communications earth & environment

Article

A Nature Portfolio journal

https://doi.org/10.1038/s43247-025-03130-2

Recovering complex ecological dynamics from time series using state-space universal dynamic equations

Jack H. Buckner[1] ✉, Zechariah D. Meunier[1], Jorge Arroyo-Esquivel[2], Nathan Fitzpatrick[3], Ariel Greiner[4], Lisa C. McManus[3] & James R. Watson[1]

Ecological systems often exhibit complex nonlinear dynamics like oscillations, chaos, and regime shifts. Universal dynamic equations have shown promise in modeling complex dynamics by combining known functional forms with neural networks that represent unknown relationships. However, these methods do not yet accommodate the forms of uncertainty common to ecological datasets. To address this limitation, we developed state-space universal dynamic equations by combining universal difference and differential equations with a state-space modeling framework, accounting for uncertainty. We tested this framework on three simulated and two empirical case studies and found that this method can recover nonlinear biological interactions that produce complex behaviors including chaos and regime shifts. Their forecasting performance is context-dependent, with the best performance on chaotic and oscillating time series. This innovative approach leveraging both ecological theory and data-driven machine learning offers a promising new way to make accurate and useful predictions of ecosystem change.

Ecological systems often exhibit complex dynamic phenomena like oscillations, chaos and nonlinear regime shifts. Quantifying these dynamics is valuable for both our understanding of these systems, and for applications of this new knowledge for improved ecosystem management. Identifying thresholds where regime shifts occur can help determine when management interventions are required to maintain desired states or to reverse undesired regime shifts[1]. Furthermore, models that capture nonlinear ecosystem dynamics can improve the ability to forecast future ecosystem states, creating opportunities for managers to proactively anticipate ecological change[2].

Time series data from long-term monitoring programs are an invaluable source of information for modeling changes in ecological systems. From these data researchers can identify the mechanisms driving ecosystem dynamics. However, modeling these data requires a flexible framework that can identify nonlinear relationships in the context of noisy observations and stochastic variation in unobserved factors. Recent advances in machine learning have produced a promising new class of models for this task: universal differential equations[3–5] for continuous-time models and universal difference equations for discrete-time models. To coin a term, we collectively call these models universal dynamic equations (UDEs). UDEs combine specific parametric functions to capture known relationships and physical constraints (i.e., ecological theory as represented mathematically) with neural networks (NNs) to learn unknown relationships directly from data.

The use of artificial NNs makes these methods promising because they can represent arbitrary nonlinear relationships in datasets and scale well in both the number of input dimensions and dataset size[6]. In principle, UDEs can improve forecasting by incorporating prior information about a system's structure through parametric functions[3]. The flexibility of NNs within UDEs also makes them valuable for inferring unknown nonlinear functions such as species interactions directly from time series data[7]. Furthermore, UDEs have the same mathematical structure as models commonly used in ecological theory, increasing their interpretability when compared to other machine learning techniques such as transformers[8], recurrent NNs, and long short-term memory networks[9].

Despite the promise of UDEs for describing nonlinear ecosystem dynamics, standard formulations of UDEs (e.g., Arroyo-Esquivel et al., 2024; Bonnaffé et al., 2021) do not account for the combination of stochastic external forcing and noisy observations common in ecological data. In fact, simple methods for training UDEs often fail when confronted with highly variable time series data that include either chaotic or stochastic dynamics[10]. To address these issues, we developed a new modeling framework called

[1]College of Earth, Ocean, and Atmospheric Sciences, Oregon State University, Corvallis, OR, USA. [2]California Department of Fish and Wildlife, West Sacramento, CA, USA. [3]Hawai'i Institute of Marine Biology, University of Hawai'i at Mānoa, Kāne'ohe, HI, USA. [4]Department of Biology, University of Oxford, Oxford, UK and Centre for Infectious Disease Dynamics, Pennsylvania State University, University Park, PA, USA. ✉e-mail: bucknejo@oregonstate.edu

state-space UDEs, where UDEs are embedded within a state-space modeling framework (*sensu* Auger-Méthé et al., 2021[11]). We hypothesized that state-space UDEs will leverage the flexibility of the UDE modeling framework for describing nonlinear local ecosystem dynamics while maintaining the ability of state-space models to accommodate the uncertainty common in ecological data.

We tested the ability of state-space UDEs to quantify nonlinear ecosystem dynamics using three simulated and two empirical datasets. The examples were chosen to cover two common forms of nonlinear dynamics observed in ecological systems: oscillating dynamics driven by trophic interactions, and alternative stable states created by self-reinforcing feedback between species and their environment. We selected one simulated and one empirical example in each of these two categories, with an additional simulated example to demonstrate how UDEs scale with large numbers of state variables, as is the case in complex ecosystems (e.g., food webs or spatial complexity). The simulated examples allowed us to test the model predictions with an unlimited number of replicates when the ground truth is known, while the empirical examples ensure that the models work under the constraints of real-world datasets. For the simulated examples, we generated datasets from: 1) a three-species food chain model[12] that exhibits non-periodic oscillations, 2) a model of nearshore rocky reefs that switch between states dominated by macroalgae and herbivores, and 3) a predator-prey metacommunity model with discrete patches that included many state variables. For the empirical examples, we used fisheries data from the RAM Legacy Database[13] that exhibits oscillations driven by bioeconomic feedback mechanisms, and a long-term agricultural monitoring dataset documenting desertification in an arid rangeland[14].

We used these examples to test the usefulness of state-space UDEs for forecasting future changes in the ecosystem's state and for making inferences about the underlying biological processes. We then compared the model's forecasting skill to two alternative methods, Gaussian process empirical dynamic models (gpEDMs), which are a powerful nonlinear forecasting tool frequently used for ecological datasets[15–17], and multivariate autoregressive state-space (MARSS) models, which fit linear dynamic models to time series data[18]. These two alternative methods allowed us to compare our method with other state-of-the-art techniques that each capture one of the two key aspects of our approach: 1) the ability to learn nonlinear relationships in time series data (gpEDMs), and 2) the use of a modeling framework to account for observation and process errors (MARSS). We also compared the state-space UDE models to two alternative methods for training UDEs, the gradient matching[19] and shooting method[5].

To highlight the value of state-space UDEs for ecology, we implemented UDE models to recover nonlinear species-species and species-environment interactions. These functional relationships underpin the complex dynamics of ecosystems[12,20], demonstrating our approach's ability to capture both the "what" and the "why" underlying changing ecosystems. Furthermore, we identified the stability and threshold parameters of alternative ecosystem states by constructing bifurcation diagrams from the trained UDE models. State-space UDEs can contribute to our understanding of ecosystems by providing a framework to quantify mechanisms leading to emergent ecosystem properties including stability, resilience, and ecological thresholds.

Results

State-space universal dynamic equation framework

UDEs are time series models that combine artificial NNs and parametric functions to describe changes in the state of a dynamical system over time. To date, most of the work on this topic has focused on continuous-time models called universal differential equations[5]. These models embed artificial NNs in the right-hand side of a system of differential equations

$$\frac{d\boldsymbol{u}}{dt} = f(\boldsymbol{u}, \boldsymbol{X}, t, NN(\boldsymbol{u}, \boldsymbol{X}); \theta), \tag{1}$$

where $\boldsymbol{u}$ is a vector of state variables, $\boldsymbol{X}$ is a vector of covariate factors, t is time, NN is the output layer of an artificial neural network, and θ is a set of parameters. These models allow the NN to learn the system's dynamics directly from data, while encoding known biological mechanisms and physical constraints through the function f (Fig. 1A, B).

Discrete-time dynamic models are also quite common in ecology, especially for seasonal environments. In these cases the NN can be embedded in the right-hand side $F(\boldsymbol{u}, \ldots)$ of a difference equation

$$\boldsymbol{u}_{t+1} = \boldsymbol{u}_t + F(\boldsymbol{u}, \boldsymbol{X}, t, NN(\boldsymbol{u}, \boldsymbol{X}); \theta), \tag{2}$$

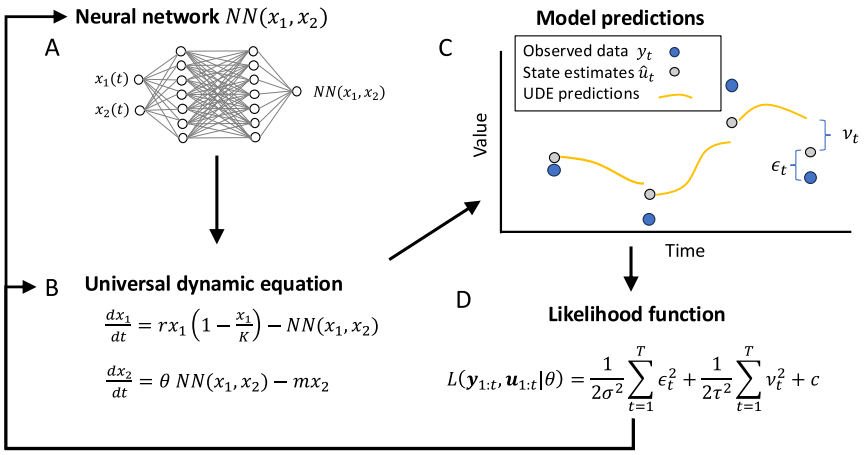

Fig. 1 | State-space universal dynamic equations schematic diagram. Schematic diagram of a state-space universal dynamic equation for a predator-prey model where x_1 is the prey and x_2 is predator. **A** Artificial NNs represent unknown relationships between state variables. **B** The NN output is combined with known functional forms (i.e., Lotka-Volterra predator-prey model with logistic growth of the prey) to build a model of the system, where $\boldsymbol{y}_t$ are the observations of the state variables x_1 and x_2 with errors $\epsilon_{1,t}$ and $\epsilon_{2,t}$. **C** The universal differential equation model predicts changes in the state of the system between observations (yellow lines). The state-space model estimates the true but unknown state of the system at each time point (gray dots), which is compared to the observed data (blue dots) to calculate observation errors ϵ_t. The differences between the predictions and the estimated states are process errors v_t. **D** The log-joint-likelihood function quantifies the performance of the model by computing the sum of squared observation and process errors. **E** An optimization algorithm trains the NN and estimates the model parameters and unknown states by maximizing either the log-joint-likelihood or log-marginal-likelihood function.

Table 1 | Types of models and training methods used in each example (Ex.) case study

Model type	Time	Training	State-space	Ex. 1	Ex. 2	Ex. 3	Ex. 4	Ex. 5
UDE (nonlinear, parametric)	Discrete	Marginal	Y		Y			Y
		Joint	Y		Y			Y
	Continuous	Matching	N			Y	Y	Y
		Shooting	N				Y	Y
		Marginal	Y			Y	Y	Y
		Joint	Y			Y	Y	Y
NODE (nonlinear, nonparametric)	Discrete	Marginal	Y	Y				Y
		Joint	Y	Y				Y
	Continuous	Matching	N	Y	Y	Y		Y
		Shooting	N	Y	Y	Y		Y
		Marginal	Y	Y	Y	Y		Y
		Joint	Y	Y	Y	Y		Y
gpEDM (nonlinear, nonparametric)	Discrete		N	Y	Y	Y	Y	
MARSS (linear, nonparametric)	Discrete		Y	Y	Y	Y	Y	

where F encodes known biological mechanisms and physical constraints in the discrete system. The state-space UDE framework provides a method for training the NNs and estimating unknown parameters in Eqs. 1, 2 on noisy and variable time series by embedding these models in a state-space modeling framework.

