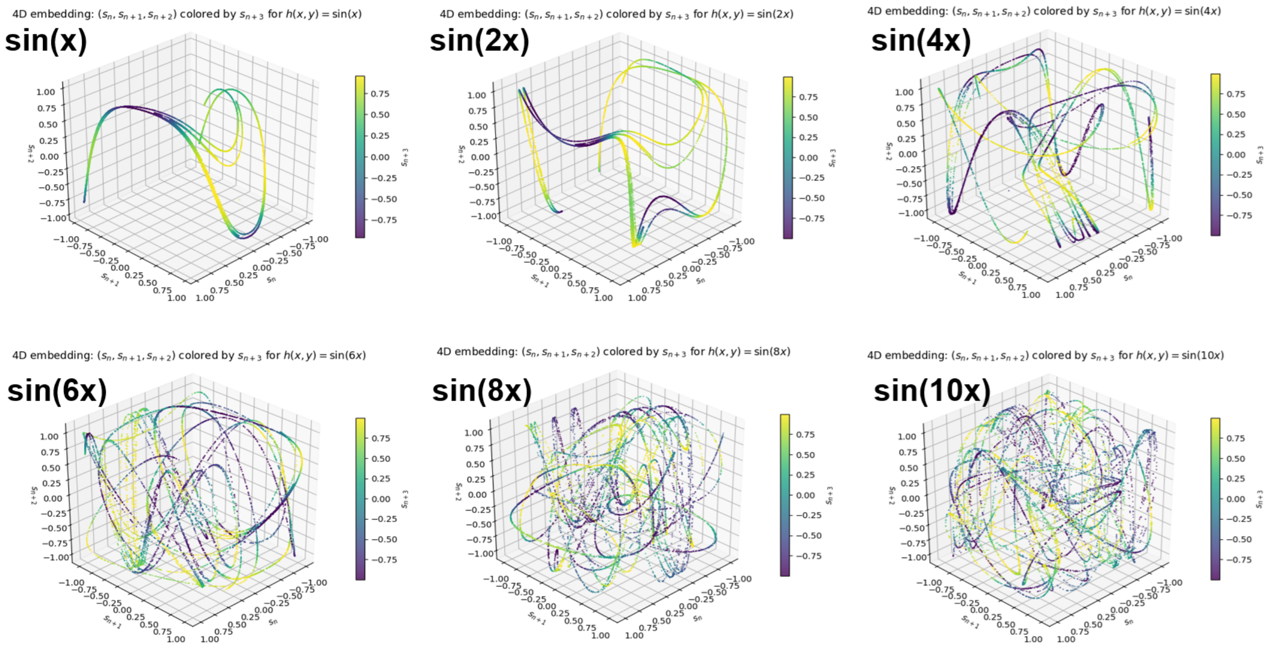

Figure 9: **4D delay reconstructions with $h(x) = \sin(kx)$, $k = 1, 2, 4, 6, 8, 10$ (left to right).** Low and moderate frequencies preserve attractor geometry, while very high frequencies fragment the structure and increase sensitivity to noise.

3.5 FNN Test Across Functions and Dimensions

In Figure 10, we evaluated embedding quality by computing the fraction of False Nearest Neighbors (FNN) as a function of embedding dimension Q for three measurement functions: $h(x,y) = \sin(6x)$, $h(x,y) = \sin(x/2) + \cos(y/2)$, and $h(x,y) = x - y$. A 5% threshold (shown as a reference line) is commonly used in practice: once the FNN fraction falls below this level, the attractor is usually well unfolded and less affected by false overlaps from projection.

- For the high frequency measurement $h(x,y) = \sin(6x)$, the FNN fraction dropped below 5% only at $Q = 3$. This shows that a higher embedding dimension was needed to untangle overlaps created by the oscillatory function.
- For both $h(x,y) = \sin(x/2) + \cos(y/2)$ and $h(x,y) = x - y$, the FNN fraction already fell below the threshold at $Q = 2$. In these cases, even a low dimensional embedding was enough for a faithful reconstruction.

These results show that the embedding dimension can depend on the measurement function. Smooth functions that preserve the overall geometry (such as $\sin(x/2) + \cos(y/2)$) unfold the attractor efficiently, while oscillatory functions (such as $\sin(6x)$) require higher dimensions to correct for artificial overlaps.

It is important to note that in real data we rarely know the true measurement function. Whenever we record a time series, such as temperature in climate studies or biomass in ecology, the observed values are already the output of some measurement function that maps the full state of the system to a single variable. Delay embedding is powerful because it can reconstruct much of the hidden geometry even when the exact measurement function is unknown. This makes the careful choice of measurement functions in synthetic experiments valuable, since it helps us understand what may or may not work when applying the method to real-world data.

In practice, this means one does not always need very high Q. With a well behaved measurement function, a modest embedding dimension can capture the dynamics. For ecological data, however, noise and limited samples may complicate reconstruction. In such cases, higher Q values may provide redundancy but also increase the risk of overfitting. Thus, both the choice of measurement function and the embedding dimension play crucial roles in successful modeling.

