PREDICTING UNOBSERVED DRIVER OF REGIME SHIFTS IN SOCIAL-ECOLOGICAL SYSTEMS WITH UNIVERSAL DYNAMIC EQUATIONS

A PREPRINT

Kunal J. Rathore*
College of Earth, Ocean, and Atmospheric Sciences,
Oregon State University, Corvallis, OR, USA
rathorek@oregonstate.edu

John H. Buckner
College of Earth, Ocean, and Atmospheric Sciences,
Oregon State University, Corvallis, OR, USA
bucknejo@oregonstate.edu

Zechariah D. Meunier
College of Earth, Ocean, and Atmospheric Sciences,
Oregon State University, Corvallis, OR, USA
meunierz@oregonstate.edu

Jorge Arroyo Esquivel
California Department of Fish and Wildlife,
West Sacramento, California, USA
arroyoesquiveljorge@gmail.com

James R. Watson
College of Earth, Ocean, and Atmospheric Sciences,
Oregon State University, Corvallis, OR, USA
james.watson@oregonstate.edu

December 13, 2025

ABSTRACT

Ecosystems around the world are anticipated to undergo regime shifts as temperatures rise and other climatic and anthropogenic perturbations erode the resilience of present-day states. Forecasting these nonlinear ecosystem dynamics can help stakeholders to prepare for the associated rapid changes. One major challenge is that regime shifts can be difficult to predict when they are driven by unobserved factors. In this paper, we advance scientific machine learning methods, specifically universal dynamic equations (UDEs), to identify changes in an unobserved bifurcation parameter and predict ecosystem regime shifts. We demonstrate this approach using simulated data created from a dynamic model of a species population experiencing loss due to unobserved extraction or harvest. This could be, for example, illegal fishing from a fishery or unreported poaching in a game reserve. We show that UDEs can accurately identify changes in the unobserved bifurcation parameter, in our case the slowly increasing harvest rate, and predict when a regime shift might occur. Compared to alternative forecasting methods, our UDE approach provides more reliable short-term predictions with fewer data. This approach provides a new set of methods for ecosystem stakeholders and managers to identify unobserved changes in key parameters that drive nonlinear change.

Keywords harvest rate, nonlinear dynamics, regime shift, population dynamics, neural networks, UDE · scientific ML

1 Introduction

Ecological forecasting can provide ecosystem stakeholders and policymakers vital actionable information on how to adapt to changing environmental conditions Dietze [2017a], Dietze et al. [2024]. Predictive models play a crucial role

*Use footnote for providing further information about author (webpage, alternative address)—*not* for acknowledging funding agencies.

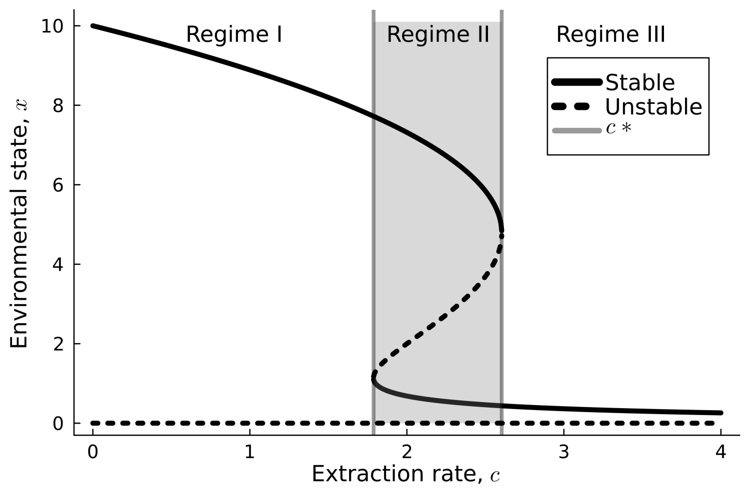

Figure 1: Bifurcation diagram illustrating the existence of three regimes in the deterministic dynamical system, where environmental state x is a function of the harvest rate c. Solid splines indicate stable states, and dotted splines show unstable states. Vertical lines mark the regime boundaries at the equilibrium harvest rates $c*$ with regime II (shaded box) exhibiting bistability. Modified from Tilman et al. [2024]

in ecological forecasting, as they can help us to understand and prepare for often large and abrupt changes in ecosystem states Clark et al. [2001], Kickert et al. [1999], Griffith and Fulton [2014]. Examples include predicting species population dynamics in response to climate change Urban et al. [2016], forecasting harmful algal blooms for effective water resource management Anderson et al. [2012], and estimating wildfire risk based on vegetation and atmospheric conditions Abatzoglou and Williams [2016]. However, one major challenge in ecological forecasting is accounting for unobserved drivers of change. This is a common problem in many scientific domains such as epidemiology, where asymptomatic carriers and unreported cases alter the dynamics of disease transmission Fraser et al. [2004]; economics, where latent variables such as consumer confidence impact macroeconomic outcomes Stock and Watson [2016]; and neuroscience, where unobserved neural states influence observed behavior and brain activity patterns Paninski et al. [2010]. This challenge is particularly critical when the unobserved quantity is a bifurcation parameter that governs rapid and large shifts in the qualitative properties of ecosystem dynamics.

Here, we consider populations that exhibit regime shifts caused by a hidden source of loss, such as illegal harvest from a fishery or poaching from a game reserve. In these systems, the harvest rate can be a bifurcation parameter May and Oster [1976] leading to unanticipated regime shifts. In such a scenario, an ecosystem with high abundance biomass (i.e., a stable regime, Fig. 1 Regime I) could pass through a transient regime (i.e., a flickering regime, Fig. 1 Regime II) to low biomass (i.e., the second stable regime, Fig. 1 Regime III). This presents an extreme challenge for ecosystem managers because these hidden drivers can push systems past critical thresholds without warning, often resulting in sudden ecosystem collapse that occurs too rapidly for conventional management interventions to prevent Scheffer et al. [2009], Biggs et al. [2009]. Forecasting methods that identify the unobserved parameters could improve forecasting accuracy, improving the management of ecosystems as they undergo rapid changes.

A wide variety of mathematical and computational methods have been developed to forecast ecosystem regime shifts, including approaches based on dynamical systems, machine learning, and statistical modeling Dakos et al. [2015], Ghadami and Epureanu [2022], Panahi et al. [2024]. Dynamical models are key tools for forecasting regime shifts, as they offer a framework for predicting critical transitions between stable states in ecosystems based on an explicit representation of key processes. These models use discrete-time difference and continuous-time differential equations to represent the underlying mechanisms driving the dynamics of the system, allowing researchers to identify early warning signals of tipping points and to improve our understanding of why these changes occur Barnosky et al. [2012], Kéfi et al. [2014]. By formally representing interactions among system components and their responses to external perturbations, dynamical models offer information on the conditions that precipitate abrupt changes, such as shifts in ecosystem states, climatic patterns, or disease outbreaks Biggs et al. [2015], May [1977].

Despite their relevance, these models have notable limitations when it comes to forecasting regime shifts. The parametrization process in dynamical models is crucial, but it can be challenging because multiple parameter sets can produce dynamics that are consistent with the data. In addition, the complexity of ecological systems characterized by numerous abiotic and biotic components, as well as direct and indirect interactions, introduces significant uncertainty

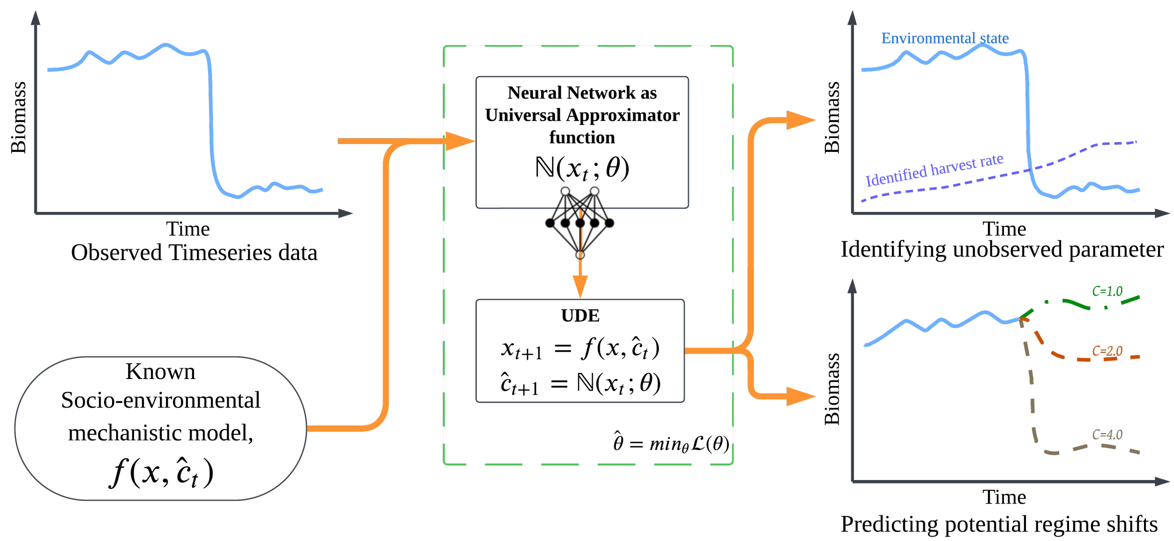

Figure 2: Workflow diagram showing the process to identify the unobserved parameter and predict regime shifts using a universal dynamic equation (UDE) model. The state-space UDE model can identify the unobserved (bifurcation) parameter and predict regime shifts in a simulated ecological population experiencing some form of harvest. We simulated 50 time-series as observed data and modeled the data using a neural network embedded in a mechanistic model of the known dynamics of the system. The UDE contains the neural network as a universal function approximator, and optimized the formulated loss function (green box). Then, the UDE model was used to estimate the unobserved harvest rate, which changes over time. Finally, we generated forecasts with different harvest values to identify potential regime shifts in the ecological population.

in the predictions of dynamical models Dietze [2017b]. Furthermore, existing parametric models used for ecological forecasting often assume predefined functional forms, which may not capture the full breadth of complexity and nonlinearity in ecosystems Clark et al. [2001]. These models are also prone to overfitting (as are most models), especially with sparse or high-dimensional data, and struggle with the uncertainty and stochastic nature of ecological processes Turchin [2013].