State-space models describe multivariate time series $\boldsymbol{y}_t$ observed at a discrete set of times $t \in 1:T$. The data $\boldsymbol{y}_t$ are noisy observations of a dynamical system with a true underlying state $\boldsymbol{u}_t$. The state-space modeling framework describes this data-generating process using a model with two parts: 1) an observation model for the relationship between the observations $\boldsymbol{y}_t$ and the state $\boldsymbol{u}_t$, and 2) a process model for the changes in the states over time. The strength of this approach lies in the model's ability to accommodate measurement errors that causes the observations $\boldsymbol{y}_t$ to differ from the underlying states $\boldsymbol{u}_t$, and unobserved factors that cause changes in the states over time $\boldsymbol{u}_t \rightarrow \boldsymbol{u}_{t+\Delta t}$ to be unpredictable (Fig. 1C).

The observation model can, in principle, represent complex relationships between the observations and states. However, in our examples, it is sufficient to assume that the observation $y_{i,t}$ of variable i at time t is equal to the underlying state $u_{i,t}$ plus a normally-distributed observation error $\epsilon_{i,t}$ with mean 0 and variance σ_i^2

$$y_{i,t} = u_{i,t} + \epsilon_{i,t}. \tag{3}$$

We use a universal dynamic equation (Fig. 1A, B) to describe changes in the system's state between observations $\boldsymbol{u}_t \rightarrow \boldsymbol{u}_{t+\Delta t}$ (Fig. 1C). For continuous-time universal differential equation models, the prediction is made by integrating Eq. 1 from t to $t + \Delta t$

$$\boldsymbol{u}_{t+\Delta t} = \boldsymbol{u}_t + \int_t^{t+\Delta t} f(\boldsymbol{u}(v), \boldsymbol{X}, v, NN(\boldsymbol{u}, \boldsymbol{X}); \theta) dv + \boldsymbol{v}_t, \tag{4}$$

where $\boldsymbol{v}_t$ captures process errors that arise from stochasticity in the underlying dynamics of $\boldsymbol{u}$ or missing information in the function f. In the discrete-time case, the prediction is obtained by adding the process error term $\boldsymbol{v}$ into Eq. 2

$$\boldsymbol{u}_{t+1} = \boldsymbol{u}_t + F(\boldsymbol{u}, \boldsymbol{X}, t, NN(\boldsymbol{u}, \boldsymbol{X}); \theta) + \boldsymbol{v}_t, \tag{5}$$

We assume that forecasting errors are independent across time and follow a multivariate normal distribution with mean 0 and covariance Σ_v.

We can derive two loss functions from the state-space modeling framework that will be useful for training the underlying UDE models: the joint-likelihood and the marginal-likelihood. The joint-likelihood calculates likelihood of the data $\boldsymbol{y}_t$ given the state estimates $\hat{\boldsymbol{u}}_t$ and a point estimate of the parameters θ_f. We can calculate the likelihood of the data given $\hat{\boldsymbol{u}}_t$ using our assumption that the errors are independent and normally-distributed by taking the product of the data likelihoods $\phi(\boldsymbol{y}_t, \hat{\boldsymbol{u}}_t, \Sigma_\epsilon)$, given the states estimates $\hat{\boldsymbol{u}}_t$. We can similarly calculate the likelihood of the states $\hat{\boldsymbol{u}}_t$ given the parameters θ by as the product of the state likelihood given the prior state $(\boldsymbol{u_{t-1}})$. The joint-likelihood is the product of these two terms

$$L_j\left(\boldsymbol{y}_{\{t\}} | \hat{\boldsymbol{u}}_t, \theta\right) = \prod_{t=1}^{T-1} \phi(\boldsymbol{y}_t, \hat{\boldsymbol{u}}_t, \Sigma_\epsilon) \phi(\hat{\boldsymbol{u}}_t, \mathrm{UDE}(\hat{\boldsymbol{u}}_{t-1}, X_t, \theta, t), \Sigma_v), \tag{6}$$

where Σ_ϵ is a diagonal matrix with $\Sigma_{\epsilon,i} = \sigma_i^2$, ϕ is the probability density function for a multivariate normal distribution, and the function UDE stands in for Eq. 4 or Eq. 5 depending on the model's formulation. UDE models can be trained using the joint-likelihood by solving for the unobserved state variables $\hat{\boldsymbol{u}}_t$ and the UDE model parameters θ to minimize the negative log-joint-likelihood of the data $\boldsymbol{y}_t$ and the states $\hat{\boldsymbol{u}}_t$.

The joint-likelihood approach is useful but it does not account for uncertainty over the estimated states $\hat{\boldsymbol{u}}_t$, leading to some bias. This bias can be eliminated by training the model with the marginal-likelihood, which equals the expectation of the log-likelihood of the observations integrating over the probability distribution of the states

$$L_m\left(\boldsymbol{y}_{\{t\}} | \theta\right) = \int_{-\infty}^{\infty} L_j\left(\boldsymbol{y}_{\{t\}} | \boldsymbol{u}_t, \theta\right) d\boldsymbol{u}_{1:T}. \tag{7}$$

Evaluating Eq. 7 requires an approximation of the distribution of the states given the observations and the parameter estimates $p(\boldsymbol{u}_t | \boldsymbol{y_{1:t}}, \theta)$. We approximate this density function with a multivariate normal distribution and calculate the mean $\hat{\boldsymbol{u}}_t$ and covariance $\Sigma_{\hat{u},t}$ terms using the unscented Kalman filter algorithm[21]. Additional details for all models and training methods are provided in the Methods and Table 1.

Recovering chaotic dynamics from a three-species food chain

We tested the ability of the state-space UDE framework to recover complex nonlinear dynamics by training a model on simulated time series from a three-species food chain ordinary differential equation system with chaotic dynamics (Fig. 2A, Hastings and Powell, 1991). We trained fully non-parametric neural ordinary differential equation (NODE) models to represent the system's dynamics. NODEs are the antecedent to UDEs, in that they simply use NNs alone to represent the dynamics of a system (i.e., a derivative), unlike UDEs which combine NNs and parametric functions to represent dynamics. NODEs reduce the influence of choices about the

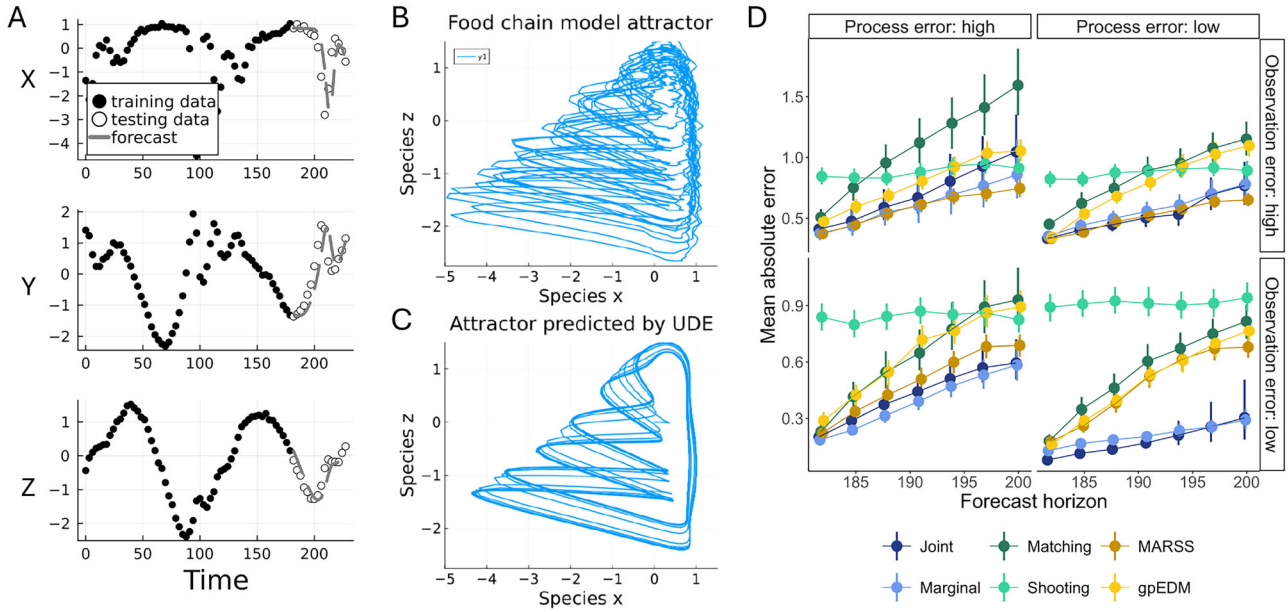

Fig. 2 | Three-species food chain model results. Performance of a state-space UDE model trained on simulated data from a three-species food chain. **A** Time series of the training and testing data compared to forecasted abundances of each species (X, Y, Z) using a continuous-time state-space UDE trained with the log-marginal-likelihood. **B** Chaotic attractor of the three-species food chain model. **C** Attractor predicted by UDE model trained on the data in panel A with the marginal-likelihood function. **D** The forecasting skill of the state-space UDE models (joint- and marginal-likelihood) compared to two alternate methods of training UDEs (gradient matching and shooting), a linear state-space model (MARSS), and an empirical dynamic model (gpEDM) under high and low amounts of process and observation error. Points are means ± bootstrapped 95% confidence intervals.

model structure on its performance, isolating the influence of the state-space formulation and training routines. We tested both continuous-time and discrete-time model formulations (Eqs. 1, 2) and the marginal- and joint-likelihood training methods (Eqs. 6, 7). We tested the model's performance under process error by adding white noise to the growth rate of the primary producer and measurements errors to the simulated abundances. Details for all models are provided in the Methods.

The state-space UDE approach was able to recover the qualitative behavior of the food chain model and generated the same pattern of irregular oscillations when forecasting as the true model (Fig. 2A–C). All formulations of the state-space UDE models (joint-likelihood, marginal-likelihood) outperformed the alternative training methods (gradient matching, shooting) at all levels of process and observation noise we tested (Fig. 2D). Based on mean absolute error, the state-space UDEs were better at forecasting than the gpEDM and MARSS models when the observation and process errors were small and had similar forecasting error compared to the gpEDM and MARSS models when both error terms were large (Fig. 2D). However, the state-space UDE models had larger mean squared forecasting errors than MARSS and gpEDM (Supplementary Fig. 1), indicating that the drop in performance was driven by a few forecasts with very large errors. The joint- and marginal-likelihood also produced similar forecasting accuracy. Forecasting performance was also similar between the discrete- and continuous-time UDE formulations (Supplementary Fig. 2).