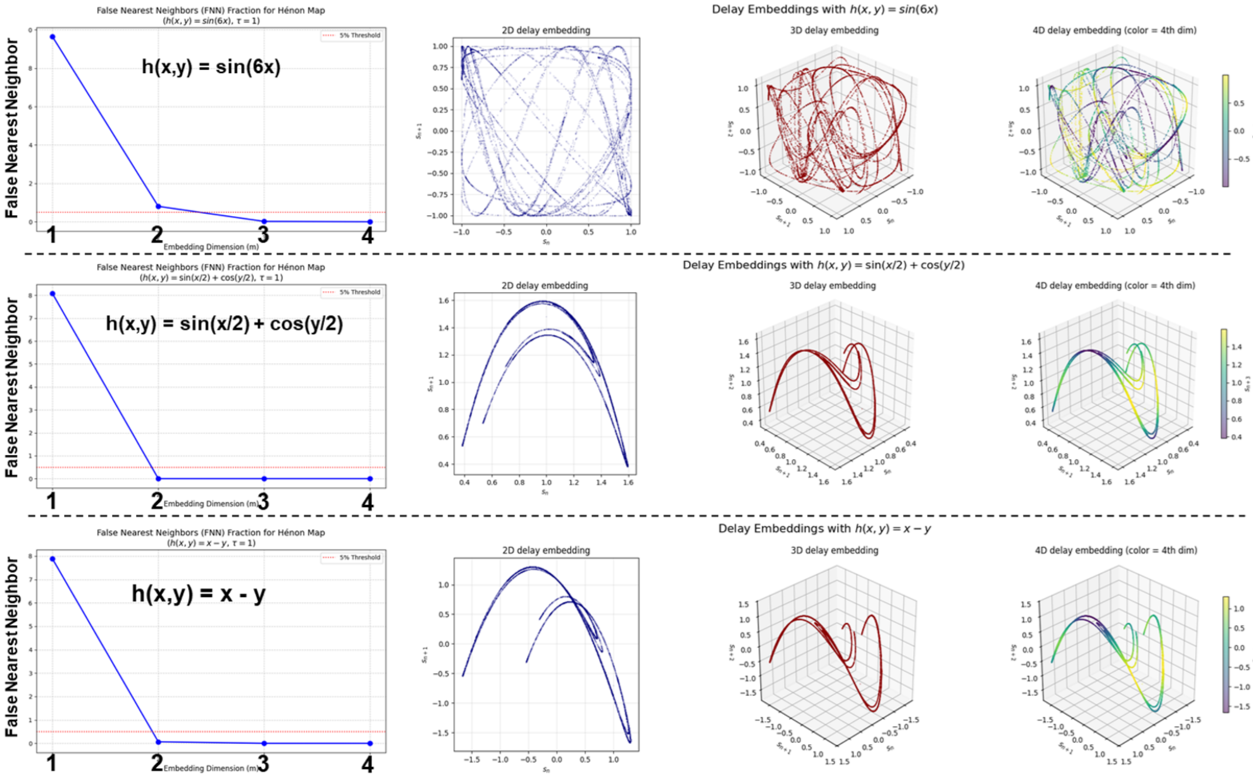

Figure 10: **FNN analysis for three measurement functions.** For $h(x, y) = \sin(6x)$, the FNN fraction falls below the 5% threshold only at $Q = 3$. For $h(x, y) = \sin(x/2) + \cos(y/2)$ and $h(x, y) = x - y$, the threshold is already met at $Q = 2$. This shows that suitable measurement functions can reduce the need for high embedding dimensions.

3.6 Comparing Reconstructions vs GPR Approximation

In this experiment, we tested how the number of training samples affects the quality of Gaussian Process Regression (GPR) reconstructions of the Hénon map. The steps were: 1) Apply the measurement function $h(x, y) = \sin(x/2) + \cos(y/2)$ to the original Hénon trajectory ($N = 2000$ points). 2) Randomly select a small subset (20, 30, or 40 points). 3) Perform delay embedding with $\tau = 1$ in 2D and 3D. 4) Train GPR models on the embedded data and generate predictions on 2000 points. 5) Compare the reconstructed attractor against the original using Wasserstein distance (computed with the POT library in Python).

- 20 samples: distances were 0.0343 (2D) and 0.0338 (3D).
- 30 samples: distances were 0.0273 (2D) and 0.0446 (3D).
- 40 samples: distances dropped to 0.0119 (2D) and 0.0236 (3D).

As shown in Figure 11, even with only 20 samples, GPR captured the global structure of the Hénon attractor, giving relatively small Wasserstein distances. Adding more samples steadily reduced the distance, showing that extra data improves the accuracy of the learned dynamics. While results varied a little between 2D and 3D embeddings, the main trend was clear: reconstructions became more accurate as sample size increased. These findings highlight the sample efficiency of GPR. Faithful approximations of chaotic attractors can be built from very sparse data, and robustness improves further as more training points are added.