In contrast, data-driven "equation-free" methods are also commonly used to make ecological forecasts. They differ from equation-based approaches in that they do not rely on predefined functional forms or assumptions about the underlying processes. One such approach, empirical dynamical modeling (EDM), uses time series data to reconstruct and infer system dynamics without assuming explicit equations Sugihara et al. [2012], Perretti et al. [2013], Munch and Brias [2024]. In addition, numerous machine learning approaches have been developed to provide predictions of ecosystem dynamics. Advanced neural network (NN) architectures, such as the recurrent neural network (RNN) and long short-term memory (LSTM) models, have shown great potential in overcoming the limitations of traditional equation-based methods Hochreiter and Schmidhuber [1997]. However, the lack of mechanistic interpretability in these methods remains a key drawback that limits their broader applicability Lipton et al. [2015]. Similarly, by not having explicit functional forms, these equation-free approaches often overlook valuable mechanistic insights that could improve forecasts, particularly when data are limited or have observation biases Perretti and Munch [2015], Greff et al. [2017].

Recently, Rackauckas et al. [2020a], Bonnaffé et al. [2021], Arroyo-Esquivel et al. [2024], Buckner et al. [2024] demonstrated the potential of scientific machine learning (SciML) methods, which combine theoretical knowledge, often represented by differential and difference equations, with data-driven neural networks to model nonlinear ecosystem dynamics. Using mathematical representations of known interactions and dynamics, in conjunction with deep neural networks, SciML bridges the gap between equation-based mechanistic approaches and equation-free data-driven approaches, offering improved forecasting skills while maintaining interpretability. Building on this foundation, we have extended the application of SciML to address two key challenges in ecosystem management: (1) identifying unobserved parameters that influence ecosystem behavior, and (2) making accurate forecasts based on this information. We note that here we focus on social-ecological systems, but this application of SciML is generalizable to other complex systems with nonlinear dynamics driven by changes in an unobserved parameter.

In this study, we use a specific form of SciML called state-space universal dynamic equations (UDEs) Buckner et al. [2024], which can embed neural networks within difference equations that represent the discrete-time dynamics of an ecosystem. The neural networks are trained on noisy time series data using a state-space modeling framework. Here we extend this approach by using the neural network component of the UDEs to estimate the unobserved bifurcation parameter of a social-ecological system and forecast a regime shift. We test this approach on simulated data from a social-ecological model of a harvested population that exhibits flickering between alternative stable states and a regime change as the harvest rate increases past a critical threshold. We assessed the UDE's ability to classify dynamic regimes by treating the predicted unobserved harvest parameter as an indicator of regime shifts and evaluated prediction accuracy using transition thresholds. We also compared the model's forecasting skill with various modern-day ecosystem forecasting methods.

2 Materials & Methods

2.1 Modeling Social-Ecological Dynamics

We generated synthetic data following the social-ecological model introduced by Tilman et al. [2024]. Although this model is simple, it effectively represents the nonlinear dynamics of many social-environmental systems, accounting for environmental stochasticity, such as the collapse of fisheries due to overfishing or kelp forests affected by sea otter populations Nicholson et al. [2024]. In the model, x represents the abundance of the target species. Population growth is modeled as logistic growth with an intrinsic growth rate r and carrying capacity K, and multiplicative shocks to the growth rate. Animals are harvested from the population at a rate that depends on a nonlinear function of their abundance x_t, a parameter c_t that determines the intensity of harvest, and a parameter h, which determines the strength of nonlinearity. Population abundances update in discrete time as follows:

$$x_{t+1} = rx_t \left(1 - \frac{x_t}{K}\right) - c\left(\frac{x_t^2}{x_t^2 + h^2}\right) + (1 + i_t)x_t \tag{1}$$

Where, time-correlated red noise that models environmental shocks is represented by i_t, with T as the time scale over which noise becomes uncorrelated, and $\eta_t \sim \mathcal{N}(0,\, \beta^2)$ as an additive element of a series of independent and identically distributed normal errors. This type of noise characterizes scenarios where random environmental fluctuations are not completely independent over time or space, but instead exhibit a certain pattern of correlation, which is common in nature Allen [2010].

$$i_t = \left(1 - \frac{1}{T}\right) i_{t-1} + \eta_t \tag{2}$$

For example, Figure 1 illustrates that for low harvest rates c, the system settles on a single high-abundance equilibrium point. In contrast, at high values of c, the system moves to a lower abundance stable state. For intermediate c values, bistability occurs, and the presence of noise induces flickering dynamics, causing the system to frequently transition between the high and low abundance basins of attraction Tilman et al. [2024], Dakos et al. [2012].

To explore the ability of state-space UDEs to identify changes in an unobserved parameter (i.e., the harvest rate c) and to predict regime shifts in abundance, we generated 50 simulated time series y_t of 1000 length, simulating sequences of population abundance x_t and noise terms i_t from equation 1. Throughout the simulations, we increased the harvest intensity c_t linearly from $[0, 4]$, which caused the system to pass from the high abundance regime, through the bistable flickering regime, to the low abundance regime (Fig. 3). We used the set of base parameters for the social-ecological model discussed in Tilman et al. [2024] for each simulation (Table 2). We incorporated observation noise into the synthetic data by simulating data points y_t from a normal distribution truncated from below at zero with mean x_t and standard deviation 1.0.

2.2 Scientific Machine Learning of Regime Shifts

We developed a SciML method to identify the unobserved bifurcation parameter c, and to forecast state variables (i.e., abundance) into the future, and anticipate potential ecosystems regime shifts (see Fig. 2 for an overview). This framework uses state-space universal dynamic equations to forecast the dynamics (indexed with t) of a system with state variables x_t and predicts the slowly changing bifurcation parameter c_t. Importantly, we assume that we have a noisy set of observations y_t, which are equal to the observed state variable x_t plus the observation error $\varepsilon_t \sim \mathcal{N}(0, 1)$. Changes in the bifurcation parameter c_t are not observed directly.

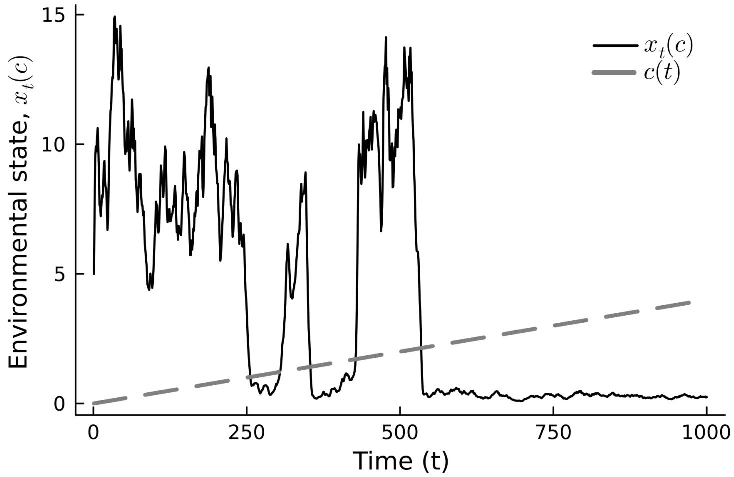

Figure 3: An example simulated time series of the environmental state (x_t) and harvest rate (c_t). The harvest rate is a monotonically increasing function with time, and its effects can be observed on the environmental state through the dynamical model of Tilman et al. [2024].

The dynamics of the system were modeled using estimates of the state variable $\hat{x}_t$. We assume that the dynamics of the observed variable can be described by a known function f. However, the dynamics of the unobserved parameter c_t are unknown. Therefore, we modeled changes in the estimates of unobserved state variable $\hat{c}$ with universal function approximator $\mathbb{N}$.

$$\hat{x}_{t+1} = f(\hat{x}_t, \hat{c}_t), \quad \hat{c}_{t+1} = \mathbb{N}(\hat{x}_t; \theta) \tag{3}$$

We used the known functional form of Tilman et al. [2024] to incorporate additional biological information to predict the change in system state (i.e., dynamics of x_t). We implemented an artificial neural network $\mathbb{N}(\hat{x}_t; \theta)$ as a universal function approximator with a fully connected five-layer architecture designed to capture complex nonlinear relationships. The network consists of a series of hidden layers, where the weights of each node in m-th layer are calculated from the nodes in $(m-1)$-th layer. Each node n undergoes a transformation through a nonlinear activation function a_n, followed by a weighted summation using parameters θ_{mn}. This process is repeated through all layers until the output layer is reached. The parameters θ are estimated by training on observed state variables and optimizing a composite likelihood function, enabling the network to accurately approximate the functional form of the dynamical system.

We formulated a composite loss function that consists of three components: dynamic loss L_{dyn}, observational loss L_{obs}, and the regularization term L_{reg}. Dynamic loss quantifies how well the model's predictions align with the expected outcomes based on the dynamical system model.

$$L_{dyn} = -\frac{1}{\Delta t \tau^2} \sum_{t=1}^{T_f - 1} \left(\left(\hat{c}_{t+1} - \mathbb{N}\left(\hat{x}_t; \theta\right) \right)^2 + \left(\hat{x}_{t+1} - f(\hat{x}_t, \hat{c}_t) \right)^2 \right) \tag{4}$$

The observational loss term measures how closely the estimated states match the data. This helps the model to capture the relationship between the input features and the observed outputs accurately.

$$L_{obs} = -\frac{1}{\sigma^2} \sum_{t=1}^{T_f} \left(y_t - \hat{x}_t \right)^2 \tag{5}$$

We applied L2 regularization, which penalizes large weights by summing squared weights across the network layers and discourages model overfitting.

$$L_{reg} = \sum_{t=l} \theta_l^2 \tag{6}$$

Finally, the objective function was defined as:

$$L_{total} = L_{dyn} + L_{obs} + \zeta L_{reg} \tag{7}$$

Here, ζ is the weight parameter to optimize the significance of regularization. In our case, experimental observations indicated that $\zeta = 0.4$ provides the best performance.