Alternative stable states in nearshore rocky reef communities

We used a model of nearshore rocky reefs to test the ability of the state-space UDE framework to recover nonlinear species interactions that produce ecological tipping points and alternative stable states. In temperate latitudes, nearshore rocky reefs can switch between a kelp forest state dominated by macroalgae and an urchin barren state where sea urchins reduce kelp populations to very low densities[22–24]. One mechanism that might cause these alternative states to persist is sea urchin behavior in response to food availability[25]. When kelp are abundant, urchins exhibit cryptic behavior in which they hide from predators and feed on drift algae detached from the canopy or understory[25]. In contrast, when kelp are sparse, urchins leave cryptic microhabitats to feed actively on live macroalgae (Fig. 3A). This change in behavior can lead to an increase in urchin grazing as kelp abundance declines (Fig. 3B), creating a feedback that can cause the system to switch between kelp-dominated and urchin barren states (Fig. 3C).

We simulated 100 time series of kelp abundance with a discrete-time model that describes changes in kelp abundance over time given the nonlinear urchin grazing functional response and autocorrelated abiotic effects on growth. To test the ability of the UDE framework to recover species interactions, we fit parametric UDE models to the simulated time series of kelp abundance with 50 observations that mirrored the structure of the simulation model and used a NN $NN(y_t; \boldsymbol{w}, \boldsymbol{b})$ to represent the unknown urchin grazing functional response. We tested the effect of including parameters of the known functional forms in the training process on model performance by comparing models where the kelp growth rates were either fixed or estimated along with the weights and biases of the NN.

The parametric state-space UDE model was able to recover the characteristic nonlinearity in the urchin grazing functional response (Fig. 3D). Including abiotic conditions $\boldsymbol{X}_t$ in the models and fixing the growth rate parameter improved the model estimates (Fig. 3D). However, in some simulations, the model predicted no relationship between kelp abundance and grazing. This is associated with training sets that span periods when the system was entirely in the kelp-dominated or urchin barren state, because the nonlinear urchin grazing function is not needed to describe a time series near a single equilibrium. The UDE models trained with either the marginal- or joint-likelihood were also able to detect the presence of alternative stable states and the ecological threshold separating them in some cases (Fig. 3F and Supplementary Figs. 3, 4). However, the models predicted the tipping point much more reliably when the kelp growth rate was fixed rather than estimated (Supplementary Fig. 5). NODE models were also constructed and they proved to have limited skill in recovering the qualitative features of the bifurcation diagram (Supplementary Fig. 6), suggesting that constraining the model with known functions was valuable for making inference about the qualitative dynamics of this system. The state-space UDE models had similar forecasting skill to the gpEDM and MARSS models (Fig. 3E). In

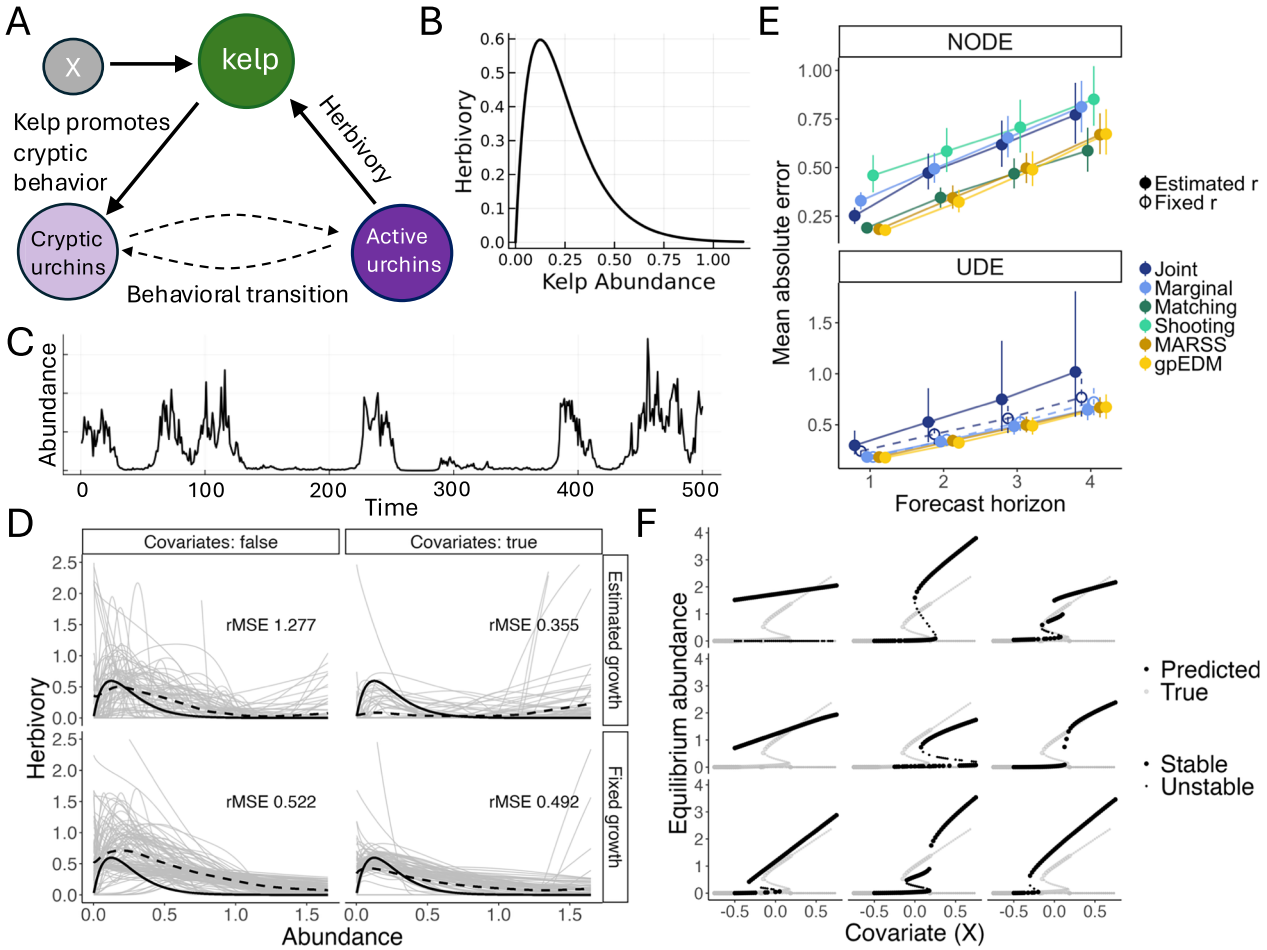

Fig. 3 | Kelp forest model results. Estimates of nonlinear functions that produce ecological tipping points between kelp forests and urchin barrens using state-space UDEs. **A** Schematic diagram of the feedback between kelp abundance and sea urchin behavior, with kelp abundance affected by environmental covariate X. **B** The functional response of urchin herbivory to kelp abundance caused by the behavioral feedback mechanism. **C** An example time series of kelp abundance illustrating the flickering behavior caused by the nonlinear urchin grazing functional response. **D** The estimated relationship between kelp abundance and urchin herbivory from the parametric UDE model with and without observation of the covariate factor X and with fixed or estimated values of the growth parameter r. The solid black lines show the true relationship, the gray lines show the individual estimated relationships, and the dashed black lines show the average of the estimated relationships. Model performance is given by root mean squared error (rMSE). **E** The forecasting skill of a parametric UDE model and nonparametric UDE model (NODE) compared to a linear state-space model (MARSS) and empirical dynamic model (gpEDM). This panel also compares the performance of the UDE models across the four alternative training routines (joint-likelihood, marginal-likelihood, gradient matching, and shooting). We did not test gradient matching and shooting with the parametric UDE model because it uses a discrete-time formulation. Points are means ± bootstrapped 95% confidence intervals. **F** Bifurcation diagrams predicted by the parametric UDE with a fixed growth rate trained separately on nine datasets simulated using different sequences of random numbers sampled for the process and observation error terms.

contrast, the state-space NODEs (joint- and marginal-likelihood) were less accurate than the alternative times series models (MARSS, gpEDM) and the NODE trained with gradient matching.

Forecasting oscillating dynamics in fisheries data

Commercial fisheries can exhibit boom and bust cycles characterized by periods of overexploitation that cause declines in the harvested stock followed by periods of low harvest and stock recovery[26] (Fig. 4A, B). These so-called bioeconomic cycles are caused by delayed feedback between the abundance of the harvested population and the incentives of firms to invest in the fishery[27]. Fisheries management can smooth out these cycling dynamics by setting harvest limits that are linked to the abundance of the population[28], creating more responsive feedback between the stock and harvest[29].

We developed a UDE model to quantify the feedback between the fish population and harvest, and to forecast the dynamics of fisheries under different management scenarios. The UDE model describes changes in the population abundance B and harvest rate H. We modeled the changes in abundance B using the logistic model minus the harvest rate H. We modeled harvest as $H = UB/q$, where U is fishing mortality and the factor q ensures the units of harvest and abundance match. We modeled changes in fishing mortality with a NN and derived a model for the resulting changes in harvest in the Methods.

We fit the UDE models to time series of harvest and abundance from two groundfish species, chilipepper rockfish (*Sebastes goodei*) and cowcod (*Sebastes levis*), sourced from the RAM Legacy Stock Assessment Database (2018). Both fisheries are from the Pacific coast of the United States and experienced major regulatory changes in the 1990s. To reflect this, we included an indicator variable $I_{t>1992}$ in the NN that takes the value of zero pre-1992 or one post-1992 when access to the fisheries was restricted and other regulations were introduced[29].

The UDE model predicted nonlinear relationships between the fish population biomass and changes in fishing mortality. Models for both groundfish species predicted that fishing effort would increase (positive derivatives) when the stocks were abundant and decline rapidly (negative derivatives) when the stocks were scarce (Fig. 4C,D) as expected by economic theory[26,27]. Although the overall shape of the relationship differed