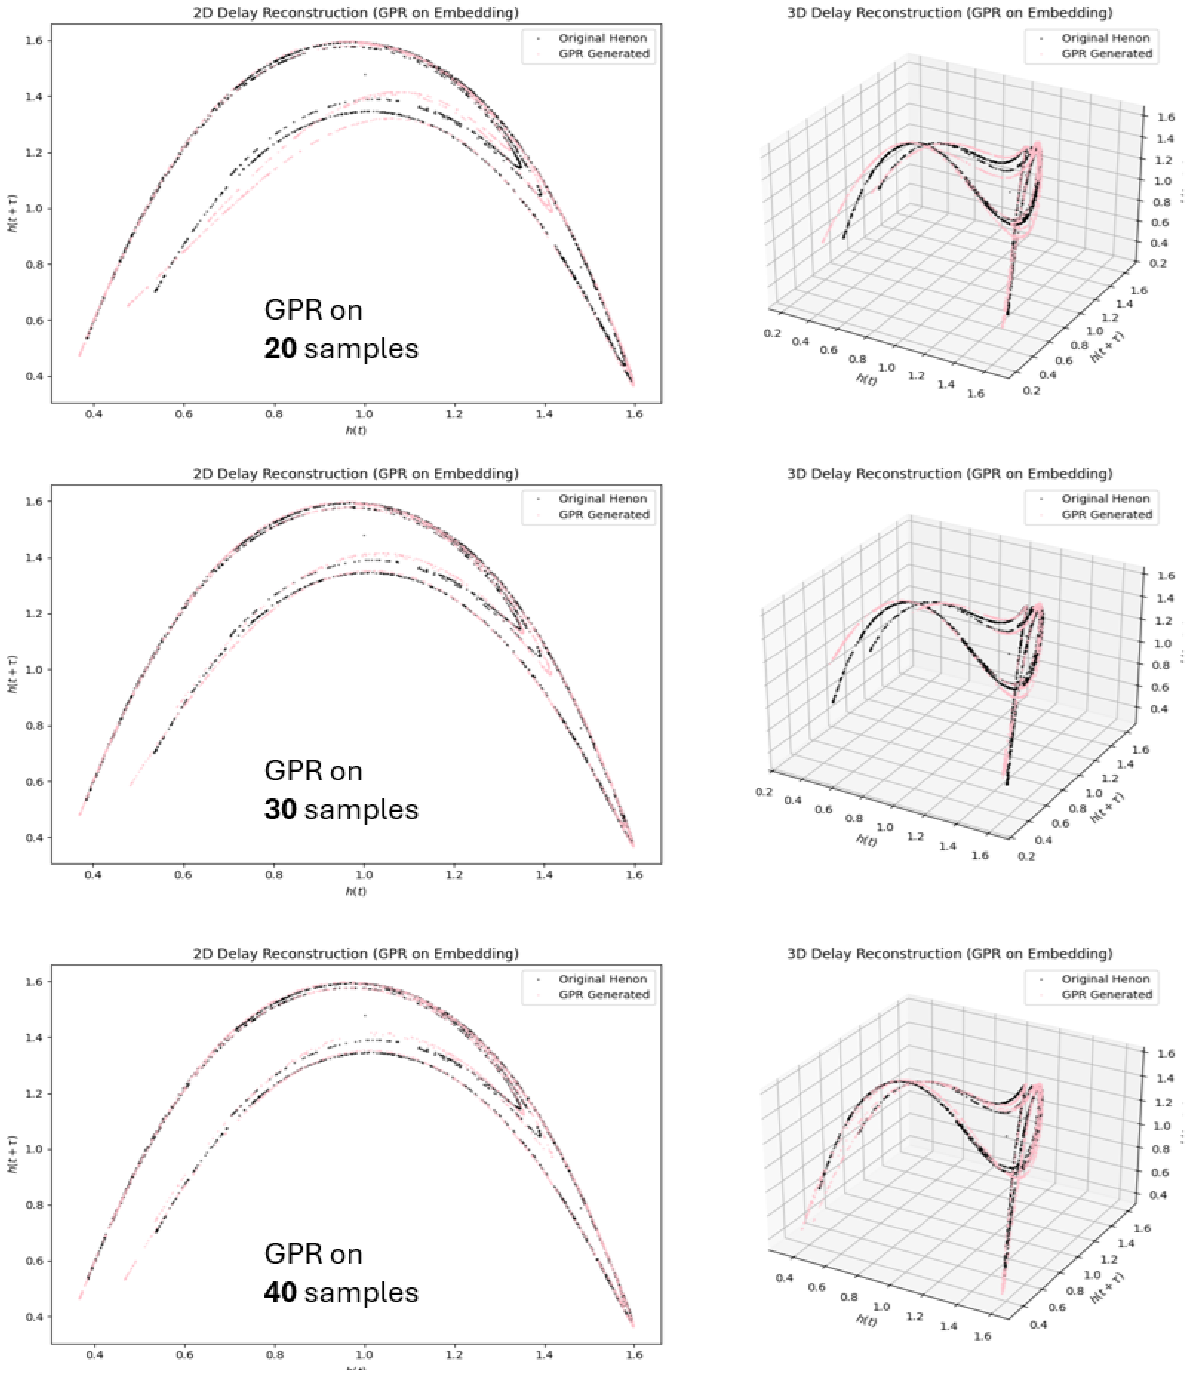

Figure 11: GPR reconstructions compared to the original Hénon attractor. Columns: 2D (left) and 3D (right) delay embeddings. Rows: training sample sizes of 20, 30, and 40. Black shows the original data, pink shows GPR predictions. Wasserstein distances decrease as sample size grows, confirming that GPR improves with more training data while still performing well with sparse samples.

3.7 Noisy Data Reconstructions and GPR

To assess robustness under measurement error, we repeated the reconstruction and GPR procedure using noisy data. The Hénon map was generated with parameters $(a, b) = (1.4, 0.3)$, and a 1D time series was obtained from the measurement function $h(x, y) = \sin(x/2) + \cos(y/2)$. Gaussian noise with standard deviation 0.05 was added to simulate real-world inaccuracies. The noisy series was then delay-embedded with $\tau = 1$ to form 2D and 3D reconstructions. From these, subsets of 40, 100, and 200 points were randomly sampled to train GPR models. Each experiment was repeated 100 times with independent resampling, and Wasserstein distances were computed between 2000 GPR-generated points and 2000 original (noise-free) points.

Figure 12 shows representative outcomes from one trial at each sample size:

- 40 samples: The GPR reconstruction recovered the broad boomerang outline but left large gaps. Many regions of the attractor were sparsely covered, resulting in rough and incomplete dynamics (Wasserstein ≈ 0.36–0.40).
- 100 samples: Reconstructions improved markedly. The orange GPR curves followed the blue noise-free attractor more closely, filling gaps and producing a smoother geometry, especially in 3D (Wasserstein ≈ 0.30).
- 200 samples: The learned attractor nearly overlapped with the true one. The global shape was well captured, and Wasserstein distances were lowest (around 0.26–0.28), indicating the best accuracy.

Across 100 repeated trials, the average Wasserstein distance decreased steadily with sample size: from 0.4171 ± 0.1768 (2D, 40 samples) to 0.3330 ± 0.1258 (2D, 200 samples), with similar trends in 3D. Normalized distances likewise improved, confirming that larger training sets yield both lower error and greater stability. Importantly, even under noisy conditions the reconstructions remained faithful to the true attractor, capturing its global boomerang shape from as few as 40 samples.

In summary, Figure 12 demonstrates that while small noisy samples already allow GPR to reproduce the overall attractor structure, increasing the number of training points significantly improves both accuracy and robustness of the reconstruction. This highlights the practical strength of GPR in recovering complicated chaotic geometries from corrupted and limited data, a situation common in real-world measurements.