2.3 SciML Model Training

We applied the state-space UDE framework (eqs. 3- 7) to the simulated data from the social-ecological population model of Tilman et al. [2024]. We trained the models on the abundance time series y_t and environmental noise term i_t, and treated the harvest rate c_t as an unobserved variable. We used the social-ecological model (eq. 1) to describe changes in the estimated abundance

$$\hat{x}_{t+1} = r\hat{x}_t\left(1 - \frac{\hat{x}_t}{K}\right) - \hat{c}_t\left(\frac{\hat{x}_t^2}{\hat{x}_t^2 + h^2}\right) + (1 + \hat{i}_t)x_t \tag{8}$$

We modeled changes in the environmental noise term and the unobserved harvest rate using a neural network.

$$\hat{i}_{t+1} = \mathbb{N}_i(\hat{x}_t, \hat{i}_t, \hat{c}_t) \tag{9}$$

$$\hat{c}_{t+1} = \mathbb{N}_c(\hat{x}_t, \hat{i}_t, \hat{c}_t) \tag{10}$$

Using these equations, we simultaneously estimated the weights and biases of the neural network and the state variables $\hat{x}_t$, $\hat{i}_t$, and $\hat{c}_t$ using the combined loss function (eq. 7).

Our UDE model was trained by optimizing the likelihood function described in Equation 7, which serves as the objective function. In training we used Adam optimizer Kingma [2014], leveraging adaptive estimates of lower-order moments for efficient convergence in stochastic objective functions. The optimization process was implemented using the Optimizers.jl Ma et al. [2021] library in Julia, with a learning rate of 0.03, achieving a balanced trade-off between the optimization time and the performance of the model. The loss function gradient was efficiently calculated using automatic differentiation with the Zygote.jl Innes [2018] package, ensuring accurate and computationally efficient updates.

The activation functions, including ReLU Nair and Hinton [2010], CELU Barron [2017], softsign Glorot and Bengio [2010], and tanh LeCun et al. [2002], were evaluated on subsets of observed data to capture the best performance to predict the unobserved parameters. Among these, the CELU activation function provided the best prediction performance based on the root mean square error (RMSE) metric. During simulation, the model also predicted abundance values and environmental noise, which were further evaluated for forecast capability.

To effectively manage neural network parameters and predict changes in the unobserved harvest rate parameter, c_t, we used the Julia packages Lux.jl Pal [2023] and Optimization.jl Dixit and Rackauckas [2023]. Lux.jl facilitated the design of a scalable neural network architecture with explicit parameter handling, improving both flexibility and efficiency in model development.

We perform training and evaluation 50 times randomly to ensure robustness and assess the stability of model performance across different θ initializations and data permutations. This helps quantify uncertainty in parameter estimation and forecasting skills under stochastic conditions.

2.4 Simulation tests

After training the model on simulated data with optimized NN parameters, we made forecasts of future population abundances using the forward prediction process. Provided the current abundance of the ecological population x_t, we used trained UDE models to estimate future populations under different regimes, specifically high abundance (regime I, starting at time step 204), intermediate flickering with bistability (regime II, starting at time step 530), and low abundance (regime III, starting at time step 800). This approach allowed us to evaluate the predictive capabilities of the models in stable and transient regimes. Once optimized, a UDE model can be used to anticipate potential regime shifts by forecasting changes in abundance.

Our experiments included forecasting the states of the system and evaluating with observed data. In addition to the methods mentioned above, we also evaluated forecasts from several benchmark approaches to provide comprehensive performance comparisons. We included the auto-regressive integrated moving average (ARIMA) method Durbin and Koopman [2012] as it represents the classical statistical approach to time series forecasting and serves as a fundamental baseline in ecological applications. ARIMA models are particularly effective for capturing linear temporal dependencies and trends in stationary data, making them an essential comparison point for forecasting time series data. We also compared performance with the random-walk null model. Long short-term memory (LSTM) networks Hochreiter and Schmidhuber [1997] were selected to represent state-of-the-art deep learning approaches for sequential data modeling. LSTMs are specifically designed to capture long-term temporal dependencies through their gating mechanisms, addressing the vanishing-gradient problem that limits traditional recurrent networks. Their demonstrated success in modeling complex nonlinear patterns in time series data makes them a natural benchmark for evaluating our approach.

The predictive performances of the universal dynamic equation with a neural network (UDE-NN) and other comparative models were evaluated using the root mean square error (RMSE) of forecasted abundances. Considering the unobserved harvest rate parameter as an indicator of the system's current regime, we also evaluated and compared classification performance to identify the most effective prediction method. We used critical transition values $c*$ (shown in Fig. 1) from the simulated mechanistic model Tilman et al. [2024] as thresholds to categorize the simulated time series of x_t into three regimes. The categorized predicted values of the unobserved parameter were evaluated with the F-score—a harmonic mean of precision (measures how many of the predicted positive results were actually correct for each regime) and recall (measures how many of the actual positive cases were correctly identified for each regime), giving a balanced metric that considers both measures.

2.5 Gaussian Processes as Surrogate Method

In addition to using neural networks as universal function approximators, we explored alternative approximation methods within the state-space dynamic equation framework, including the Gaussian processes (GPs) approximation technique Olivier et al. [2021]. GPs have been extensively validated for differential equation problems similar to our UDE framework Medeiros et al. [2025], Raissi et al. [2019a, 2017], provide superior uncertainty quantification capabilities that are critical for scientific applications Yang et al. [2020], and offer the flexibility to incorporate physical constraints through specialized kernel design Solin and Särkkä [2020]. Furthermore, the theoretical equivalence between infinitely wide neural networks and Gaussian processes Lee et al. [2017] makes this comparison particularly meaningful for understanding the relative advantages of artificial neural networks and Bayesian methods in the scientific machine learning literature.

We incorporated GPs into the difference equation (DE) framework (for brevity, we use the acronym DE-GP for this method). This approach leverages the kernel-based framework of GPs to provide smooth, continuous, and locally adaptive representations while maintaining computational tractability Turner et al. [2010].

The system evolves through interactions between current-state variables and hidden variables. The state variables capture both the directly observed system conditions and the environmental noise influences. While the dynamics governing the observable state variables follow a known mechanistic form, the evolution of the latent components remains unknown and must be inferred from data. To model this hidden structure, we integrated GPs into the DE framework, allowing the Gaussian process component to learn the unknown latent dynamics from the available state information.

We used a likelihood function along with equations [4, 5], substituting the neural network with Gaussian processes as the function approximator to ensure a well-calibrated probabilistic model. The model parameters were optimized using a gradient-based technique. This formulation effectively combines physical system modeling and data-driven inference, enabling data-adaptable predictions and offering a probabilistic framework to assess uncertainty in complex systems with partially known dynamics. We implemented this approach using the AbstractGPs.jl Widmann et al. [2024] and KernelFunction.jl Galy-Fajou et al. [2024] packages in the Julia programming language.

3 Results

3.1 Unobserved parameter identification

Our universal dynamic equation framework successfully recovered the linear increasing trend in the unobserved harvest rate in the simulated examples (Fig. 4). The state-space UDE-NN method predicted the values of the unobserved parameters with the lowest error compared to alternative methods, evaluated using the RMSE metric.

As the statistical properties of data change in different regimes, we evaluated the prediction performance of the UDE-NN and DE-GP methods by training with different data sizes (N) in three different regimes. In Regime I (Fig. 5, N=204), the UDE-NN approach achieved significantly lower RMSE values, demonstrating effective learning with limited data. The UDE-NN approach also excelled in capturing unobserved parameters under flickering dynamics in Regime II (Fig. 5, N=530). Our approach also identified the unobserved parameter to outperform DE-GP in Regime III (Fig. 5, N=800),

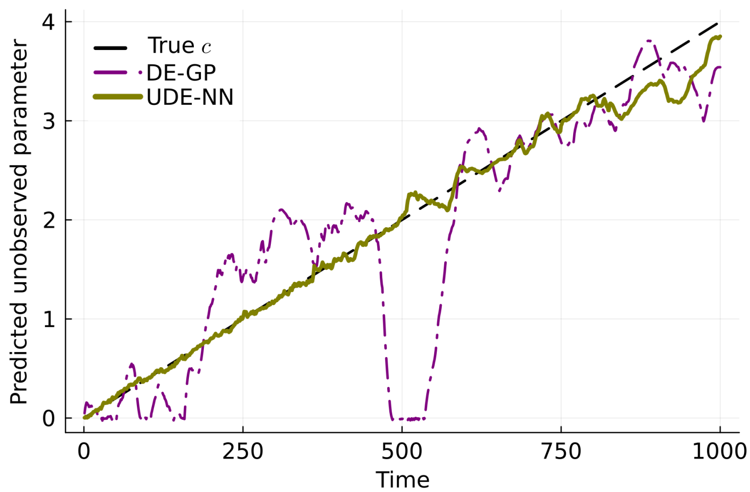

Figure 4: Identification of the unobserved parameter c_t from one of the random experiments. The actual parameter values (black) are compared with estimates from two methods: DE-Gaussian process (purple), and UDE-NN (olive).