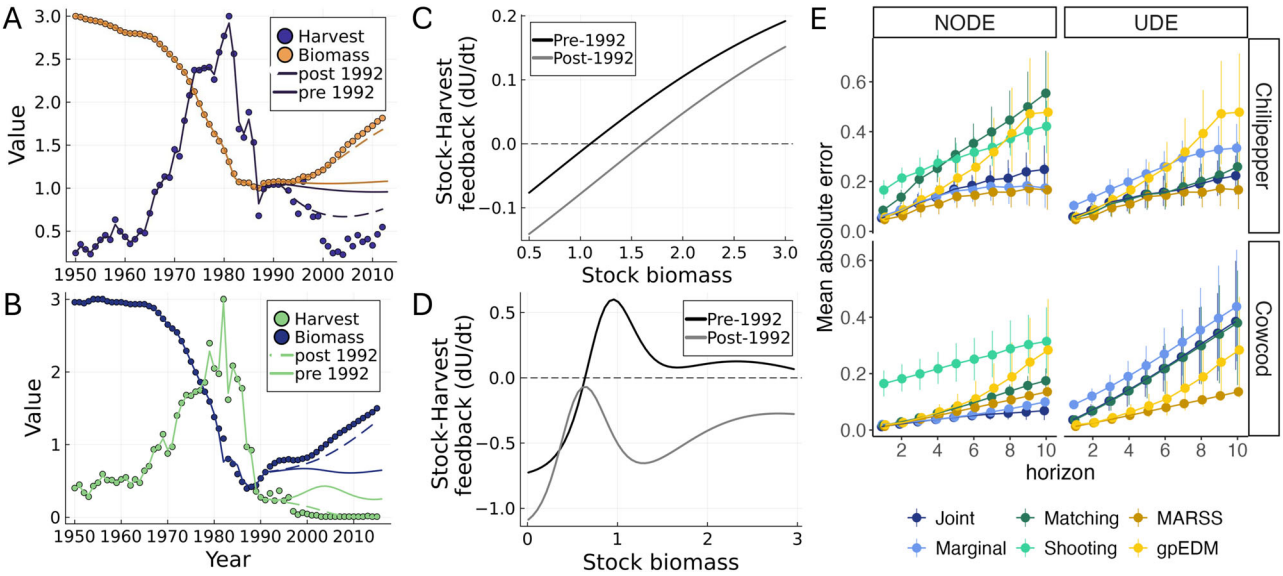

Fig. 4 | Fisheries bioeconomic model results. Forecasts of state-space UDE models for two groundfish fisheries. Time series of stock and harvest biomass with forecasts under the pre- and post-1992 management regimes for (**A**) chilipepper rockfish and (**B**) cowcod. Estimated relationships between the harvestable stock and the change in fishing effort over time for (**C**) chilipepper rockfish and (**D**) cowcod. **E** Forecasting skill of UDE models trained with the two state-space methods (joint- and marginal-likelihood) or shooting and gradient matching algorithms compared to a Gaussian process empirical dynamic model (gpEDM) and a multivariate autoregressive state-space model (MARSS) over one-to-ten-year time horizons. The shooting algorithm did not converge for the UDE model so those forecasting results are omitted. Points are means ± bootstrapped 95% confidence intervals.

substantially between the two stocks, the stock-harvest feedback functions for each groundfish were qualitatively similar before and after the change in regulations in 1992. Interestingly, the stock size where the change in fishing effort became negative was larger after 1992 for chilipepper rockfish, whereas the change in fishing effort became entirely negative for all stock sizes of cowcod. This indicates that fishing effort became more responsive to fish population abundance after the change in management for chilipepper rockfish and was reduced uniformly for cowcod, which is consistent with the change in regulations that occurred for these species[30–32].

Both the state-space UDE and NODE models had similar forecasting skill to the MARSS model and outperformed the gpEDM and alternative UDE training algorithms at forecasting the chilipepper rockfish time series (Fig. 4E). For the cowcod time series, the state-space NODE models had the highest forecasting performance of all models tested over longer time horizons (5–10 years). The state-space UDE models had more limited forecasting skill for the cowcod fishery, suggesting that the increased flexibility of the NODE model was important for capturing key features of the system's dynamics.

Identifying ecological thresholds in an arid rangeland

We tested the ability of state-space UDEs to recover ecological thresholds in an empirical dataset using long-term monitoring data from the Jornada Experimental Range (JER) in the Chihuahuan Desert of New Mexico, USA[14] (Fig. 5A, B). Plant communities in arid rangelands like the JER can exhibit rapid changes between states dominated by perennial grasses, woody forbs, and bare ground[33] (Fig. 5C). Transitions between community states can be caused by the management of livestock grazing, abiotic factors like temperature and precipitation, and local feedbacks between local plant abundances and soil conditions[34,35]. For example, grasses can increase the ability of rainfall to infiltrate soils and reduce horizontal transport of water and nutrients. This results in a homogeneous distribution of soil moisture that promotes the grass-dominated community (Fig. 5C). However, if exogenous drivers cause grass densities to decline, then infiltration will decrease, resulting in a heterogenous distribution of soil moisture that promotes a shrub-dominated state[35,36] (Fig. 5C).

The JER dataset documents a regime shift from a grass-dominated state to a shrub-dominated state starting in the mid-1950s when drought conditions caused an abrupt decline in grass densities, followed by shrub encroachment[37] (Fig. 5A, B). We developed two state-space UDE models of community composition in the JER to assess their ability to recover the feedback mechanisms that produced the regime shift. The first model (the competition UDE) uses a NN to estimate species interactions and known functions to capture competition for space. The second model (the plant-soil feedback UDE) uses known functional forms to mechanistically represent the feedback between plant community composition and soil conditions hypothesized to cause regime changes in arid rangelands.

The plant-soil feedback UDE predicted a positive effect of grass density and a negative effect of shrubs on rainfall infiltration (Fig. 5D), matching the patterns reported in the literature (e.g., Schlesinger et al., 1990). This relationship stabilizes alternative grass-dominated and shrub-dominated states separated by a fold bifurcation at low precipitation (Fig. 5E). The more flexible competition UDE also identified a fold bifurcation as a function of precipitation (Fig. 5E) along with more complex relationships between the other two species groups and rainfall. The state-space UDE models were able to recover the presence of alternative stable states in the system using data collected before the observed regime change in the 1950s (Fig. 5F). The plant-soil feedback UDE made more consistent predictions of the threshold precipitation separating these two alternative regimes.

State-space UDEs were the best forecasters in some cases. The more flexible competition UDE generally had better forecasting performance than the plant-soil feedback UDE and the joint-likelihood training routine performed better than the marginal-likelihood for this dataset. The state-space UDE models improved forecasting accuracy the most compared to the alternative training routines, gpEDM, and MARSS models when the regime changes in the 1950s were included in the testing set (short dataset, Fig. 5G). The gradient matching training routine also had high forecasting skill when the regime change was not included in the testing set (full dataset, Fig. 5G).

Scaling to high-dimensional predator-prey datasets with spatial structure

The increase in forecasting accuracy of state-space UDEs compared to the gradient matching and shooting algorithms comes at the expense of computational time, especially when the dataset includes many state variables. We tested the ability of the state-space UDE framework to scale to high-

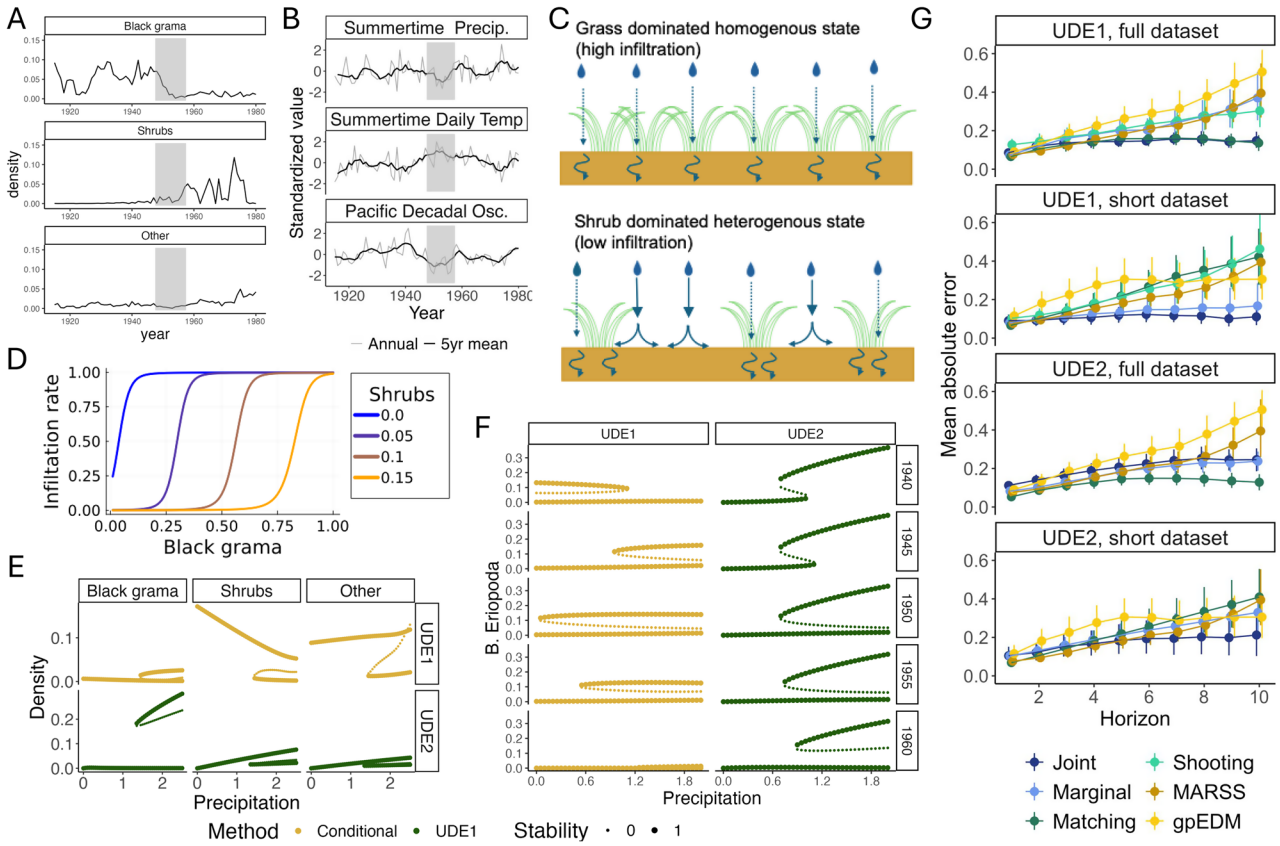

Fig. 5 | Arid rangeland model results. UDE models of alternative stable states in an arid rangeland. **A** Time series of densities of the dominant grass species black grama (*Bouteloua eriopoda*), shrubs, and other vegetation species from sites with sandy soils in the Jornada Experimental Range in the Chihuahuan Desert, New Mexico. The period of community transition (gray shading) from a grass- to shrub-dominated state occurred around the year 1950. **B** Time series of summer precipitation, summer temperature, and the Pacific Decadal Oscillation (PDO). The rangeland community transition highlighted in gray corresponds with decreased precipitation, increased temperatures and aridity, and a persistent negative PDO. **C** Schematic diagrams of the rainfall infiltration of soil under the grass-dominated and shrub-dominated states. Under high grass densities, the velocity of falling precipitation and horizontal transport along the ground is reduced, resulting in relatively even soil infiltration and moisture levels. In the shrub-dominated state, rainfall landing on bare ground can flow more easily, concentrating in pools and producing heterogenous soil moisture. **D** The estimated relationship between black grama and shrub densities and the rate of rainfall soil infiltration. **E** Bifurcation diagrams of the trained competition model (UDE1) and plant-soil feedback model (UDE2), showing the equilibrium abundance of each species group as a function of precipitation. **F** Competition and plant-soil feedback UDE model predictions of equilibrium grass density as a function of precipitation based on subsets of data starting in 1915 and ending in the year indicated by the panel title. **G** The forecasting performance of the two UDE models trained with the two state-space training methods, gradient matching and shooting algorithms compared to gpEDM and MARSS models. Points are means ± bootstrapped 95% confidence intervals.

dimensional datasets by training models on data simulated from a spatially structured predator-prey model (Supplementary Fig. 15), where we could vary the number of locations included in the model and hence the modeled dataset. Increasing the number of state variables might influence the efficiency of the training routine either by increasing the number of inputs to the NN, or by increasing the number of unobserved states that need to be estimated in the state-space modeling framework. To isolate these factors, we compared the performance of two models: 1) a fully nonparametric model (NODE) that used every state variable as an input to a NN, and 2) a structured model (UDE) that used a NN to estimate species interactions at each location in the dataset and used parametric functions to represent movement of species between locations. In the latter case, the number of state variables in the model increased with the number of sites but the complexity of the NN remained constant. We found that both the marginal- and joint-likelihood formulations of the state-space UDE models scaled to higher dimensional datasets less efficiently than the shooting and gradient matching methods (Table 2). The marginal-likelihood approach scales the least efficiently because it computes estimates of the mean and covariance for each estimated state when calculating the loss, which scales $O(n^2)$. Conversely, the joint-likelihood approach only calculates the mean, which scales $O(n)$. The discrete-time models trained much more quickly and scaled to higher dimensional datasets more efficiently than the continuous-time models because they do not require ODE solvers to evaluate the loss function (Table 2).