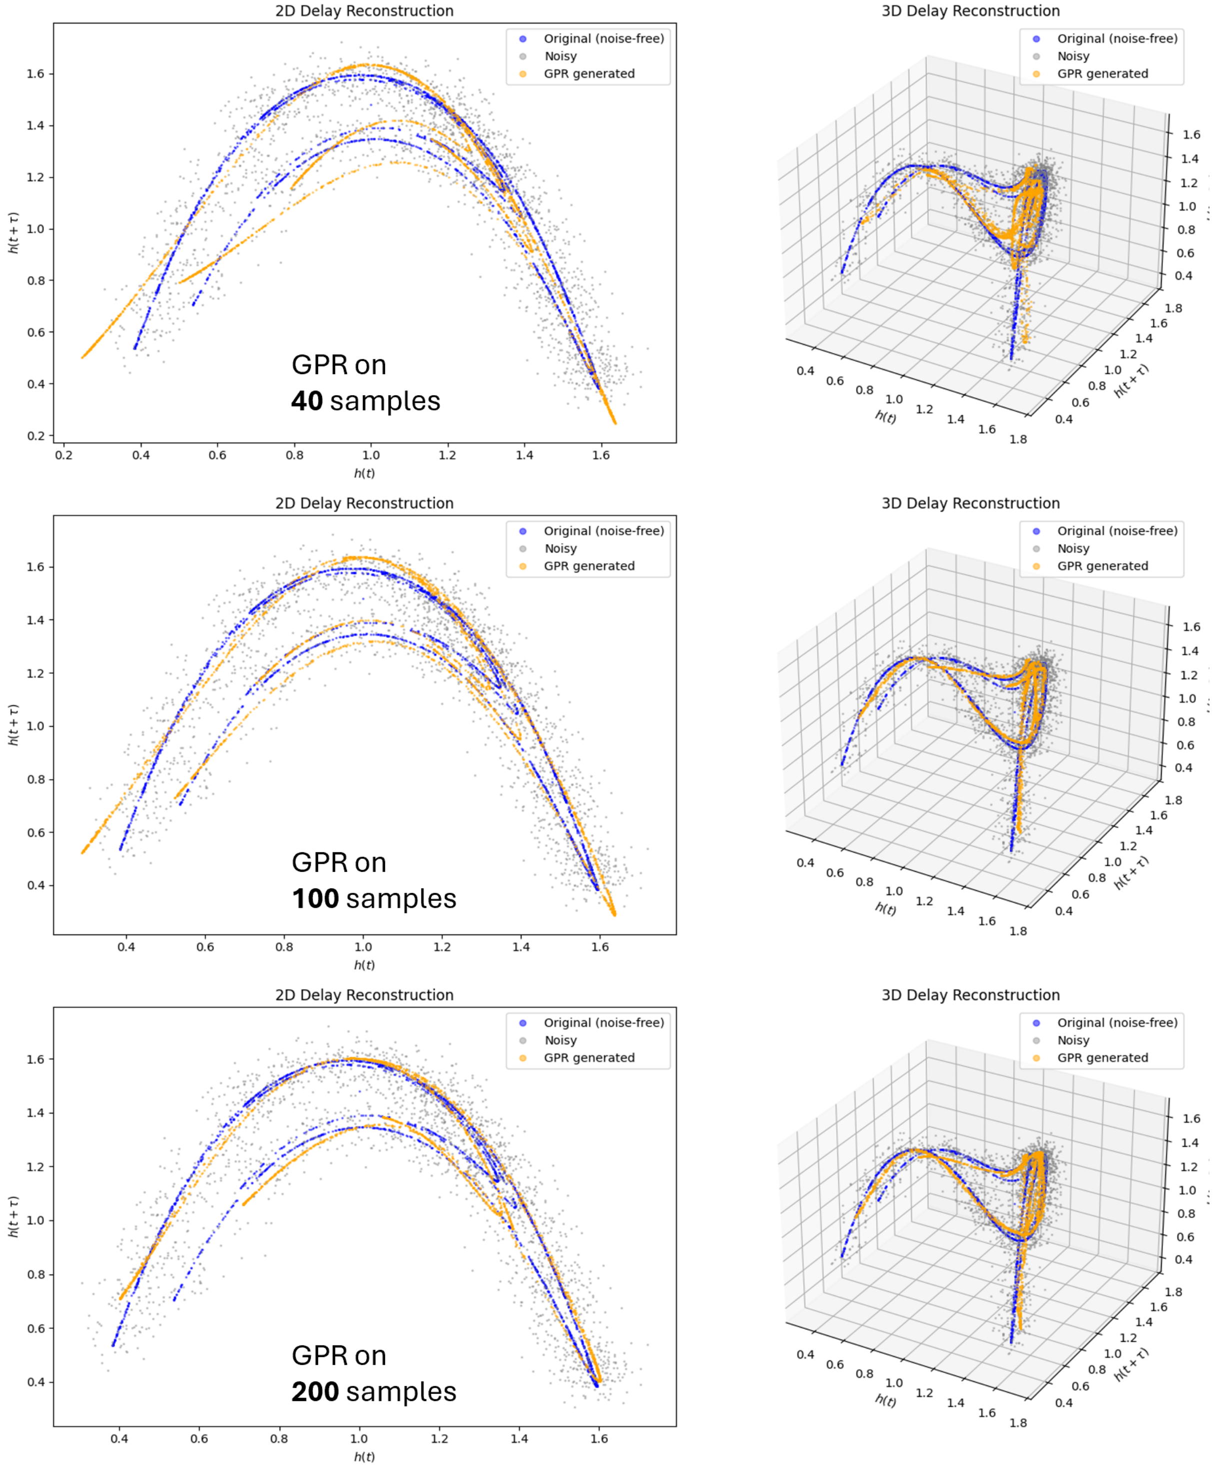

Figure 12: Delay embeddings and GPR reconstructions from noisy data. Training with 40, 100, and 200 samples (rows) in 2D (left) and 3D (right). Blue = original noise-free attractor, gray = noisy data, orange = GPR reconstruction. Reconstructions improve and Wasserstein distances decrease as sample size increases, while even the smallest sample size recovers the essential chaotic structure.

4 Real-World Ecological Data

We next applied the delay embedding and GPR framework to ecological time series. The dataset comes from a long-term biological control study of the invasive aquatic fern *Salvinia molesta* in the Northern

Territory of Australia (Schooler et al. 2011). Over several years, four field sites (Jabiluka, Jaja, Minggung, and Island Billabong) were surveyed for salvinia biomass and for damage caused by the biocontrol weevil *Cyrtobagous salviniae*.

Schooler et al. (2011) noted that the field data were highly irregular, variable, and nonlinear, but did not label them as "chaotic." In our study, we interpret the observed fluctuations as potentially showing chaos-like dynamics and use delay embedding with GPR to test whether this structure can be reconstructed from the ecological time series. This approach connects theory from dynamical systems with real biocontrol data, linking mathematical modeling to applied ecology.

4.1 Visualization of Site Time Series

We visualized the time series of salvinia biomass and weevil damage for the four sites in the study. Because Jaja and Minggung had long gaps and many missing records, we excluded them from further analysis and focused on the two most complete sites: Jabiluka and Island Billabong. The main focus of our analysis is the biomass of the invasive fern *Salvinia molesta*, so biomass (log-transformed dry weight) was used as the basis for delay embedding and GPR modeling.