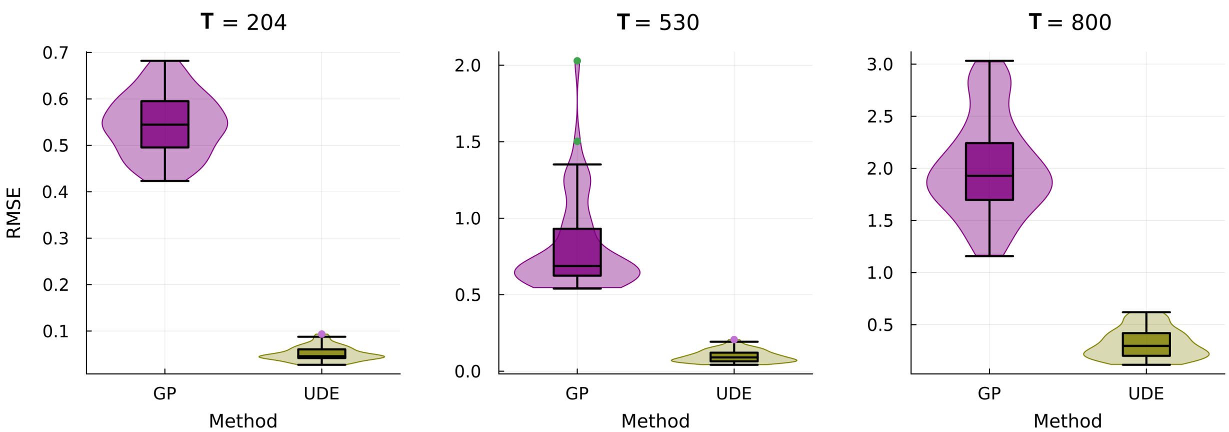

Figure 5: Violin-boxplots of root mean square error (RMSE) values for the predicted unobserved parameter $\hat{c}_t$ across different regimes over 50 random experiments. Results are given for high abundance regime I (N=204), intermediate flickering regime II (N=530), and low abundance regime III (N=800). The bottom and top of each box represent the first and third quartiles, respectively. The horizontal line inside the box indicates the median. Whiskers extend to the smallest and largest values within 1.5 times the interquartile range, and points beyond the whiskers represent outliers. The olive boxes represent the universal dynamic equation with a neural network (UDE-NN), and the purple boxes represent DE-Gaussian process (GP) models. Median RMSE values from UDE-NN are lowest across the regimes and methods. Regime I has relatively lower RMSE values compared to other regimes.

highlighting the effectiveness of neural networks in approximating nonlinear relationships. The Mann-Whitney U-test demonstrated a highly significant difference in RMSE performance between methods ($U = 0.0, p < 1e - 17$). The UDE-NN method consistently achieved lower RMSE values compared to DE-GP, and the test results showed complete separation between the two distributions.

In the classification context, the highest values of the F-score indicated that UDE-NN outperforms the DE-GP approximation technique in identifying the correct regime influenced by unobserved parameters (harvest rates) in the system (Fig. 6). This consistent prediction ability across different regimes highlights the robustness of UDE-NN in accurately identifying critical transitions and making reliable predictions, even in complex and noisy environments.

However, the distribution of F-scores varies substantially between ecological regimes, reflecting differences in classification ability under changing system dynamics (Fig. 6). In regime I (high abundance), the F-scores are tightly clustered, with the first quartile exceeding 0.75 and the upper whiskers approaching 1.0, indicating consistently high precision and recall. Regime II (flickering/bistable) exhibits moderate dispersion, with the first quartile around 0.75 and the upper

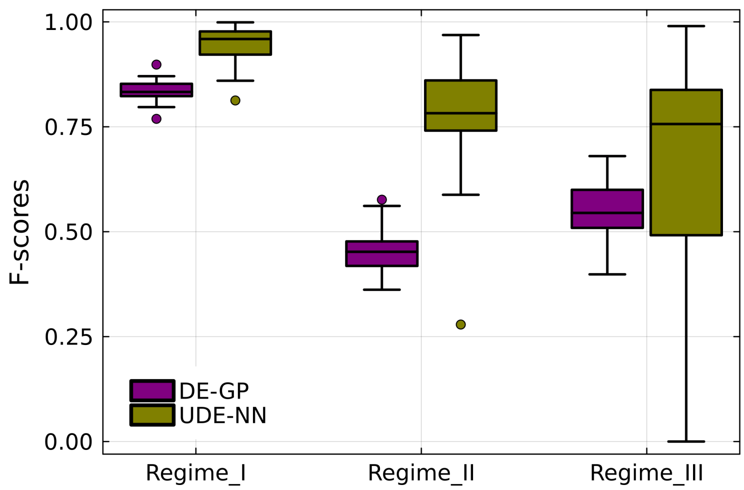

Figure 6: Classification-based predictions for the unobserved parameter $\hat{c}_t$, identifying the current categorical regime of the system (I, II, III) based on the predicted values. Boxplot shows F-score performance ranges (the higher, the better) in different regimes. Box color indicates methods, DE-Gaussian process (GP) in purple, and our method, UDE-NN in olive.

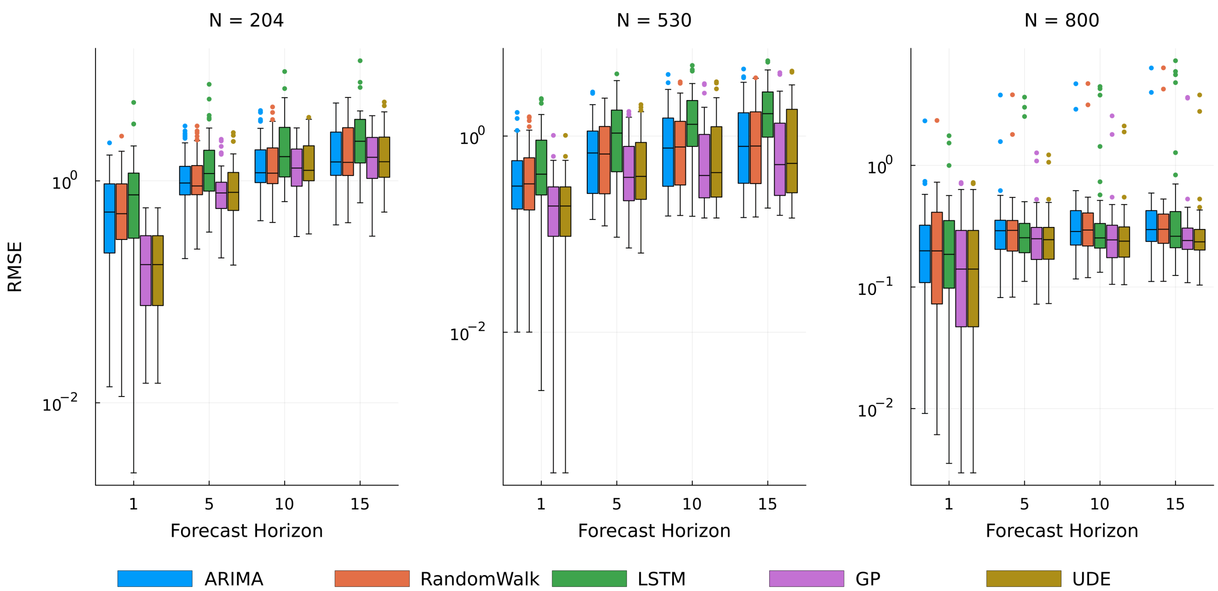

Figure 7: Comparative forecast performance using the root mean square error (RMSE) metric across different regimes and forecast horizons. Blue, orange, green and purple boxes represent ARIMA, RandomWalk, LSTM, and DE-Gaussian process methods, respectively, and olive boxes correspond to the UDE-NN method. Stacked boxplots of the performance metric after training on 204 data points (Regime I), 530 data points (Regime II), and 800 data points (Regime III).

values still reaching 1.0, suggesting reasonable classification confidence despite transient dynamics. In regime III (low abundance), the median F-score remains relatively high (~ 0.75), but the whiskers extend fully from 0 to 1, highlighting extreme variability and occasional misclassifications. This regime-dependent spread in F-score distributions underscores the challenge of making reliable predictions under nonlinear transitions and ecological degradation.

Table 1: Quantitative comparison of forecasts using root mean square error (RMSE). Median RMSE values are reported across different training scenarios based on the number of training data points, approximation techniques within the difference equation (DE) and universal dynamic equation (UDE) frameworks, and forecast horizons (1, 5, 10, or 15 time steps). The lowest RMSE values per training set are in bold.

Training datasize (N)	Method	RMSE-1	RMSE-5	RMSE-10	RMSE-15	Mean training time (s)
204	ARIMA	0.5234	0.9572	1.1842	1.4818	17.21
204	RandomWalk	0.5063	0.9003	**1.1733**	**1.4685**	1.05
204	LSTM	0.7487	1.1642	1.6540	2.2764	58.05
204	GP	**0.1760**	**0.7787**	1.3053	1.6314	533.11
204	UDE	**0.1760**	0.7886	1.2447	1.4871	246.34
530	ARIMA	0.3093	0.673	0.7556	0.7862	19.52
530	RandomWalk	0.3257	0.6564	0.7743	0.7918	1.13
530	LSTM	0.4090	1.0727	1.3187	1.6924	327.31
530	GP	**0.1938**	**0.3799**	**0.3965**	**0.5121**	3769.40
530	UDE	**0.1938**	0.3883	0.424	0.5269	617.21
800	ARIMA	0.1986	0.2903	0.2859	0.2968	20.73
800	RandomWalk	0.1982	0.2923	0.2941	0.2985	1.14
800	LSTM	0.1856	0.2541	0.2534	0.2614	785.05
800	GP	**0.1405**	0.2495	0.2442	0.2413	9330.91
800	UDE	**0.1405**	**0.2451**	**0.2383**	**0.2353**	904.33

3.2 Near-term forecasts of regime shifts

The forecast performance of the implemented approaches was evaluated for varying horizons (time-steps) {1, 5, 10, 15} using the RMSE metric (Fig. 7). Methods using known dynamics (UDE-NN and DE-GP) consistently showed the best performance values for a one-step-ahead forecast (horizon = 1) and slightly improved performance with lower median RMSE values for the forecast horizon of 5 time-steps (Fig. 7, N=530). Both methods showed relatively better performance than the alternative methods in the flickering regime II. Mann-Whitney U-test shows significant differences ($p < 1e-8$) in RMSE performance between methods that use known dynamics (UDE-NN and DE-GP) and those without known dynamics (ARIMA, Random-Walk, LSTM) for one-step forecasts under regime I and regime II. All the methods performed poorly in forecasting at longer horizons, which was expected given the highly stochastic nature of this system.

The median RMSE values of the experiment with more training data points (Table 1, N=800) indicate that the UDE-NN method performed better than alternative methods, while the DE-GP method prevailed in performance with intermediate data size (Table 1, N=530). Overall, the experiment suggested that methods using known dynamics outperform other methods in short-term forecasting capabilities, and have potential for improvement in long-term forecasting capabilities.