Discussion

In this study, we showed that state-space UDEs can recover nonlinear interactions that produce complex ecosystem dynamics even in the presence of environmental stochasticity and noisy ecological data. As a result, this new class of models can represent a wide range of dynamic phenomena in ecology and conservation including chaos, regime shifts, management interventions, and population cycles. Our findings from the Jornada Experimental Range dataset even suggest that state-space UDEs can identify alternative stable states and ecological thresholds before a regime shift occurs. Because of their ability to recover nonlinear dynamics from noisy data, state-space UDEs can be valuable both for forecasting and as interpretive tools for inferring interactions between organisms and their environment.

State-space UDEs perform well at these tasks by smoothing the dataset to account for measurement errors and estimating the nonlinear functions that determine the system's dynamics while accounting for uncertainty. Accounting for imperfect measurements (observation errors) can reduce overfitting to training data by limiting the influence of noise from the observation process on the trained model. Accounting for uncertainty

Table 2 | The estimated time in minutes required to train UDE models on the spatial predator-prey datasets for 1000 iterations of the Adam gradient descent algorithm using the joint-likelihood, marginal-likelihood, gradient matching, and shooting training methods

Model time	Model type	Training method	Scaling	Rate (min/1000 iter.)	Timing 10 states (min/1000 iter.)	Timing 20 states (min/1000 iter.)
Continuous	NODE	Joint	2.83	0.015	11.1	68.0
		Marginal	NA	NA	NA	NA
		Matching	**0.24**	0.173	**0.299**	**0.343**
		Shooting	3.14	**0.001**	1.17	7.97
	UDE	Joint	1.15	0.358	5.11	10.8
		Marginal	2.61	0.235	96.3	NA
		Matching	**0.88**	**0.057**	**0.377**	**0.837**
		Shooting	1.1	0.074	0.936	2.04
Discrete	NODE	Joint	**0.27**	**0.195**	**0.352**	**0.411**
		Marginal	1.33	0.246	6.32	12.4
	UDE	Joint	**0.72**	**0.132**	**0.750**	**1.24**
		Marginal	1.71	0.463	23.4	79.2

The scaling and rate columns estimate how much computation time increases as function of the number of sites included in the dataset. Cells with NA values indicate models that ran too slowly to assess their performance. Cells with bold font performed the fastest per model type and time.

(process errors) in the model's predictions reduces the likelihood that this smoothing process removes real variation in the system's states that cannot be predicted from previous observations and covariates alone.

The marginal-likelihood training routine allows the level of process error to be directly estimated from the data and should, in theory, reduce biases in the parameter estimates by accounting for uncertainty in the underlying state variables. However, in practice we found that the joint-likelihood approach also performed well at forecasting and inference tasks and scaled to high-dimensional datasets more efficiently than the marginal-likelihood approach. The two alternative UDE training routines (gradient matching and shooting) generally performed more poorly than the state-space training methods. The shooting method trains the model by simulating a solution to the ODE model over the full length of the training set, causing it to attribute real variation in the time series that cannot be explained by nonlinear dynamics to observation error. This limitation can be reduced by simulating solutions to the ODE over smaller sections of the dataset (e.g., Turan and Jäschke, 2022[10]). These methods can provide some computational advantages and might be the better choice when process errors are small. The gradient matching algorithm could be a more promising approach than shooting-based methods because it can accommodate some process error and does not require the use of ODE solvers, which greatly reduces computational costs. However, it smooths the time series when it does not include information about the dynamics of the system, which could explain why it had lower forecasting accuracy than state-space UDEs in most of our examples. This limitation can be addressed in NODE models by training the smoothing function and the NN weights and biases simultaneously[7,38]. To our knowledge, this has not yet been applied to the more general case of universal dynamic equations but would be a promising direction for future research.

UDEs may perform well at several tasks that are critical for understanding ecosystem dynamics because of their ability to combine prior information from known functions and parameters with functions and parameters learned from the data. This combination allows UDEs to quantify species-species interactions (e.g., urchin herbivory of kelp) and species-environment interactions (e.g., feedback between rainfall and grass density) in an ecologically-meaningful way. In these cases, known functions are required for the NN function to be ecologically interpretable. Moreover, nonlinear feedback mechanisms are often critical for driving complex patterns of change in ecological systems[1,12,20], highlighting the value of our method for understanding drivers of change or resilience. In our examples, we illustrated how UDEs can be used to estimate a nonlinear functional response that stabilizes alternative ecosystem states, how fisheries management policy can strengthen feedback between population abundance and harvest rate to increase the resilience of the ecosystem, and how nonlinear effects of species on their abiotic environments can lead to regime shifts. Our state-space UDE framework illuminates these case studies by providing a robust method for training the models on ecosystem monitoring data.

Universal dynamic equations can also be valuable for modeling systems with multiple dynamic regimes (e.g., alternative stable states). In the kelp forest example, including a model for kelp growth and the associated growth rate in the absence of urchin herbivory allowed the UDE to identify the presence of a tipping point more reliably in the system's dynamics. In the Jornada rangeland example, we found that a UDE model explicitly encoding the key feedback mechanisms between the plant community and soil conditions produced more consistent estimates of the ecosystem's tipping point over time. These findings also indicate that UDE models can be valuable for monitoring and experimental design. These two examples both highlighted key mechanisms that need to be identified to confirm the presence of bistability in the systems: the growth rate in the kelp model, and the plant-soil interaction in the rangeland case study. Both can be directly measured in the field, which would affect the credence given to the predictions of regime shifts made by the UDE models trained on empirical data.

Adding known functions to the state-space UDE models had mostly negative effects on forecasting skill in our examples. This is consistent with prior studies on forecasting that show fully nonparametric models can outperform correct parametric models[2], and in similar recent experiments with standard UDE models[3]. Taken together, these results suggest that adding parametric functional forms to UDE models is primarily useful for inference tasks and recovering qualitative features of the system's dynamics like alternative stable states. This can be important for using other methods for anticipating large and rapid regime shifts, such as the use of early-warning signals based on critical slowing down[39,40] or flickering[41].

Despite their potential, the usefulness of state-space UDEs is likely context-dependent. We found that improvements in forecasting skill compared to other methods were most pronounced over longer forecasting horizons and when nonlinear dynamics drove large changes in the ecosystem's state. We found the largest increase in forecasting skill in the three-species food chain and fisheries datasets, and the two exceptions to this prove the rule: 1) UDEs outperformed the linear state-space model in the rangeland dataset when we modified our cross-validation routine to include the regime change in the testing set, and 2) the MARSS models worked similarly well to the UDEs in the fisheries datasets that exhibited more regular oscillating patterns that could be captured by a linear model. Finally,

we expect that state-space UDEs will provide limited additional value for both forecasting and inference on time series that remain near a single equilibrium point, since local dynamics near equilibria can always be approximated with a linear model.

Given UDE's potential value for making inferences about nonlinear species interactions from time series data, an important future direction will be quantifying model uncertainty. This could, in principale, be achieved with Bayesian methods including stochastic gradient Langevin dynamics, or related Monte-Carlo methods which have already been successfully applied to neural ordinary differential equations[7,42]. Another limitation is the ability of the models to scale to high-dimensional datasets, particularly when training with the marginal-likelihood. However, it may be possible to mitigate this limitation by replacing the Kalman filter method we used to approximate the marginal-likelihood with a particle filtering algorithm[43]. Furthermore, even though ecosystems can have hundreds of species, it is common for analysts to select a subset of highly abundant or otherwise important species to be modeled (e.g., Langendorf et al., 2025[44]) or to aggregate species into functional groups for analysis (e.g., Meunier et al., 2024[45]), Thus, analyses of complex ecosystems often focus on fewer than 20 species, well within the computational limits of state-space UDEs.

Developing effective means of training UDE models on ecosystem monitoring data can open up many new applications in ecology beyond those presented here. In our case studies, we used the model structure to encode known mechanisms governing a single population's dynamics and physical constraints on space. However, models of ecosystem-level processes might include constraints on the conservation of matter and energy. UDEs could also be used for food web models by leveraging NNs to capture the context dependence and modularity of trophic interactions[44]. Finally, UDEs could be valuable for models of systems with spatial connectivity, where spatial linkages are known or can be estimated from hydrodynamic models[46], but in which local-scale nonlinear population dynamics are unknown.

Altogether, state-space UDEs are a powerful new method for uncovering nonlinear ecosystem dynamics in the context of multiple forms of uncertainty common to ecological datasets. UDE models link what we know to be true from observational and experimental studies to flexible mathematical functions; these models can then be trained with time series data from long-term environmental monitoring. As a result, they can encode mechanisms identified from empirical studies and estimate their influence on system dynamics in a quantitatively rigorous framework. By utilizing multiple sources of information in this way, state-space UDEs can be a valuable tool for synthesizing ecosystem data and, under the right conditions, forecasting ecological change.

Methods

Training state-space UDEs

We trained the state-space UDE models by maximizing the log-likelihood functions (Eqs. 6, 7) using the Adam[47] gradient descent algorithm implemented in the Optimizers.jl[48] library in the Julia programming language[49]. We computed the gradients of the loss function via automatic differentiation using Zygote.jl[50]. For continuous-time models, we numerically solved the integral in Eq. 4 using the Tsit-5 algorithm[51] implemented in DifferentialEquations.jl[52] and use DiffEqFlux.jl[5] for compatibility between the ODE solvers and the automatic differentiation. We implemented this training procedure using the UniversalDiffEq.jl library.

We use feedforward NNs with a single hidden layer in each model and included an L2 regularization penalty term for the weights of the network to control the complexity of the functions it learned. In the two empirical examples, we select the regularization weight λ by comparing the performance of the same model trained with different values of λ using a leave-future-out cross-validation routine.

Both the marginal- and joint-likelihood methods require values of the observation and process error terms. When training with the joint-likelihood, these terms we chose to fixed these terms at pre-defined values. Both terms can be estimated when training with the marginal-likelihood function, but in our examples, we fix the observation error Σ_ϵ and estimated the process error Σ_ν to reduce the number of free parameters estimated in the training procedure and because Σ_ϵ can often be estimated directly from the sampling procedure.