For context, we also plotted the weevil damage index (logit-transformed when available), which shows herbivory pressure from *Cyrtobagous salviniae*. However, consistent with Schooler et al. (2011), these damage series are irregular and noisy, and they explain biomass dynamics much less effectively. For this reason, we restricted predictive modeling to salvinia biomass, while keeping the damage plots to provide ecological background.

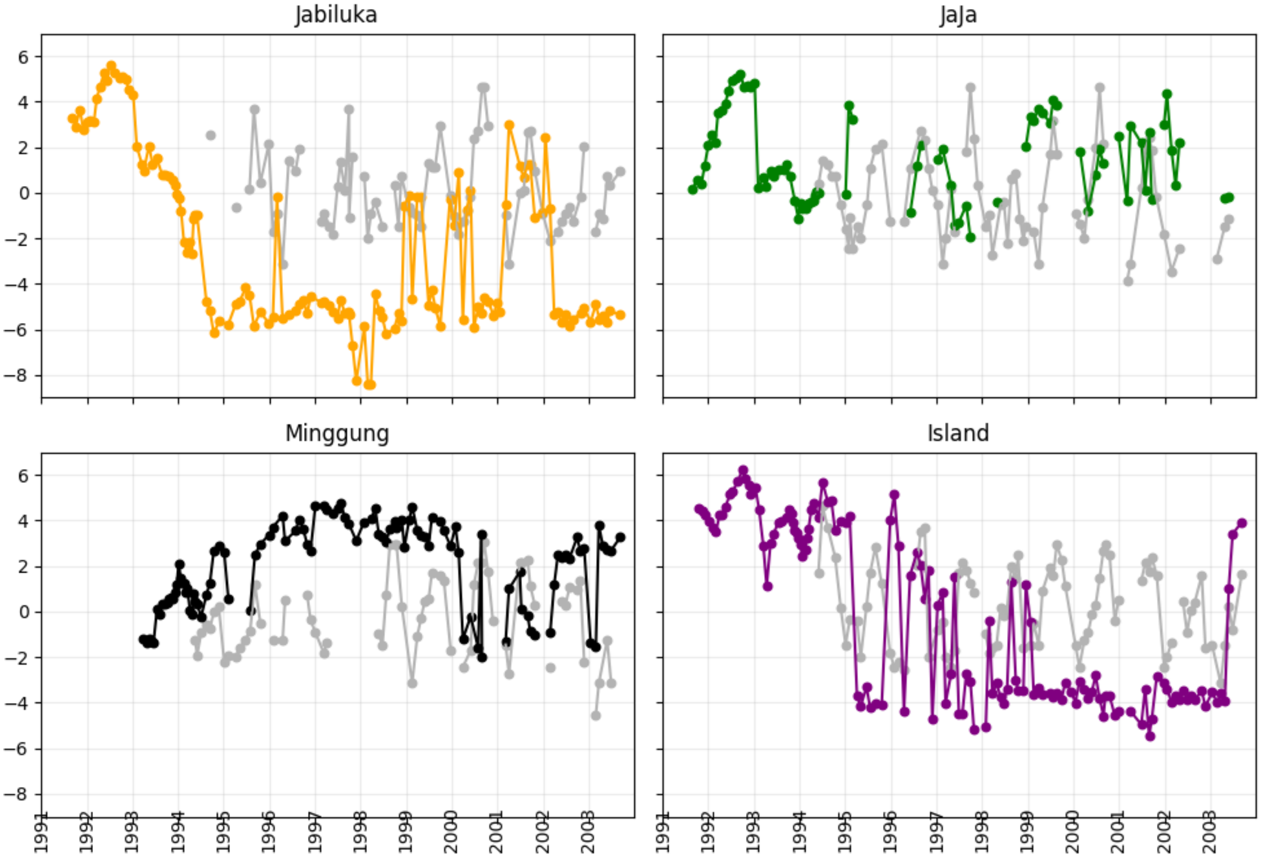

Figure 13: **Time series of salvinia biomass and damage across four sites.** Colored lines show log biomass of *Salvinia molesta*, gray lines show logit weevil damage. Gaps indicate missing observations.

4.2 Resampling and Interpolation

Ecological surveys of salvinia biomass were not taken at regular intervals. Sometimes there were multiple measurements within one month, while in other cases there were gaps of nearly two months. Such irregular spacing makes it difficult to apply delay embedding, which requires data on a regular time grid. To address this, as shown in Figure 14, we resampled each site to a uniform 30-day grid and then applied linear interpolation across neighboring points to fill moderate gaps:

$$x^{\text{interp}}(t) = x(t_i) + \frac{t - t_i}{t_{i+1} - t_i}\big(x(t_{i+1}) - x(t_i)\big), \quad t_i \le t \le t_{i+1}.$$

This produced evenly spaced time series suitable for embedding while still preserving the main trends of the raw data. Interpolation was applied only to moderate gaps. Long absences of data were left as breaks in the plots, reflecting true gaps in field sampling.