4 Discussion

Predicting regime shifts in systems where there are unobserved drivers remains a significant challenge in many domains Dakos et al. [2015], Hamilton [1989], Akbal [2024]. Our study demonstrates that UDEs (both based on neural networks and Gaussian processes) can effectively predict unobserved driving parameters in social-environmental systems experiencing regime shifts. The framework successfully estimated unobserved harvest rates from simulated abundance data alone while incorporating governing equations, achieving three key outcomes: 1) more accurate predictions of unobserved parameter changes (indicated by lower RMSE values), 2) successful regime identification based on predicted parameters (with high F-scores), and 3) competitive short-term forecasting performance compared to alternative approaches.

The UDE method excels when systems exhibit explicit environmental stochasticity, where random fluctuations and unreported external forces create challenges for traditional forecasting approaches. By incorporating known mechanistic dynamics, UDEs can distinguish between stochastic noise and meaningful signals, enabling more reliable predictions even with noisy observations. The approach is particularly powerful when systems contain unknown or unreported influencing parameters that are difficult to estimate through direct observation. These unobserved drivers can emerge from multiple sources: technological and measurement limitations including sensor failures and analytical methodology constraints Rocke et al. [2003], Decorte et al. [2024], microclimate heterogeneity creating fine-scale environmental variations that influence species interactions Mislan and Helmuth [2008], temporal scale mismatches where ecological processes operate over different timescales than monitoring programs Winkler et al. [2021], subsurface processes that remain largely unobservable Condon et al. [2021], legacy effects from historical disturbances that continue to influence contemporary ecosystem dynamics Krause et al. [2020], economic and logistical constraints that limit monitoring coverage Sparrow et al. [2020], and cryptic ecological processes such as undiscovered species interactions Bickford et al. [2007], Uusitalo et al. [2018]. These unobserved drivers can also emerge from illegal activities such as unreported fishing events, poaching, or other clandestine resource extraction activities that are deliberately concealed from monitoring systems. Crucially, when such hidden parameters significantly influence system regime transitions, UDEs can capture these effects within a mechanistic framework that maintains biological and physical realism.

The incorporation of known dynamics provides several advantages over purely data-driven approaches. Combining mechanistic models with data-driven approaches offers a powerful means of addressing regime shift predictions in dynamical systems. These mechanistic constraints help the neural network components learn system behavior faster and produce more reliable results by constraining the solution space to physically plausible outcomes. Furthermore, incorporating established ecological or physical principles enhances the interpretability of the model, allowing researchers and managers to understand not only what the model predicts, but also why those predictions make sense within known theoretical frameworks.

Our UDE approach builds on and extends the valuable foundations established by current methods for predicting unobserved parameters in complex systems. Empirical dynamic modeling Munch and Brias [2024] and statistical approaches have demonstrated significant success in identifying patterns in observed data and have been instrumental in advancing our understanding of complex system dynamics. Similarly, existing modeling approaches Agan et al. [2021], Átilla Colombo Ferreguetti et al. [2018] to estimating unobserved harvest or poaching rates have made important contributions to conservation science.

Sparse system identification methods such as SINDy Brunton et al. [2016] and non-parametric approaches including Gaussian process methods Alvarez et al. [2012] have contributed valuable techniques to infer system dynamics from limited observations. In parallel, scientific machine learning approaches have developed distinct paradigms for discovering and modeling unobserved dynamics in complex systems. Physics-informed neural networks (PINNs) have demonstrated success in incorporating known physical laws while learning unknown parameters from sparse data Raissi et al. [2019b], Karniadakis et al. [2021]. Neural ordinary differential equations (NODEs) Chen et al. [2018] provide powerful frameworks for learning continuous-time dynamics directly from observations, while universal differential equation approaches Rackauckas et al. [2020b] have shown promise in combining mechanistic knowledge with machine learning for scientific discovery. These methods have proven particularly effective either when the underlying system structure is well-characterized or when sufficient data are available to learn complex mappings.

However, these approaches face inherent challenges when dealing with systems in which unobserved driving parameters exhibit their own complex temporal dynamics. While traditional methods excel at pattern recognition and strategic optimization, they typically rely on sparse and potentially biased observation data, making it difficult to model the dynamic adversarial behaviors that characterize many social-environmental systems. Our UDE framework complements these existing approaches by explicitly modeling and predicting the temporal evolution of hidden drivers within a mechanistic context, offering a distinct methodological contribution that addresses scenarios where traditional approaches may be limited.

Recent advances in reservoir computing for extracting slowly time-varying parameters Tokuda and Katori [2024] and learning governing equations of unobserved states Grigorian et al. [2025] have demonstrated the feasibility of inferring hidden dynamics from time series data. Our UDE framework extends these contributions through three key advances: 1) explicit incorporation of unobserved parameters within mechanistic model structures, 2) direct integration of known mechanistic dynamics into the fast dynamics model, and 3) application to regime-shifting systems between stable equilibrium points. The latter is a more challenging inference problem than the chaotic dynamics typically studied, as stable phases provide fewer informative signals about underlying parameter changes.

The broader Scientific Machine Learning landscape has seen significant advances through symbolic regression Angelis et al. [2023], physics-informed neural networks (PINNs) Cuomo et al. [2022], and Gaussian process-based modeling Girard and Murray-Smith [2005]. Our UDE framework builds on these innovations while specifically targeting the challenge of predicting the temporal evolution of unobserved driving forces. The neural network approximators within UDEs leverage the universal approximation capabilities demonstrated by these previous works, extending their application to high-dimensional systems with multiple interacting variables where complex nonlinear relationships between observed and unobserved components can be captured within mechanistic frameworks.

Despite its effectiveness, the UDE approach faces limitations that define the boundaries of its application and highlight areas for future development. The method relies heavily on a deep understanding of process-based or mechanistic models, which may not always be readily available for complex systems. Researchers must have sufficient knowledge of the dynamics of the underlying system to specify appropriate governing equations, which limits applicability in domains where the mechanistic understanding remains incomplete. The approach also requires robust approximation techniques to effectively capture the intricacies of the nonlinear or chaotic dynamics characteristic of ecological and stochastic systems Buckner et al. [2024]. Careful handling of feedback mechanisms proves crucial for accurate modeling, as mismanagement of these interactions can lead to erroneous conclusions or unstable model behavior. Current computational limitations become particularly pronounced when scaling to high-dimensional datasets with hundreds of interacting components, where complexity and computational demands increase significantly.

While our results show that UDEs do not consistently outperform alternative approaches in short-term forecast accuracy, this finding aligns with recent insights on the limitations of forecast-skill-based model selection. Boettiger [2022] demonstrates that selecting models based solely on forecast accuracy can create a 'forecast trap', where seemingly superior statistical performance leads to worse real-world management outcomes. This occurs because forecast accuracy alone fails to capture the non-uniqueness of models and their differential impacts on management objectives. Our UDE framework addresses this challenge by prioritizing mechanistic understanding and parameter interpretability over pure forecast performance. By explicitly modeling unobserved driving parameters within established ecological frameworks, UDEs provide actionable insights into the causal mechanisms underlying regime shifts, information that may be more valuable for management intervention than marginally-improved forecast accuracy. This approach aligns with recommendations to promote broader sets of models rather than selecting based on forecast skill alone, emphasizing the importance of mechanistic insight for robust decision-making in social-environmental systems.

Future research should prioritize improving long-term forecasting ability, particularly for systems with regime shifts or strong stochastic dynamics, where prediction accuracy typically degrades over extended time horizons. Developing more efficient algorithms, advanced neural network architectures, and high-performance computing strategies will be essential for scaling to complex real-world systems. Furthermore, incorporating uncertainty quantification methods would enhance the reliability of the framework and provide confidence limits for management decisions based on the mechanistic foundation that distinguishes UDEs from purely statistical approaches.

The UDE approach shows potential for addressing conservation challenges, though validation with real-world data would be needed. In wildlife conservation, the method could potentially improve population assessments by estimating illegal harvest rates from monitoring data. This capability might help wildlife managers better understand the true impact of poaching on population dynamics and assess the effectiveness of existing conservation interventions.

For broader ecosystem management, the framework offers pathways toward early warning systems for regime shift detection across multiple domains. The approach can be extended to predict regime shifts in other complex systems, including political transitions, Earth system dynamics, and epidemiological patterns. Political regime shifts, often driven by economic instability, social movements, or institutional breakdowns, exhibit nonlinear dynamics that could be captured using hybrid mechanistic data-driven models. Similarly, regime shifts in the Earth system such as those quantified in the planetary boundaries framework Rockström et al. [2009], Steffen et al. [2015] and research on cascading tipping points Rocha et al. [2018]—highlight the interconnected nature of environmental thresholds. Furthermore, epidemiological applications Brauer et al. [2019], such as predicting the spread of infectious diseases, could benefit from this approach by integrating mechanistic models of disease transmission with real-time data assimilation.

However, realizing these applications requires several key advances. Scaling computational efficiency for high-dimensional systems remains critical for handling the complexity of real-world scenarios. Developing robust methods for uncertainty quantification would also provide essential confidence measures for management decisions.

This work provides a new method for quantifying hidden dynamics in complex systems by directly predicting changes in unobserved parameters that drive regime shifts. This capability provides actionable information for management interventions and would be of great potential use for effective governance of social-ecological systems. The ability to predict not only what will happen but also the causal mechanisms of change can mean the difference between proactive management and reactive crisis response. As global systems face unprecedented pressures and regime shifts become increasingly common, tools that can peer into hidden mechanisms of change offer hope for more effective stewardship of social-ecological systems. The UDE framework has the potential to provide policy makers and resource managers with improved capabilities to design optimal control strategies. Future efforts could focus on scaling these methods to high-dimensional systems and validating their performance across diverse real-world applications. Ultimately, through new methods like SciML, we can improve our ability to predict and prepare for critical transitions in complex systems.

Appendix

A Mechanistic model parameters and data pre-processing

We simulated the Tilman et al. [2024] model to generate time series data with different random initializations, by solving the discrete state-space equation 1. We considered the harvest rate as an unobserved parameter that monotonically increases between the suggested range [0,4] in Tilman et al. [2024]. This process provides time-series observations $\hat{x}, \hat{i}$ corresponding to the environmental state and the term of noise in the system. The parameter values are listed in table 2.