Approximating the marginal-likelihood

The marginal-likelihood (Eq. 7) cannot be computed directly and instead must be approximated numerically. The common approaches are Laplace approximations and filtering algorithms including Kalman and particle filters. We used an extension of the Kalman filter for nonlinear models called the unscented Kalman filter[21]. This approach approximates the distribution of the states at time t given the observations up to time t and the parameter estimates $p(\boldsymbol{u}_t|\boldsymbol{y}_{1:t-1},\theta)$ with multivariate normal distributions. The algorithm uses an iterative procedure to calculate mean $\hat{\boldsymbol{u}}_t$ and covariance $\Sigma_{u,t}$ at each time given the mean and covariance of the states in the previous time step $(\hat{\boldsymbol{u}}_{t-1},\ \Sigma_{u,t-1})$ by first propagating uncertainty through the dynamic model to get an initial estimate $(\hat{\boldsymbol{u}}'_t,\ \Sigma'_{u,t})$ and then applying Bayes' formula to update these estimates by conditioning on the observation $\boldsymbol{y}_t$ yielding the final estimates $(\hat{\boldsymbol{u}}_t,\ \Sigma_{u,t})$. The algorithm is initialized at the first time point in the dataset using the observed value as the initial mean $(\hat{u}_0 = y_0)$ and the observation errors as the initial covariance $(\Sigma_{u,0} = \Sigma_\epsilon)$. We then calculated the approximate negative log-marginal-likelihood of each observation by integrating the likelihood of the data given the states over the distribution of the states

$$L_m(\boldsymbol{y}_t|\theta) = \int_{-\infty}^{\infty} \phi(\boldsymbol{y}_t|\boldsymbol{u},\Sigma_\epsilon)\phi(\boldsymbol{u}|\hat{\boldsymbol{u}}_t,\Sigma_{u,t})\,d\boldsymbol{u}. \tag{8}$$

Under this approximation, $\boldsymbol{y}_t$ is the sum of two normal random variables, the states $\boldsymbol{u}_t$ and observation errors $\boldsymbol{\epsilon}_t$, so the marginal distribution of $\boldsymbol{y}_t$ is multivariate normal with mean $\hat{\boldsymbol{u}}_t$ and covariance $\Sigma_{u,t} + \Sigma_\epsilon$

$$\phi(\boldsymbol{y}_t) = \phi(\boldsymbol{y}_t|\hat{\boldsymbol{u}}_t,\Sigma_{u,t} + \Sigma_\epsilon). \tag{9}$$

We calculated the negative log-marginal-likelihood by summing the log-marginal-likelihood of each observation

$$\begin{aligned} l_m(\boldsymbol{y}_t|\theta) = &-\frac{1}{2}\sum_{t=1}^{T}(\boldsymbol{y}_t - \hat{\boldsymbol{u}}_t)^{\mathrm{T}}(\Sigma_{u,t} + \Sigma_\epsilon)(\boldsymbol{y}_t - \hat{\boldsymbol{u}}_t) \\ &-\frac{1}{2}\log\left(\mathrm{Det}(\Sigma_{u,t} + \Sigma_\epsilon)\right) - \frac{d}{2}\log(2\pi) \end{aligned} \tag{10}$$

Gradient matching algorithm

We implemented a simple version of the gradient matching algorithm (Ellner et al., 2002), which works by fitting smoothing functions $\boldsymbol{u}(t)$ to the time series data and then trains the UDE model by comparing the right-hand side of the UDE model (Eq. 1) to the derivatives of the smoothing functions $\boldsymbol{u}\prime(t)$. Our implementation uses a regularization smoothing algorithm implemented in the DataInterpolations.jl library[53], which smooths each dimension of the time series $y_{i,t}$ using a ridge regression[54], where the regularization parameter λ is chosen by generalized cross-validation[55]. We calculated the loss by comparing the predictions of the right-hand side of the UDE model (Eq. 1) to the derivatives of the smoothing splines at the time of each observation in the dataset with the squared error

$$L_{gm}(\boldsymbol{u}(t)|\theta) = \sum_{t\in T}\sum_{i=1}^{d}\left(u'_i(t) - f_i(\boldsymbol{u}(t),X(t),t,\theta)\right)^2. \tag{11}$$

We minimized the loss function using the Adam gradient descent algorithm with gradients calculated by automatic differentiation using Zygote.jl[50].

Shooting algorithm

The shooting algorithm trains UDE models by minimizing a loss function calculated by numerically solving the ODE model (Eq. 1) over the time span of the dataset to estimate the value of the system's states at each point in the time series $\hat{\boldsymbol{u}}(t)$. In our implementation, we solved the system of ODEs using the Tsit-5 algorithm. We then compared these estimated states to the observations using the squared error loss

$$L_s\left(\boldsymbol{y}_{\{t\}}|\theta,\hat{\boldsymbol{u}}_{t_0}\right)=\sum_{t\in \mathrm{T}}\sum_{i=1}^{d}\left(y_{i,t}-\hat{u}_i(t)\right)^2. \tag{12}$$

We minimized the loss function using the Adam gradient descent algorithm with gradients calculated by automatic differentiation using Zygote.jl[50] and DiffEqFlux.jl for compatibility with the ODE solver[5].

Regularization

To prevent the NNs included in the models from overfitting to the training set we added an L2 penalty term over the weights of the NN to the loss function weighted by a factor λ

$$L\left(\boldsymbol{y}_{\{t\}}|\theta,\lambda\right)=L\left(\boldsymbol{y}_{\{t\}}|\theta\right)+\lambda\sum_{i=1}^{m}w_i^2. \tag{13}$$

We trained the models without regularization ($\lambda=0$) in the simulated examples to reduce the number of factors being tested. In the empirical applications, we selected the regularization weight using leave-future-out cross-validation to estimate the model's performance on out-of-sample data. We searched over a grid of values ranging from 10^{-4} to 10^{8} increasing in powers of ten. The values chosen for each model in the fisheries and Jornada rangeland examples are given in Supplementary Methods 3 and 4, respectively.

Cross-validation

We estimated the state-space UDE model performance on out-of-sample data using leave-future-out cross-validation. Leave-future-out cross-validation uses the beginning of the dataset to train the model and uses the final observations as testing data. The algorithm constructs multiple testing sets by sequentially reducing the length of the training set, removing the final observation in the training set and adding it to the testing set. We calculated the model's performance by forecasting from the estimated value of the state variables at the final time point in the training data $\hat{u}_T$ over the full length of the testing set, which we varied in length to test the model's performance on different forecasting horizons. We quantified the accuracy of the forecasts using the mean absolute error (MAE) and mean squared error (MSE). We calculated uncertainty in the estimates of MAE and MSE by bootstrapping 95% confidence intervals using the Hmisc library in R[56].

Evaluating model performance

In examples one through four, we sought to validate the state-space UDE framework's ability to recover nonlinear dynamics from time series data and compare its forecasting ability to alternative methods for training UDEs and time series models. Example five was strictly used to test the computational efficiency of the training algorithms on large datasets. In examples one through four, we developed a state-space UDE model with a structure appropriate for the dataset and tested the model's ability to make inference from the training data by comparing estimated functional forms to either the ground truth in the case of the simulated examples, or qualitative features documented in the literature in the case of the empirical examples. We then compared the forecasting skill of the state-space UDE model to models with the same structure trained with alternative methods. Specifically, if the primary UDE contained known functional forms, then we also tested a fully nonparametric UDE model called a neural ordinary differential equation (NODE), along with the gpEDM and MARSS models. We assessed forecasting skill in the simulated datasets by forecasting observations left off the end of the time series and averaging the forecasting errors over many simulated datasets. For the empirical examples, we used leave-future-out cross-validation to estimate forecasting skill. Note that not all combinations of models and training routines are compatible with each of the example datasets. For example, models with discrete-time formulations cannot be trained with the gradient-matching technique. The models we evaluated in each case study are listed in Table 1.

Example 1: Three-species food chain simulation

We extended the Hastings and Powell (1991) three-species food chain model to include a source of stochasticity by adding a multiplicative white noise term with variance σ_x^2 to the basal resource yielding a system of three stochastic differential equations:

$$\frac{dx}{dt}=x(1-x)-\frac{a_1xy}{1+b_1x}+\sigma_x x\mathrm{d}W_t \tag{14}$$

$$\frac{dy}{dt}=\frac{a_1xy}{1+b_1x}-\frac{a_2yz}{1+b_2y}-d_1y+0\mathrm{d}W_t \tag{15}$$

$$\frac{dz}{dt}=\frac{a_2yz}{1+b_2y}-d_2z+0\mathrm{d}W_t. \tag{16}$$

We simulated time series with 75 observations sampling at an interval of 3.333 units of model time. We standardized the time series to have mean 0 and variance 1 before adding normally-distributed measurement error with mean 0 and standard deviation σ_e. The remaining model parameters (Supplementary Table 1) were chosen so that the model would exhibit chaotic dynamics in the absence of the white noise term. All models were trained on the first 60 observations and the last 15 were used to test the model's forecasting ability. Hyperparameters used in the model training procedure are listed in Supplementary Table 2.

We simulated 100 total datasets with four combinations of process and observation noise: high process noise $\sigma_x=0.1$ and high observation noise $\sigma_e=0.25$, low process noise $\sigma_x=0.025$ and low observation noise $\sigma_e=0.05$, and the two mixed cases. We compared the forecasting skill of the UDE model to a Gaussian process empirical dynamic model (gpEDM, Munch and Rogers, 2024[17]), and a multivariate autoregressive state-space model (MARSS, Holmes et al., 2012[18]). The full details of these alternative models are provided in Supplementary Methods 1. We evaluated the forecasting skill of each model by forecasting using the trained model from the final time point in the training set over the full length of the testing set. We calculated the MAE and MSE between the forecasts and the testing data along with 95% bootstrapped confidence intervals of these metrics using the Hmisc library[56].

Example 2: Kelp forest model simulation

We simulated time series of kelp abundance y_t using a discrete-time model that accounted for density-dependent growth r, kelp carrying capacity K, urchin herbivory $h(y_t)$, observed abiotic factors X_t, and unobserved abiotic factors v_t. We modeled a density-dependent effect on growth using the Ricker model with additive effects for herbivory and abiotic conditions:

$$y_{t+1}=y_te^{r\left(1-y_t/K\right)-h\left(y_t\right)+X_t+v_t}. \tag{17}$$

We selected a functional form for urchin functional response to kelp that increases at very low kelp densities up to a maximum value, and then declines at high and intermediate kelp densities to match the qualitative patterns of the kelp-urchin behavioral feedback:

$$h\left(y_t\right)=\alpha y_t\exp\left(-\beta y_t\right). \tag{18}$$

The abiotic conditions X_t and v_t were modeled as autoregressive processes (AR-1) with equal variance. The model parameters were selected to produce time series that regularly flickered between the kelp forest and

urchin barren states. We simulated 100 time series of length 55 using the first 50 observations as a training set and the next five as a testing set. Model parameters used to simulated the datasets are listed in Supplementary Table 3.