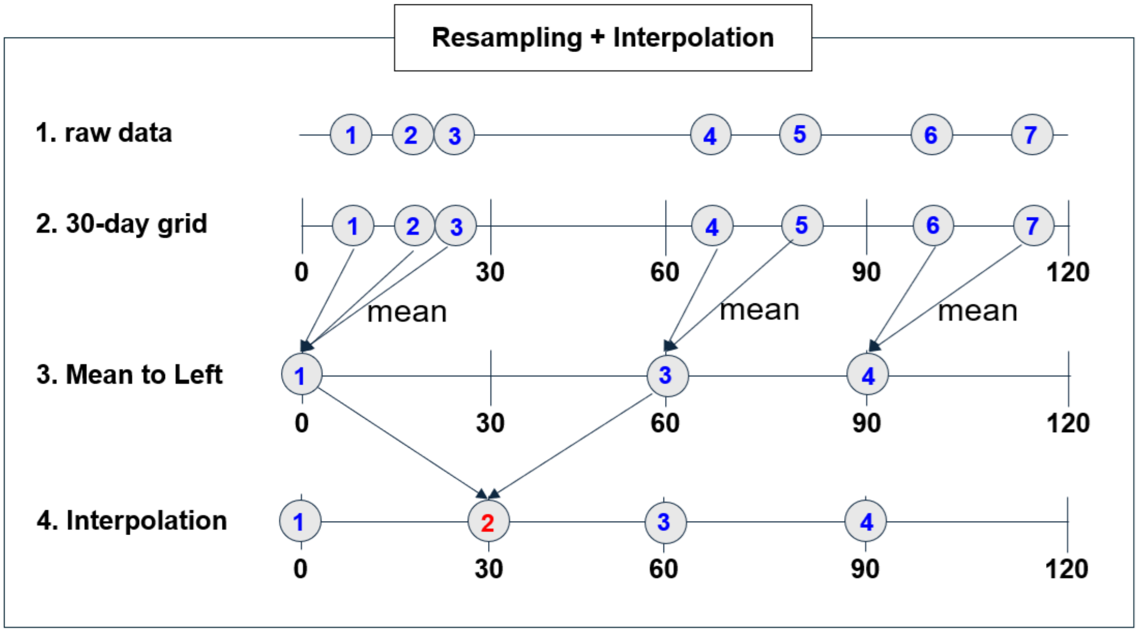

Figure 14: **Resampling and interpolation procedure.** Step 1: raw observations (numbered circles) collected at irregular intervals. Step 2: align observations to a uniform 30-day grid. Step 3: when multiple observations occur within a bin, aggregate them with a mean-to-left rule. Step 4: apply linear interpolation across adjacent grid points to fill moderate gaps. This pipeline produces evenly spaced data for delay embedding, while preserving the overall patterns of the raw series.

4.3 Delay Embedding and GPR Forecasts

We applied Takens' delay embedding and trained Gaussian Process Regression (GPR) models on the two most complete sites, Jabiluka and Island. For each site, we tested embedding dimensions $Q = 2, 3$ and lags $\tau = 1, 2$, training on 70% of the resampled data and forecasting the remaining 30%. Performance was measured using both R^2 and the Wasserstein distance between predicted and observed trajectories.

For Jabiluka, the Wasserstein distances were: $(Q, \tau) = (2, 1) : 2.591$, $(2, 2) : 2.719$, $(3, 1) : 3.311$, $(3, 2) : 3.413$. For Island, the distances were higher: $(2, 1) : 5.966$, $(2, 2) : 6.022$, $(3, 1) : 7.401$, $(3, 2) : 7.514$. In both sites, larger embedding dimensions Q gave consistently higher distances, and increasing τ also caused a smaller increase. These results suggest that broader experiments over a wider range of Q and τ values will be needed to confirm the trend.

Even so, the reconstructions captured the essential dynamics. In Figures 15 and 16, the red vertical line marks the train/test split. Predictions (black) follow the test data (gray) closely, preserving both short-term fluctuations and long-term trends. This shows that the delay embedding and GPR framework can generalize beyond the training data, even with noisy ecological series.

One more observation is that the embedding reconstructions (middle and right panels) show that predicted attractors preserve the overall shape of the actual delay spaces, although finer details are smoothed as expected with limited data. This balance between local predictive accuracy and global geometric fidelity illustrates both the strengths and the limits of the method for ecological time series.