Table 2: Tilman et al. [2024] mechanistic model parameters

Variable	Value	Description
x_t	0-20	Current environmental state (x_0 = 10.0)
i_t		Auto-correlated red noise
T	30	Timescale over which noise becomes uncorrelated
η_t	0	i.i.d. normal error term
β	0.07	The standard deviation of η's
r	1	Resource growth rate
K	10	Resource carrying capacity
c_t	.0-4.0	harvest rate
h	1	harvest half-saturation constant

We used trained models to forecast changes in population abundances as a function of different harvest rates. Figure 8. illustrates the predicted abundances, and the different colored lines corresponding to different harvest rate scenarios. The results highlight the potential for recovery of the system with lower harvest rates [figure 8. Panel D.] regardless of the current regime. It also highlights that adopting lower harvest rates can prevent irreversible collapse and enable recovery, whereas higher harvest rates drive the system toward the unstable regime II or the low-abundance regime III.

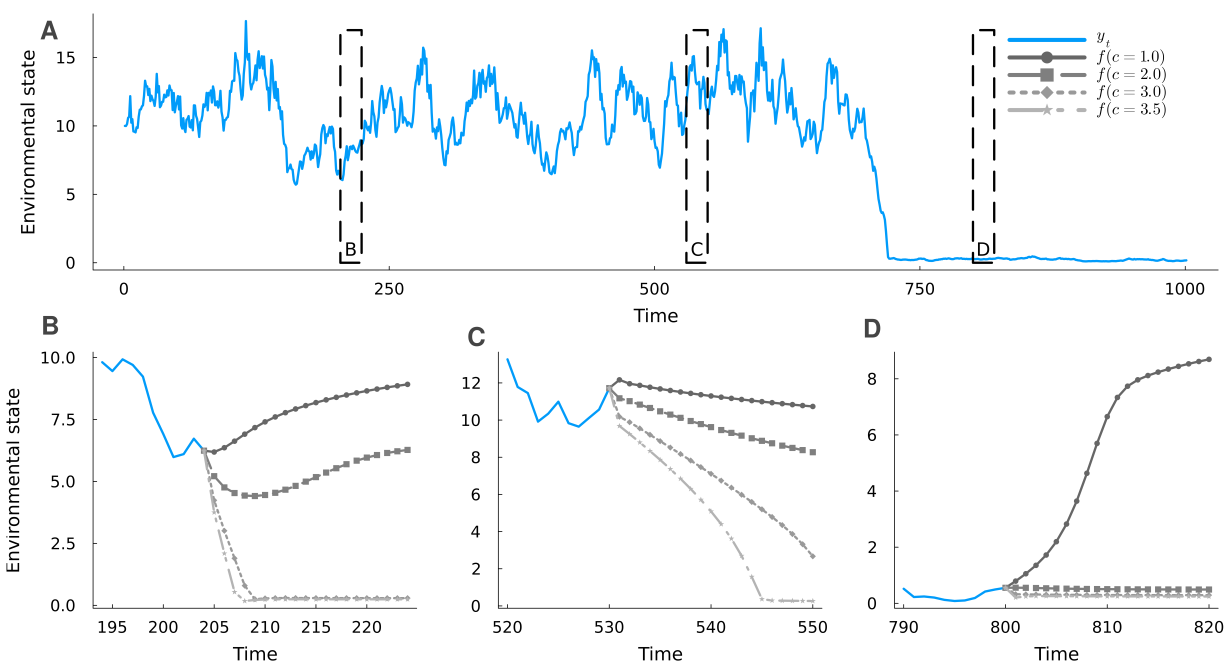

Figure 8: Forecasts of population abundances under different harvest rates c in each regime. An example time series of a noisy environmental state y_t simulated for 1000 time steps (A). The dashed boxes show insets for regime I forecasts (B), regime II forecasts (C), and regime III forecasts (D) under increasing harvest rates.

Figure 9 shows the comparisons of the forecast uncertainty between different methods. The UDE-NN model, with an embedded neural network, demonstrates a narrower confidence interval, indicating reduced uncertainty in forecasting.

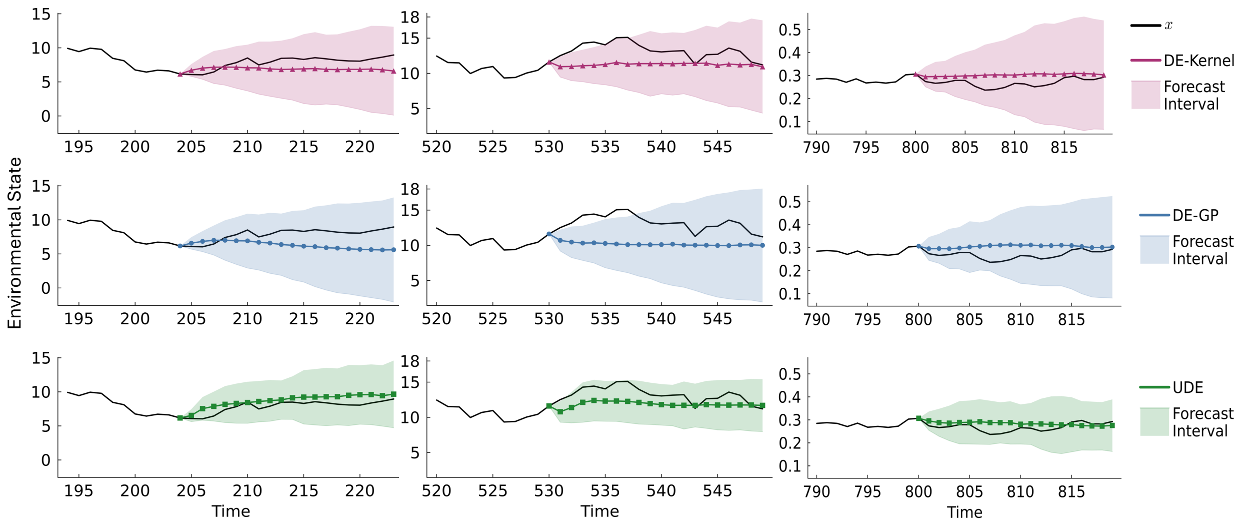

Figure 9: Forecasts and confidence intervals for the environmental state x across different regimes using comparative methods. Column 1 displays predictions for regime I, column 2 for regime II, and column 3 for regime III. Average forecast values are represented as points and lines, with confidence intervals given as shaded regions. The methods are Kernel-based parameter approximation, Gaussian process (GP) augmentation, and universal dynamic equation (UDE). Note the different y-axis scales per column.

In contrast, other DE methods exhibit increasing uncertainty as the forecast horizon is extended. Additionally, the results suggest that the algorithms had more uncertainty in Regime II, which is the system's flickering regime.

Data Availability

The code used to generate synthetic data, models, and their analysis is available in the Github repository `https://github.com/kjrathore/C_Star` .

Author contributions

Conceptualization: KJR, JHB, JRW, JAE, ZDM; Software: KJR, JHB, JAE, ZDM; Formal Analysis: KJR, JHB; Original Writing Draft: KJR, JHB; Writing Review and Editing: JHB, ZDM, JAE, JRW; Funding Acquisition: JAE, JRW Supervision: JRW; Project administration: JRW

Funding

This research was supported by the National Science Foundation awards #2233982 and #2233983 to JRW and LCM on Model-Enabled Machine Learning to Predict Ecosystem Regime Shifts.

Acknowledgments

This paper is a product of the model-enabled machine learning for ecology working group, which includes the authors of the paper, Lisa McManus, Ariel Greiner, Nathan Fitzpatrick, Cheyenne Jarman, and Emerson Arehart, all of whom provided valuable contributions to the intellectual environment that led to this paper. We also thank Chris Rackauckas for help working with Julia Scientific Machine Learning tools and the Hawai'i Institute of Marine Biology for hosting a workshop where the ideas for this project were developed.

References

Michael Dietze. Ecological forecasting. In *Ecological Forecasting*. Princeton University Press, 2017a.

Michael Dietze, Ethan P White, Antoinette Abeyta, Carl Boettiger, Nievita Bueno Watts, Cayelan C Carey, Rebecca Chaplin-Kramer, Ryan E Emanuel, SK Morgan Ernest, Renato J Figueiredo, et al. Near-term ecological forecasting for climate change action. *Nature Climate Change*, pages 1–9, 2024.

James S Clark, Steven R Carpenter, Mary Barber, Scott Collins, Andy Dobson, Jonathan A Foley, David M Lodge, Mercedes Pascual, Roger Pielke Jr, William Pizer, et al. Ecological forecasts: an emerging imperative. *science*, 293 (5530):657–660, 2001.

Ronald N Kickert, Giorgio Tonella, Alexander Simonov, and Sagar V Krupa. Predictive modeling of effects under global change. *Environmental Pollution*, 100(1):87–132, 1999.

Gary P Griffith and Elizabeth A Fulton. New approaches to simulating the complex interaction effects of multiple human impacts on the marine environment. *ICES Journal of Marine Science*, 71(4):764–774, 2014.

Mark C Urban, Greta Bocedi, Andrew P Hendry, J-B Mihoub, Guy Pe'er, Alex Singer, JR Bridle, LG Crozier, Luc De Meester, William Godsoe, et al. Improving the forecast for biodiversity under climate change. *Science*, 353(6304): aad8466, 2016.

Donald M Anderson, Allan D Cembella, and Gustaaf M Hallegraeff. Progress in understanding harmful algal blooms: paradigm shifts and new technologies for research, monitoring, and management. *Annual review of marine science*, 4:143–176, 2012.

John T Abatzoglou and A Park Williams. Impact of anthropogenic climate change on wildfire across western us forests. *Proceedings of the National Academy of sciences*, 113(42):11770–11775, 2016.

Christophe Fraser, Steven Riley, Roy M Anderson, and Neil M Ferguson. Factors that make an infectious disease outbreak controllable. *Proceedings of the National Academy of Sciences*, 101(16):6146–6151, 2004.