The UDE models used the same structure as the simulation model (Eq. 17), but we replaced the urchin grazing rate with a NN and applied a log-transform:

$$\log(u_{t+1}) = \log(u_t) + r\left(1 - \frac{u_t}{K}\right) - NN(\log(u_t), X_t). \quad (19)$$

We evaluated Eq. 19 with and without environmental covariates to determine the influence of known abiotic conditions on model performance. The hyperparameters used to fit the model are listed in Supplementary Table 4. In addition to this UDE model, we fit MARSS models, gpEDMs, and continuous-time nonparametric UDE models (NODEs) trained with each of the four UDE fitting procedures. Detailed descriptions of the alternative models are given in Supplementary Methods 2.

We evaluated the urchin functional responses predicted by the UDEs by integrating the squared difference between the true and predicted functions over the range of biologically meaningful values, which we defined from zero kelp abundance to 1.5 times the kelp carrying capacity

$$\text{score} = \sqrt{\frac{2}{3K}\int_0^{3K/2} \left(h(u) - NN(\log(u))\right)^2 du}. \quad (20)$$

We calculated the forecasting skill of each model by forecasting from the end of the training set to the end of the testing set. We calculated the MAE and MSE between the forecasts and the testing data along with 95% bootstrapped confidence intervals of these metrics using the Hmisc library.

Example 3: Fisheries empirical study

We modeled changes in the biomass of the harvest population B with a logistic growth model with growth rate r, carrying capacity K, and harvest H:

$$\frac{dB}{dt} = rB\left(1 - \frac{B}{K}\right) - qH. \quad (21)$$

where q is used to ensure the units of harvest and stock biomass match. We modeled changes in the rate of harvest as a function of changes in the biomass of the population and changes in fishing mortality rate U, which is equal to the fraction of the stock harvested per unit time $U = qH/B$. We used a NN to describe the changes in fishing mortality over time as a nonlinear function of the stock and an indicator variable $I_{t>1992}$ for the regulatory changes that began in 1992:

$$\frac{dU}{dt} = U \times NN(B,\ I_{t>1992}). \quad (22)$$

We found that multiplying the NN output by the current fishing mortality rate, so that the NN represents the proportion change in fishing mortality, produced better predictions in cross-validation tests. Finally, we derived a model for harvest using the product rule:

$$\frac{dH}{dt} = \frac{dB}{dt}\frac{U}{q} + \frac{dU}{dt}\frac{B}{q} = \frac{dB}{dt}\frac{U}{q} + NN(B, I_{t>1992})H. \quad (23)$$

We fit separate models to the chilipepper rockfish and cowcod abundance and harvest time series. For each species, we compared the forecasting ability of the parametric UDE model trained with each of the four algorithms to a linear state-space MARSS model, a Gaussian process empirical dynamic model, and nonparametric UDE (NODE), described in detail in Supplementary Methods 3. Forecasting skill was evaluated using leave-future-out cross-validation with 15 testing sets. Hyperparameters for the NODE and UDE models were selected by cross validation (Supplementary Figs. 7, 8) and are listed in Supplementary Table 5.

Example 4: Jornada Experimental Range empirical study

We analyzed long-term monitoring data from permanent quadrats in the Jornada Experimental Range, sourced from Christensen et al. (2021). The data include observations from 89 quadrats spanning 71 years with records of 126 species. We restricted our analysis to quadrats with high initial densities of the perennial grass *B. eriopoda*. These quadrats were selected by a cluster analysis (Supplementary Figs. 9, 10). The remaining quadrats exhibited distinct dominant species groups, temporal dynamics, and soil textures. We chose to restrict our analysis to three species groups: *B. eriopoda*, shrubs that were primarily honey mesquite (*Prosopis glandulosa*), and others, because species other than *B. eriopoda* and shrubs including *P. glandulosa* occurred infrequently and in low abundances (Supplementary Table 6, Supplementary Figs. 9–11). A detailed explanation of the data preprocessing is provided in Supplementary Methods 4, including an interpolation procedure used to estimate the abundance of each species group in quadrats that were missing in a given sampling period (Supplementary Fig. 12).

The state-space UDE models for the Jornada Experimental Range each tracked the abundances of the dominant perennial grass species *B. eriopoda* X_g, shrubs X_s, and other species X_r. The change in abundance of species group i, X_i, included a term for growth into empty space $g_i(\boldsymbol{X}, P, T)$ that depended on the abundance of each species group $\boldsymbol{X}$, precipitation P, temperature T, and a constant mortality rate m_i. These effects are captured by a system of ODEs:

$$\frac{dX_i}{dt} = g_i(\boldsymbol{X}, P, T)\left(A - X_g - X_s - X_r\right) - m_i X_i. \quad (24)$$

where A is the total area of space described by the model.

The competition and plant-soil feedback UDE models differed in how they constructed the function $g_i(\boldsymbol{X}, P, T)$. The competition UDE used a NN to capture species interactions and linear terms for precipitation and temperature. The final value was transformed by the softplus function $\log(1 + \exp(x))$ to ensure the growth rate always had a positive value:

$$g_i(\boldsymbol{X}, P, T) = \text{softplus}\left(NN(\boldsymbol{X}) + \beta_{i,P}P + \beta_{i,T}T\right). \quad (25)$$

The plant-soil feedback model explicitly incorporated the effect of plant community composition on soil moisture by adding two additional dynamic variables to the model: the quantity of homogenously-distributed soil moisture M and the quantity of heterogeneously-distributed soil moisture H. The quantities of soil moisture were determined by the rate at which precipitation infiltrates the soil minus the rate of evaporation. We assumed that water either infiltrates the soil directly and enters the homogenous pool M or concentrates on the surface and enters the heterogenous pool H. We modeled the rate at which rainfall enters the homogenous pool using a NN that takes the abundance of each species group as inputs and applies the logistic transform to the output to obtain a value between zero and one. Given these assumptions, the levels of the soil moisture pools can be expressed with a system of ODEs:

$$\frac{dM}{dt} = NN(\boldsymbol{X})P - (E_{M,0} + \xi_{\mathrm{M}}\mathrm{T})M \quad (26)$$

$$\frac{dH}{dt} = (1 - NN(\boldsymbol{X}))P - (E_{H,0} + \xi_{\mathrm{H}}\mathrm{T})H, \quad (27)$$

where evaporation is a linear function of temperature with slope ξ_j and intercept $E_{j,0}$. The growth rate of each species group is proportional to the

sum of each soil moisture pool weighted by $r_{i,j}$, where i indexes over species groups and j indexes over the two soil moisture pools:

$$g_i(\boldsymbol{X}, P, T) = \left(r_{i,M}M + r_{i,H}H\right). \tag{28}$$

We assumed that the dominant grass species *B. eriopoda* only benefits from the homogenous soil moisture and therefore fix $r_{g,H} = 0$. We chose to approximate the values of M and H with their equilibrium values. This is often called a fast-slow approximation and assumes the dynamics of soil moisture stocks are fast relative to the other state variables, allowing us to derive a model that only depends on the current values of the plant community composition $\boldsymbol{X}$, temperature T, and precipitation P (Supplementary Methods 4). We fit both models using the five-year running averages of temperature and precipitation following a prior study in this system which found that smoothed climate variables were better predictors of *B. eriopoda* abundances than raw climate variables (Christensen et al. 2023). Hyperparameters were slected using cross validation (Supplementary Figs. 13, 14) and are listed in Supplementary Table 7. We compared the two UDE models trained with all four algorithms to gpEDM and MARSS models using leave-future-out cross-validation with 20 testing sets. We repeated the leave-future-out cross-validation tests twice, once using the full dataset constructing the testing sets from the end of the training set forward. However, this procedure did not include the regime change that occurred in the 1950s in any of the testing sets. To evaluate the model's ability to forecast during the regime change, we repeated the cross-validation tests on a dataset ending in the 1960s.

Example 5: Metacommunity predator-prey simulation

The primary goal of this dataset and analysis was to test the computational costs of scaling the UDE models to datasets with large numbers of state variables, a goal that does not necessarily require an empirical or even a realistic simulated dataset. However, training UDE models can depend on dynamics in the dataset because the training algorithm relies on an adaptive ODE solver. Because of this, we chose to simulate data from a predator-prey model with several patches connected by dispersal. This allowed us to simulate biologically plausible time series with an arbitrary number of state variables simply by adding additional patches to the simulation model. The model describes the abundance of a predator species P and prey species N. The prey population grows according to the logistic model with growth rate r in the absence of the predator and is consumed by the predator according to a type II functional response with attack rate α and handling time h. New predators are born at a rate proportional to prey consumption by a factor ϕ, have morality rate m, and disperse between patches according to dispersal matrix D (Supplementary Fig. 15) where each element $d_{i,j}$ describes the rate of dispersal from patch i to patch j. We assume both populations experience multiplicative shocks W to their growth rate with variance τ^2

$$\frac{dN_i}{dt} = rN_i\left(1 - N_i\right) - \frac{\alpha N_i P_i}{1 + hN_i} + \tau N_i dW_t \tag{29}$$

$$\frac{dP_i}{dt} = \frac{\alpha\phi N_i P_i}{1 + hN_i} - mP - \sum_{j\neq i} d_{i,j}P_j + \sum_{j\neq i} d_{j,i}P_i + \tau P dW_t. \tag{30}$$

We trained two models on these simulated datasets: a fully nonparametric model (NODE) that used each state variable as an input to a NN

$$\frac{d\boldsymbol{u}}{dt} = NN(\boldsymbol{u}, \theta), \tag{31}$$

and a model that used a NN to describe the local interactions between the predator and prey and used the parametric model for dispersal

$$\frac{dN_i}{dt} = NN_N\left(N_i, P_i, \theta\right) \tag{32}$$

$$\frac{dP_i}{dt} = NN_P\left(N_i, P_i, \theta\right) - \sum_{j\neq i} d_{i,j}P_j + \sum_{j\neq i} d_{j,i}P_i. \tag{33}$$

To estimate the models' training times, we simulated datasets of length 30 and trained the models for 20 iterations each and then multiplied by 50 to estimate the time required to train the model for 1000 iterations. Parameters for the simulations are given in Supplementary Table 8. We estimated the scaling by fitting a simple linear regression model to a log-log plot of the number of state variables and the training time. Note that we do not compare the timing to gpEDM and MARSS models for this example, because our aim is to understand how the state-space method influences training times when compared to other UDE methods.

Data availability

The data used in this paper were obtained from two publicly accessible data repositories: The Jornada Experimental Range LTER (https://doi.org/10.6073/pasta/63bfa45df4858db674bf37b52ee5ff44) and the RAM Legacy Stock Assessment Database https://doi.org/10.17616/R34D2X.

Code availability

Software to implement the State-space universal dynamic equation method presented in this papers has been published by the authors on the Julia programming language registry (https://github.com/JuliaRegistries/General) and on GitHub (https://github.com/Jack-H-Buckner/UniversalDiffEq.jl).