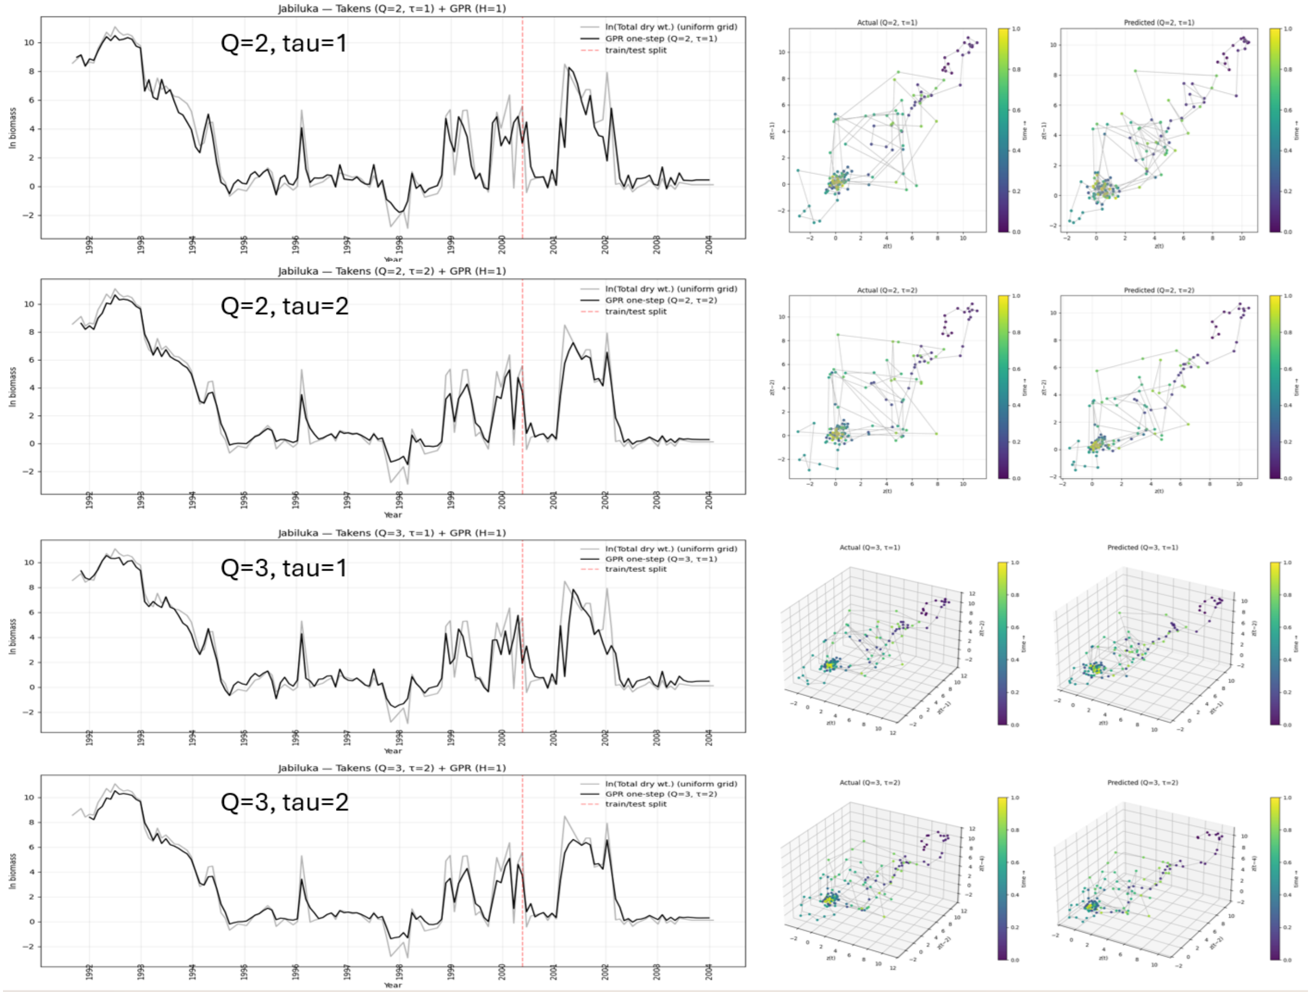

Figure 15: **Delay embedding and GPR reconstruction of salvinia biomass at Jabiluka.** Left: time series with train/test split (red line). Black = GPR predictions, gray = resampled biomass. Middle and right: actual vs. predicted embeddings at selected (Q, τ) values. Reported Wasserstein distances: (2,1)=2.591, (2,2)=2.719, (3,1)=3.311, (3,2)=3.413.

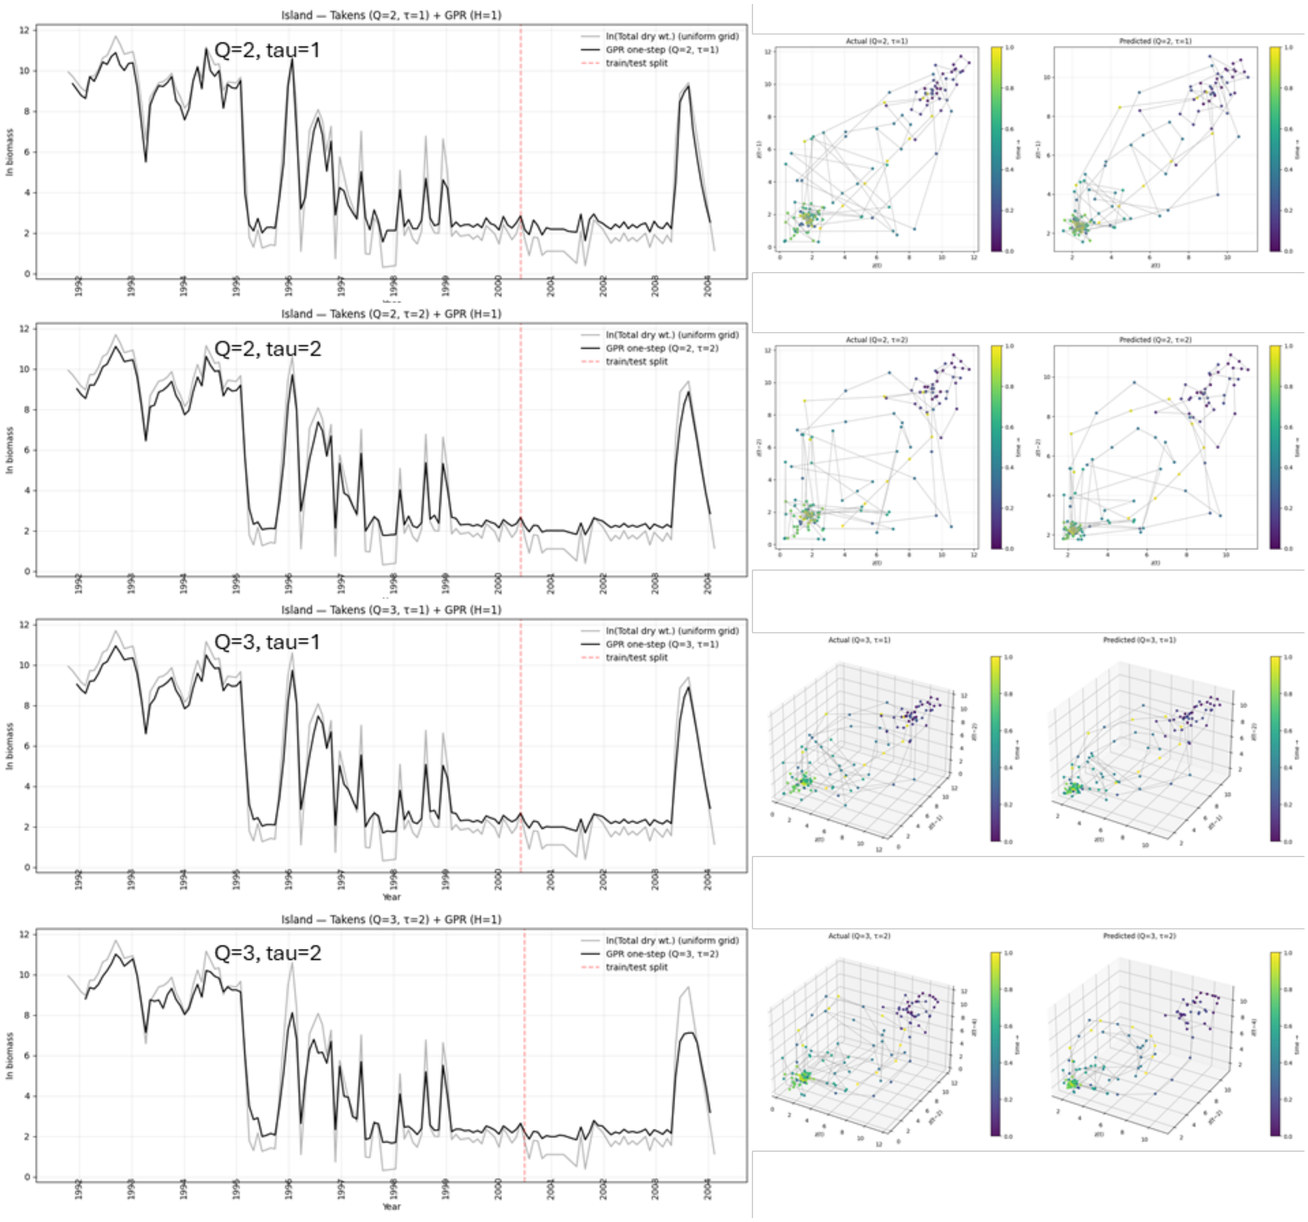

Figure 16: **Delay embedding and GPR reconstruction of salvinia biomass at Island.** Left: time series with train/test split (red line). Black = GPR predictions, gray = resampled biomass. Middle and right: actual vs. predicted embeddings at selected (Q, τ) values. Reported Wasserstein distances: (2,1)=5.966, (2,2)=6.022, (3,1)=7.401, (3,2)=7.514.