James H Stock and Mark W Watson. Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics. *Handbook of Macroeconomics*, 2:415–525, 2016.

Liam Paninski, Yashar Ahmadian, Daniel G Ferreira, Shinya Koyama, Kamiar Rahnama Rad, Michael Vidne, Joshua Vogelstein, and Wei Wu. A new look at state-space models for neural data. *Journal of Computational Neuroscience*, 29(1):107–126, 2010.

Andrew R. Tilman, Elisabeth H. Krueger, Lisa C. McManus, and James R. Watson. Maintaining human wellbeing as socio-environmental systems undergo regime shifts. *Ecological Economics*, 221:108194, 7 2024. ISSN 0921-8009. doi:10.1016/J.ECOLECON.2024.108194.

Robert M May and George F Oster. Bifurcations and dynamic complexity in simple ecological models. *The American Naturalist*, 110(974):573–599, 1976.

Marten Scheffer, Jordi Bascompte, William A Brock, Victor Brovkin, Stephen R Carpenter, Vasilis Dakos, Hermann Held, Egbert H Van Nes, Max Rietkerk, and George Sugihara. Early-warning signals for critical transitions. *Nature*, 461(7260):53–59, 2009.

Reinette Biggs, Stephen R Carpenter, and William A Brock. Turning back from the brink: Detecting an impending regime shift in time to avert it. *Proceedings of the National Academy of Sciences*, 106(3):826–831, 2009.

Vasilis Dakos, Stephen R Carpenter, Egbert H van Nes, and Marten Scheffer. Resilience indicators: prospects and limitations for early warnings of regime shifts. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 370(1659):20130263, 2015.

Amin Ghadami and Bogdan I Epureanu. Data-driven prediction in dynamical systems: recent developments. *Philosophical Transactions of the Royal Society A*, 380(2229):20210213, 2022.

Shirin Panahi, Ling-Wei Kong, Mohammadamin Moradi, Zheng-Meng Zhai, Bryan Glaz, Mulugeta Haile, and Ying-Cheng Lai. Machine learning prediction of tipping in complex dynamical systems. *Phys. Rev. Res.*, 6: 043194, November 2024. doi:10.1103/PhysRevResearch.6.043194. URL `https://link.aps.org/doi/10.1103/PhysRevResearch.6.043194`.

Anthony D Barnosky, Elizabeth A Hadly, Jordi Bascompte, Eric L Berlow, James H Brown, Mikael Fortelius, Wayne M Getz, John Harte, Alan Hastings, Pablo A Marquet, et al. Approaching a state shift in earth's biosphere. *Nature*, 486 (7401):52–58, 2012.

Sonia Kéfi, Vishwesha Guttal, William A Brock, Stephen R Carpenter, Aaron M Ellison, Valerie N Livina, David A Seekell, Marten Scheffer, Egbert H Van Nes, and Vasilis Dakos. Early warning signals of ecological transitions: methods for spatial patterns. *PloS one*, 9(3):e92097, 2014.

Reinette Biggs, Maja Schlüter, and Michael L Schoon. Principles for building resilience: sustaining ecosystem services in social-ecological systems. 2015.

Robert M May. Thresholds and breakpoints in ecosystems with a multiplicity of stable states. *Nature*, 269(5628): 471–477, 1977.

Michael C. Dietze. *Ecological Forecasting*. Princeton University Press, 2017b. ISBN 9780691160573. URL http://www.jstor.org/stable/j.ctvc7796h.

Peter Turchin. *Complex population dynamics: a theoretical/empirical synthesis (MPB-35)*. Princeton university press, 2013.

George Sugihara, Robert May, Hao Ye, Chih hao Hsieh, Ethan Deyle, Michael Fogarty, and Stephan Munch. Detecting causality in complex ecosystems. *Science*, 338(6106):496–500, 2012. doi:10.1126/science.1227079. URL https://www.science.org/doi/abs/10.1126/science.1227079.

Charles T Perretti, George Sugihara, and Stephan B Munch. Nonparametric forecasting outperforms parametric methods for a simulated multispecies system. *Ecology*, 94(4):794–800, 2013.

Stephan B. Munch and Antoine Brias. Empirical dynamic programming for model-free ecosystem-based management. *Methods in Ecology and Evolution*, 15(4):769–778, February 2024. ISSN 2041-210X. doi:10.1111/2041-210x.14302. URL http://dx.doi.org/10.1111/2041-210X.14302.

Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9:1735–1780, 11 1997. doi:10.1162/neco.1997.9.8.1735.

Zachary C. Lipton, John Berkowitz, and Charles Elkan. A critical review of recurrent neural networks for sequence learning, 2015. URL https://arxiv.org/abs/1506.00019.

Charles T Perretti and Stephan B Munch. On estimating the reliability of ecological forecasts. *Journal of Theoretical Biology*, 372:211–216, 2015.

Klaus Greff, Rupesh K. Srivastava, Jan Koutnik, Bas R. Steunebrink, and Jurgen Schmidhuber. Lstm: A search space odyssey. *IEEE Transactions on Neural Networks and Learning Systems*, 28(10):2222–2232, 10 2017. ISSN 2162-2388. doi:10.1109/tnnls.2016.2582924. URL http://dx.doi.org/10.1109/TNNLS.2016.2582924.

Christopher Rackauckas, Yingbo Ma, Julius Martensen, Collin Warner, Kirill Zubov, Rohit Supekar, Dominic Skinner, Ali Ramadhan, and Alan Edelman. Universal differential equations for scientific machine learning. 1 2020a. URL http://arxiv.org/abs/2001.04385.

Willem Bonnaffé, Ben C Sheldon, and Tim Coulson. Neural ordinary differential equations for ecological and evolutionary time-series analysis. *Methods in Ecology and Evolution*, 12(7):1301–1315, 2021.

Jorge Arroyo-Esquivel, Christopher A Klausmeier, and Elena Litchman. Using neural ordinary differential equations to predict complex ecological dynamics from population density data. *Journal of the Royal Society Interface*, 21(214): 20230604, 2024.

Jack Buckner, Zechariah Meunier, Jorge Arroyo Esquivel, Nathan Fitzpatrick, Ariel Greiner, Lisa McManus, and James Watson. Recovering complex ecological dynamics from time series using state-space universal dynamic equations, 10 2024.

Teri E. Nicholson, Loren McClenachan, Kisei R. Tanaka, and Kyle S. Van Houtan. Sea otter recovery buffers century-scale declines in California kelp forests. *PLOS Climate*, 3(1):1–18, January 2024. doi:10.1371/journal.pclm.0000290. URL https://doi.org/10.1371/journal.pclm.0000290. Publisher: Public Library of Science.

Linda JS Allen. *An introduction to stochastic processes with applications to biology*. CRC press, 2010.

Vasilis Dakos, Stephen R. Carpenter, William A. Brock, Aaron M. Ellison, Vishwesha Guttal, Anthony R. Ives, Sonia Kéfi, Valerie Livina, David A. Seekell, Egbert H. van Nes, and Marten Scheffer. Methods for detecting early warnings of critical transitions in time series illustrated using simulated ecological data. *PLoS ONE*, 7, 7 2012. ISSN 19326203. doi:10.1371/journal.pone.0041010.

Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Yingbo Ma, Vaibhav Dixit, Michael J Innes, Xingjian Guo, and Chris Rackauckas. A comparison of automatic differentiation and continuous sensitivity analysis for derivatives of differential equation solutions. In *2021 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–9. IEEE, 2021.

Mike Innes. Don't unroll adjoint: Differentiating ssa-form programs. *arXiv preprint arXiv:1810.07951*, 2018.

Vinod Nair and Geoffrey E. Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ICML'10, pages 807–814, Madison, WI, USA, June 2010. Omnipress. ISBN 978-1-60558-907-7.

Jonathan T. Barron. Continuously Differentiable Exponential Linear Units, 4 2017. URL http://arxiv.org/abs/1704.07483. arXiv:1704.07483.

Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings, 2010.

Yann LeCun, Léon Bottou, Genevieve B Orr, and Klaus-Robert Müller. Efficient backprop. In *Neural networks: Tricks of the trade*, pages 9–50. Springer, 2002.

Avik Pal. Lux: Explicit Parameterization of Deep Neural Networks in Julia, April 2023. URL https://doi.org/10.5281/zenodo.7808904. If you use this software, please cite it as below.

Vaibhav Kumar Dixit and Christopher Rackauckas. Optimization.jl: A unified optimization package, March 2023. URL https://doi.org/10.5281/zenodo.7738525.

James Durbin and Siem Jan Koopman. *Time series analysis by state space methods*, volume 38. OUP Oxford, 2012.

Audrey Olivier, Michael D. Shields, and Lori Graham-Brady. Bayesian neural networks for uncertainty quantification in data-driven materials modeling. *Computer Methods in Applied Mechanics and Engineering*, 386:114079, December 2021. ISSN 0045-7825. doi:10.1016/j.cma.2021.114079. URL https://www.sciencedirect.com/science/article/pii/S0045782521004102.

Lucas P. Medeiros, Darian K. Sorenson, Bethany J. Johnson, Eric P. Palkovacs, and Stephan B. Munch. Revealing unseen dynamical regimes of ecosystems from population time-series data. *Proceedings of the National Academy of Sciences*, 122(24):e2416637122, 2025. doi:10.1073/pnas.2416637122. URL https://www.pnas.org/doi/abs/10.1073/pnas.2416637122.

M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019a. ISSN 0021-9991. doi:https://doi.org/10.1016/j.jcp.2018.10.045. URL https://www.sciencedirect.com/science/article/pii/S0021999118307125.

Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Machine learning of linear differential equations using gaussian processes. *J. Comput. Phys.*, 348:683–693, November 2017. ISSN 0021-9991. doi:10.1016/j.jcp.2017.07.050. URL https://doi.org/10.1016/j.jcp.2017.07.050.

Liu Yang, Dongkun Zhang, and George Em Karniadakis. Physics-informed generative adversarial networks for stochastic differential equations. *SIAM Journal on Scientific Computing*, 42(1):A292–A317, 2020. doi:10.1137/18M1225409. URL https://doi.org/10.1137/18M1225409.