Received: 16 October 2024; Accepted: 11 December 2025;
Published online: 31 January 2026

References

1. Suding, K. N. & Hobbs, R. J. Threshold models in restoration and conservation: a developing framework. *Trends Ecol. Evol.* **24**, 271–279 (2009).
2. Perretti, C. T., Sugihara, G. & Munch, S. B. Nonparametric forecasting outperforms parametric methods for a simulated multispecies system. *Ecology* **94**, 794–800 (2013).
3. Arroyo-Esquivel, J., Klausmeier, C. A. & Litchman, E. Using neural ordinary differential equations to predict complex ecological dynamics from population density data. *J. R. Soc. Interface* **21**, 20230604 (2024).
4. Bonnaffé, W., Sheldon, B. C. & Coulson, T. Neural ordinary differential equations for ecological and evolutionary time-series analysis. *Methods Ecol. Evol.* **12**, 1301–1315 (2021).
5. Rackauckas, C. et al. Universal differential equations for scientific machine learning. Preprint at. http://arxiv.org/abs/2001.04385 (2021).
6. Cheridito, P., Jentzen, A. & Rossmannek, F. Efficient approximation of high-dimensional functions with neural networks. *IEEE Trans. Neural Netw. Learn. Syst.* **33**, 3079–3093 (2022).
7. Bonnaffé, W. & Coulson, T. Fast fitting of neural ordinary differential equations by Bayesian neural gradient matching to infer ecological interactions from time-series data. *Methods Ecol. Evol.* **14**, 1543–1563 (2023).
8. Vaswani, A. et al. Attention is all you need. Preprint at. http://arxiv.org/abs/1706.03762 (2023).
9. Hochreiter, S. & Schmidhuber, J. Long short-term memory. *Neural Comput.* **9**, 1735–1780 (1997).
10. Turan, E. M. & Jäschke, J. Multiple shooting for training neural differential equations on time series. *IEEE Control Syst. Lett.* **6**, 1897–1902 (2022).
11. Auger-Méthé, M. et al. A guide to state–space modeling of ecological time series. *Ecol. Monogr.* **91**, e01470 (2021).
12. Hastings, A. & Powell, T. Chaos in a three-species food chain. *Ecology* **72**, 896–903 (1991).

13. RAM Legacy Stock Assessment Database. RAM Legacy Stock Assessment Database v4 Zenodo https://doi.org/10.5281/ZENODO.2542919 (2018).
14. Christensen, E. et al. Quadrat-based monitoring of desert grassland vegetation at the Jornada Experimental Range, New Mexico, 1915–2016. *Ecology* **102**, e03530 (2021).
15. Medeiros, L. P., Sorenson, D. K., Johnson, B. J., Palkovacs, E. P. & Munch, S. B. Revealing unseen dynamical regimes of ecosystems from population time-series data. *Proc. Natl. Acad. Sci. USA*. **122**, e2416637122 (2025).
16. Munch, S. B., Rogers, T. L. & Sugihara, G. Recent developments in empirical dynamic modelling. *Methods Ecol. Evol.* **14**, 732–745 (2023).
17. Munch, S. & Rogers, T. GPEDM: Gaussian process regression for empirical dynamic modeling. https://tanyalrogers.github.io/GPEDM/reference/GPEDM-package.html (2024).
18. Holmes, E., E., Ward, E., J. & Wills, K. MARSS: multivariate autoregressive state-space models for analyzing time-series data. *R J.* **4**, 11 (2012).
19. Ellner, S. P., Seifu, Y. & Smith, R. H. Fitting population dynamics models to time-series data by gradient matching. *Ecology* **83**, 2256–2270 (2002).
20. Rosenzweig, M. L. & MacArthur, R. H. Graphical representation and stability conditions of predator-prey interactions. *Am. Nat.* **97**, 209–223 (1963).
21. Wan, E. A. & Van Der Merwe, R. The unscented Kalman filter for nonlinear estimation. in *Proc. IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No.00EX373)* 153–158. https://doi.org/10.1109/ASSPCC.2000.882463 (IEEE, 2000).
22. Estes, J. A. & Duggins, D. O. Sea otters and kelp forests in Alaska: generality and variation in a community ecological paradigm. *Ecol. Monogr.* **65**, 75–100 (1995).
23. Konar, B. & Estes, J. A. The stability of boundary regions between kelp beds and deforested areas. *Ecology* **84**, 174–185 (2003).
24. Rogers-Bennett, L. & Catton, C. A. Marine heat wave and multiple stressors tip bull kelp forest to sea urchin barrens. *Sci. Rep.* **9**, 15050 (2019).
25. Harrold, C. & Reed, D. C. Food availability, sea urchin grazing, and kelp forest community structure. *Ecology* **66**, 1160–1169 (1985).
26. Wilen, J. E. Common property resources and the dynamics of overexploitation: the case of the North Pacific Fur Seal. *Mar. Resour. Econ.* **33**, 217–243 (2018).
27. Smith, V. L. On models of commercial fishing. *J. Polit. Econ.* **77**, 181–198 (1969).
28. Free, C. M. et al. Harvest control rules used in US federal fisheries management and implications for climate resilience. *Fish Fish.* https://doi.org/10.1111/faf.12724 (2022).
29. Walters, C. J. & Hilborn, R. Adaptive control of fishing systems. *J. Fish. Res. Board Can.* **33**, 145–159 (1976).
30. Pacific Fisheries Management Council. *Pacific Coast Goundfish Fishery Management Plan*, 159 (NOAA Fisheries, Pacific Fisheries Management Council, 2019).
31. Dick, E. J. & He, X. *Status of Cowcod (Sebastes Levis) in 2019* (NOAA Fisheries, 2019).
32. Field, J. C., Beyer, S. G. & He, X. *Status of the Chilipepper Rockfish, Sebastes Goodei, in the California Current for 2015* (NOAA Fisheries, 2015).
33. Briske, D. D., Fuhlendorf, S. D. & Smeins, F. E. State-and-transition models, thresholds, and rangeland health: a synthesis of ecological concepts and perspectives. *Rangel. Ecol. Manag.* **58**, 1–10 (2005).
34. Briske, D. D., Fuhlendorf, S. D. & Smeins, F. E. Vegetation dynamics on rangelands: a critique of the current paradigms. *J. Appl. Ecol.* **40**, 601–614 (2003).
35. Schlesinger, W. H. et al. Biological feedbacks in global desertification. *Science* **247**, 1043–1048 (1990).
36. D'Odorico, P. et al. Vegetation–microclimate feedbacks in woodland–grassland ecotones. *Glob. Ecol. Biogeogr.* **22**, 364–379 (2013).
37. Christensen, E. M., James, D. K., Randall, R. M. & Bestelmeyer, B. T. Abrupt transitions in a southwest USA desert grassland related to the Pacific Decadal Oscillation. *Ecology* **104**, e4065 (2023).
38. Treven, L., Wenk, P., Dörfler, F. & Krause, A. Distributional gradient matching for learning uncertain neural dynamics models. Preprint at. https://doi.org/10.48550/arXiv.2106.11609 (2021).
39. Boettiger, C., Ross, N. & Hastings, A. Early warning signals: the charted and uncharted territories. *Theor. Ecol.* **6**, 255–264 (2013).
40. Scheffer, M. et al. Early-warning signals for critical transitions. *Nature* **461**, 53–59 (2009).
41. Tilman, A. R., Krueger, E. H., McManus, L. C. & Watson, J. R. Maintaining human wellbeing as socio-environmental systems undergo regime shifts. *Ecol. Econ.* **221**, 108194 (2024).
42. Dandekar, R. et al. Bayesian neural ordinary differential equations. Preprint at. http://arxiv.org/abs/2012.07244 (2022).
43. Doucet, A., Godsill, S. & Andrieu, C. On sequential Monte Carlo sampling methods for Bayesian filtering. in *Readings In Unobserved Components Models* (eds Harvey, A. C. & Proietti, T.) 418–441.https://doi.org/10.1093/oso/9780199278657.003.0022 (Oxford University PressOxford, 2005).
44. Langendorf, R. E. et al. Dynamic and context-dependent keystone species effects in kelp forests. *Proc. Natl. Acad. Sci. USA*. **122**, e2413360122 (2025).
45. Meunier, Z. D., Hacker, S. D. & Menge, B. A. Regime shifts in rocky intertidal communities associated with a marine heatwave and disease outbreak. *Nat. Ecol. Evol.* **8**, 1285–1297 (2024).
46. Watson, J. R., Kendall, B. E., Siegel, D. A. & Mitarai, S. Changing seascapes, stochastic connectivity, and marine metapopulation dynamics. *Am. Nat.* **180**, 99–112 (2012).
47. Kingma, D. P. & Ba, J. Adam: a method for stochastic optimization. Preprint at. http://arxiv.org/abs/1412.6980 (2014).
48. Dixit, V. K. & Rackauckas, C. Optimization.jl: A Unified Optimization Package. https://doi.org/10.5281/ZENODO.7738525 (2023).
49. Bezanson, J., Edelman, A., Karpinski, S. & Shah, V. B. Julia: a fresh approach to numerical computing. *SIAM Rev.* **59**, 65–98 (2017).
50. Innes, M. Don't unroll adjoint: differentiating SSA-form programs. Preprint at. http://arxiv.org/abs/1810.07951 (2019).
51. Tsitouras, C. Runge–Kutta pairs of order 5(4) satisfying only the first column simplifying assumption. *Comput. Math. Appl.* **62**, 770–775 (2011).
52. Rackauckas, C. & Nie, Q. DifferentialEquations.jl – a performant and feature-rich ecosystem for solving differential equations in Julia. *JORS* **5**, 15 (2017).
53. Bhagavan, S., De Koning, B., Maddhashiya, S. & Rackauckas, C. DataInterpolations.jl: fast interpolations of 1Ddata. *JOSS* **9**, 6917 (2024).
54. Stickel, J. J. Data smoothing and numerical differentiation by a regularization method. *Comput. Chem. Eng.* **34**, 467–475 (2010).
55. Golub, G. H., Heath, M. & Wahba, G. Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics* **21**, 215–223 (1979).
56. Harrell, F. E. Jr Hmisc. *Harrell Misc.* **5**, 2–3 (2003).

Acknowledgements

This research was supported by the National Science Foundation awards #2233982 and #2233983 to JRW and LCM on Model Enabled Machine Learning for Predicting Ecosystem Regime Shifts. This paper is a product of the model-enabled machine learning for ecology working group, which includes the authors of the paper as well as Cheyenne Jarman, Kunal Rathore, and Emerson Arehart, all of whom provided valuable contributions to the intellectual environment that led to this paper. We would also like to

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.

Author contributions

J.H.B.: Conceptualization, formal analysis, methodology, software, writing-original draft. Z.D.M.: Conceptualization, methodology, software, writing-review and editing. J.A.E.: Conceptualization, funding acquisition, software, writing-review and editing. N.F.: methodology, software, writing-review and editing. A.G.: Conceptualization, writing-review and editing. L.C.M.: Funding acquisition, project administration, writing-review and editing. JRW: Conceptualization, funding acquisition, project administration, writing-review and editing.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s43247-025-03130-2.

Correspondence and requests for materials should be addressed to Jack H. Buckner.

Peer review information *Communications Earth and Environment* thanks the anonymous reviewers for their contribution to the peer review of this work. Primary Handling Editors: Lifen Jiang and Mengjie Wang. [A peer review file is available].