5 Results and Discussion

Our experiments show that the delay embedding and GPR framework can capture important aspects of chaotic dynamics and extend to ecological data. On the synthetic logistic and Hénon maps, the method recovered the global structure of the attractors even when training data were very limited. Increasing the sample size reduced Wasserstein distances and improved predictive accuracy, demonstrating that GPR is sample efficient while still benefiting from more data.

The Hénon map experiments also highlighted how measurement functions, delays, and embedding dimensions influence reconstruction quality. Smooth measurement functions required lower embedding dimensions and produced clear embeddings, while oscillatory functions created tangling that demanded higher dimensions. Similarly, small delays provided clean reconstructions, whereas larger delays distorted the geometry. These results provide practical guidelines for choosing embedding parameters when the underlying system is unknown, as in most real data.

When applied to ecological time series of *Salvinia molesta* biomass, the framework produced forecasts that followed both local fluctuations and overall trends. Even with noisy and irregular data, GPR models generalized beyond the training sets. At the same time, the reconstructions showed that fine-scale structures were smoothed, reflecting the limits of sparse field data. This balance between predictive accuracy and structural fidelity illustrates both the promise and the challenges of applying chaos-inspired methods to real-world ecological systems.

6 Conclusion and Future Work

We presented a study that combines delay embedding with Gaussian Process Regression to reconstruct and forecast chaotic dynamics. On synthetic maps, the approach recovered attractor geometry with relatively few samples, and on ecological data it captured meaningful predictive structure. These results suggest that tools from dynamical systems can provide insight into ecological processes when paired with modern machine learning.

There are several directions for future work. First, more systematic testing across a wider range of embedding dimensions, delays, and kernels could clarify the tradeoffs observed here. Second, extending the method to multivariate data, such as biomass together with environmental drivers, could provide a richer view of ecological dynamics. Third, exploring other machine learning models alongside GPR may help balance accuracy and efficiency for larger datasets. Finally, developing statistical tools to better detect chaos in ecological data would strengthen the link between theory and application.

In particular, for the *Salvinia molesta* dataset, it will be important to test a broader set of embedding dimensions Q and delays τ. The False Nearest Neighbors (FNN) test can be used to choose Q, while the Average Mutual Information (AMI) method can guide the selection of τ. Combining these approaches with predictive performance measures such as R^2 provides a path to identifying optimal embedding parameters. With these values in hand, Wasserstein distances between predicted and observed trajectories can be evaluated more systematically, giving a clearer view of how delay embedding and GPR capture ecological dynamics.

Overall, this study demonstrates that combining delay embedding with GPR is a promising way to bridge mathematical modeling and applied ecology. With further refinement, this framework may contribute to building early warning tools and improving the understanding of complex ecological dynamics.

References

[1] Schooler, S. S., Salau, B., Julien, M. H., & Ives, A. R. (2011). Alternative stable states explain unpredictable biological control of *Salvinia molesta* in Kakadu. *Nature*, 470, 86–89. doi:10.1038/nature09735.

[2] Day, S. & Storch, L. S. (2022). Takens delay variant distributions: Connecting time delay embedding with probability measures. *Physica D: Nonlinear Phenomena*, 436, 133325. doi:10.1016/j.physd.2022.133325.

[3] Day, S., Kalies, W. D., & Storch, L. S. (2021). State space reconstruction and Gaussian process regression as a flexible framework for forecasting chaotic dynamics. *Chaos*, 31(12), 123123. doi:10.1063/5.0067371.

[4] Sprott, J. C. (2003). *Period Doubling Cascades from Data*. Department of Physics, University of Wisconsin–Madison. Preprint.

[5] Ives, A. R. & Dakos, V. (2012). Detecting dynamical changes in nonlinear time series using locally linear state space models. *Ecosphere*, 3(6), 1–15. doi:10.1890/ES11-00347.1.

[6] Forgoston, E., Day, S. L., de Ruiter, P. C., Doelman, A., Hartemink, N., Hastings, A., Hemerik, L., Hening, A., Hofbauer, J., Kéfi, S., Kessler, D. A., Klauschies, T., Kuehn, C., Lim, X., Moore, J. C., Morriën, E., Neutel, A.-M., Pantel, J., Schreiber, S. J., Shaw, L. B., Shnerb, N., Siero, E., Storch, L. S., Thorne, M. A. S., van de Leemput, I., van Velzen, E., & Weinan, E. (2023). Stability and fluctuations in complex ecological systems. *arXiv preprint*, arXiv:2306.07447.

[7] Storch, L. S. & Day, S. L. (2022). Topological early warning signals: Quantifying varying routes to extinction in a spatially distributed population model. *Journal of Theoretical Biology*, 549, 111259. doi:10.1016/j.jtbi.2022.111259.

[8] Munch, S. B. (2017). A Bayesian perspective on nonlinear forecasting for ecology. *NOAA Technical Memorandum NMFS-SWFSC-582*, U.S. Department of Commerce.

[9] Krakovská, A., Melicherčík, M., and Budáčová, H. (2022). State space reconstruction techniques and the accuracy of prediction. *Communications in Nonlinear Science and Numerical Simulation*, 110, 106401. doi:10.1016/j.cnsns.2021.106401.

[10] Okuno, H., Hirata, Y., and Suzuki, H. (2020). Forecasting high-dimensional dynamics exploiting suboptimal embeddings. *Scientific Reports*, 10, 1017. doi:10.1038/s41598-019-57255-4.