Arno Solin and Simo Särkkä. Hilbert space methods for reduced-rank gaussian process regression. *Statistics and Computing*, 30(2):419–446, 2020.

Jaehoon Lee, Yasaman Bahri, Roman Novak, Samuel S. Schoenholz, Jeffrey Pennington, and Jascha Narain Sohl-Dickstein. Deep neural networks as gaussian processes. *ArXiv*, abs/1711.00165, 2017. URL https://api.semanticscholar.org/CorpusID:3708505.

Ryan Turner, Marc Deisenroth, and Carl Rasmussen. State-space inference and learning with gaussian processes. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 868–875. JMLR Workshop and Conference Proceedings, 2010.

David Widmann, Will Tebbutt, st, Théo Galy-Fajou, Sharan Yalburgi, Hong Ge, Simone Carlo Surace, david vicente, Jacob Vaverka, John Skovbekk, Nathanael Bosch, Niklas Schmitz, Ross Viljoen, Tom Wright, Vikram, and andreaskoher. Juliagaussianprocesses/abstractgps.jl: v0.5.21, March 2024. URL https://doi.org/10.5281/zenodo.10801764.

Théo Galy-Fajou, David Widmann, Sharan Yalburgi, Will Tebbutt, st, Isak Falk, Simone Carlo Surace, Steffen Ridderbusch, Tom Wright, Hong Ge, Sebastian Khan, Pietro Monticone, Letif Mones, david vicente, Johannes Giersdorf, Julia TagBot, Ross Viljoen, Simon Schölly, Tor Erlend Fjelde, and Kaan Öcal. JuliaGaussianProcesses/KernelFunctions.jl: v0.10.64, July 2024. URL https://doi.org/10.5281/zenodo.12793559.

James D Hamilton. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica: Journal of the econometric society*, pages 357–384, 1989.

Omer Faruk Akbal. Regime-switching factor models and nowcasting with big data. *IMF Working Papers*, 2024(190): A001, 2024. doi:10.5089/9798400286407.001.A001. URL https://www.elibrary.imf.org/view/journals/001/2024/190/article-A001-en.xml.

David M. Rocke, Blythe Durbin, Machelle Wilson, and Henry D. Kahn. Modeling uncertainty in the measurement of low-level analytes in environmental analysis. *Ecotoxicology and Environmental Safety*, 56(1):78–92, 2003. ISSN 0147-6513. doi:https://doi.org/10.1016/S0147-6513(03)00052-6. URL https://www.sciencedirect.com/science/article/pii/S0147651303000526.

Thomas Decorte, Steven Mortier, Jonas J Lembrechts, Filip JR Meysman, Steven Latré, Erik Mannens, and Tim Verdonck. Missing value imputation of wireless sensor data for environmental monitoring. *Sensors*, 24(8):2416, 2024.

K. A. S. Mislan and B. Helmuth. Microclimate. In Sven Erik Jørgensen and Brian D. Fath, editors, *Encyclopedia of Ecology*, pages 2389–2393. Academic Press, Oxford, January 2008. ISBN 978-0-08-045405-4. doi:10.1016/B978-008045405-4.00520-6. URL https://www.sciencedirect.com/science/article/pii/B9780080454054005206.

Klara J. Winkler, Marie C. Dade, and Jesse T. Rieb. Mismatches in the Ecosystem Services Literature—a Review of Spatial, Temporal, and Functional-Conceptual Mismatches. *Current Landscape Ecology Reports*, 6(2):23–34, June 2021. ISSN 2364-494X. doi:10.1007/s40823-021-00063-2. URL https://doi.org/10.1007/s40823-021-00063-2.

Laura E Condon, Stefan Kollet, Marc FP Bierkens, Graham E Fogg, Reed M Maxwell, Mary C Hill, Harrie-Jan Hendricks Fransen, Anne Verhoef, Anne F Van Loon, Mauro Sulis, et al. Global groundwater modeling and monitoring: Opportunities and challenges. *Water Resources Research*, 57(12):e2020WR029500, 2021.

Andreas Krause, Almut Arneth, P Anthoni, and A Rammig. Legacy effects from historical environmental changes dominate future terrestrial carbon uptake. *Earth's Future*, 8(10):e2020EF001674, 2020.

Ben D. Sparrow, Will Edwards, Samantha E.M. Munroe, Glenda M. Wardle, Greg R. Guerin, Jean-Francois Bastin, Beryl Morris, Rebekah Christensen, Stuart Phinn, and Andrew J. Lowe. Effective ecosystem monitoring requires a multi-scaled approach. *Biological Reviews of the Cambridge Philosophical Society*, 95(6):1706–1719, December 2020. ISSN 1464-7931. doi:10.1111/brv.12636. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7689690/.

David Bickford, David J. Lohman, Navjot S. Sodhi, Peter K. L. Ng, Rudolf Meier, Kevin Winker, Krista K. Ingram, and Indraneil Das. Cryptic species as a window on diversity and conservation. *Trends in Ecology & Evolution*, 22(3): 148–155, March 2007. ISSN 0169-5347. doi:10.1016/j.tree.2006.11.004. URL https://www.sciencedirect.com/science/article/pii/S0169534706003703.

Laura Uusitalo, Maciej T. Tomczak, Bärbel Müller-Karulis, Ivars Putnis, Neda Trifonova, and Allan Tucker. Hidden variables in a Dynamic Bayesian Network identify ecosystem level change. *Ecological Informatics*, 45:9–15, May 2018. ISSN 1574-9541. doi:10.1016/j.ecoinf.2018.03.003. URL https://www.sciencedirect.com/science/article/pii/S1574954117301747.

Suzanne W. Agan, Adrian Treves, and Lisabeth L. Willey. Estimating poaching risk for the critically endangered wild red wolf (canis rufus). *PLOS ONE*, 16(5):1–20, 05 2021. doi:10.1371/journal.pone.0244261. URL https://doi.org/10.1371/journal.pone.0244261.

Átilla Colombo Ferreguetti, Juliane Pereira-Ribeiro, Jayme Augusto Prevedello, Walfrido Moraes Tomás, Carlos Frederico Duarte Rocha, and Helena Godoy Bergallo. One step ahead to predict potential poaching hotspots: Modeling occupancy and detectability of poachers in a neotropical rainforest. *Biological Conservation*, 227:133–140, 2018. ISSN 0006-3207. doi:https://doi.org/10.1016/j.biocon.2018.09.009. URL https://www.sciencedirect.com/science/article/pii/S0006320718306396.

Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, April 2016. doi:10.1073/pnas.1517384113. URL https://www.pnas.org/doi/10.1073/pnas.1517384113. Publisher: Proceedings of the National Academy of Sciences.

Mauricio A Alvarez, Lorenzo Rosasco, Neil D Lawrence, et al. Kernels for vector-valued functions: A review. *Foundations and Trends® in Machine Learning*, 4(3):195–266, 2012.

M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, February 2019b. ISSN 0021-9991. doi:10.1016/j.jcp.2018.10.045. URL https://www.sciencedirect.com/science/article/pii/S0021999118307125.

George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, June 2021. ISSN 2522-5820. doi:10.1038/s42254-021-00314-5. URL https://www.nature.com/articles/s42254-021-00314-5. Publisher: Nature Publishing Group.

Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, pages 6572–6583, Red Hook, NY, USA, December 2018. Curran Associates Inc.

Christopher Rackauckas, Yingbo Ma, Julius Martensen, Caleb Warner, Kirill Zubov, Rohit Supekar, Dominic Skinner, and Ali Ramadhan. Universal differential equations for scientific machine learning. *arXiv preprint arXiv:2001.04385*, 2020b.

Keita Tokuda and Yuichi Katori. Prediction of unobserved bifurcation by unsupervised extraction of slowly time-varying system parameter dynamics from time series using reservoir computing. *Frontiers in Artificial Intelligence*, 7:1451926, 2024.

Gevik Grigorian, Sandip V George, and Simon Arridge. Learning governing equations of unobserved states in dynamical systems. *Physica D: Nonlinear Phenomena*, 472:134499, 2025.

Dimitrios Angelis, Filippos Sofos, and Theodoros E Karakasidis. Artificial intelligence in physical sciences: Symbolic regression trends and perspectives. *Archives of Computational Methods in Engineering*, 30(6):3845–3865, 2023.

Salvatore Cuomo, Vincenzo Schiano Di Cola, Fabio Giampaolo, Gianluigi Rozza, Maziar Raissi, and Francesco Piccialli. Scientific machine learning through physics–informed neural networks: Where we are and what's next. *Journal of Scientific Computing*, 92(3):88, 2022.

Agathe Girard and Roderick Murray-Smith. Gaussian processes: Prediction at a noisy input and application to iterative multiple-step ahead forecasting of time-series. *Switching and Learning in Feedback Systems: European Summer School on Multi-Agent Control, Maynooth, Ireland, September 8-10, 2003, Revised Lectures and Selected Papers*, pages 158–184, 2005.

Carl Boettiger. The forecast trap. *Ecology Letters*, 25(7):1655–1664, 2022.

Johan Rockström, Will Steffen, Kevin Noone, Åsa Persson, F Stuart Chapin, Eric F Lambin, Timothy M Lenton, Marten Scheffer, Carl Folke, Hans Joachim Schellnhuber, et al. A safe operating space for humanity. *nature*, 461(7263): 472–475, 2009.

Will Steffen, Katherine Richardson, Johan Rockström, Sarah E Cornell, Ingo Fetzer, Elena M Bennett, Reinette Biggs, Stephen R Carpenter, Wim De Vries, Cynthia A De Wit, et al. Planetary boundaries: Guiding human development on a changing planet. *science*, 347(6223):1259855, 2015.

Juan C Rocha, Garry Peterson, Örjan Bodin, and Simon Levin. Cascading regime shifts within and across scales. *Science*, 362(6421):1379–1383, 2018.

Fred Brauer, Carlos Castillo-Chavez, Zhilan Feng, et al. *Mathematical models in epidemiology*, volume 32. Springer, 2019.