Predator-Prey Dynamics Modeling Using Neural SDEs and UDEs: SciML Modeling in Theoretical and Real-World Ecological Systems

John Trixie M. Ocampo[1], Romie C. Mabborang[2]
jtmocampo2021@plm.edu.ph[1], rcmabborang@plm.edu.ph[2]
Pamantasan ng Lungsod ng Maynila (University of the City of Manila)
Manila, Philippines

Abstract: *Accurately modeling predator-prey dynamics is vital for ecological understanding, but traditional models often fall short in capturing real-world complexity and randomness. Scientific Machine Learning (SciML) offers hybrid approaches, merging mechanistic knowledge with data-driven techniques. This study provides a comprehensive comparison of two such methods: Neural Stochastic Differential Equations (SDEs) and Universal Differential Equations (UDEs). Both frameworks were applied to model predator-prey interactions using noisy synthetic Lotka-Volterra data and empirical algae-rotifer time series. We embedded a learnable Lotka-Volterra structure within each framework, augmented by neural networks designed to capture model discrepancies or learn drift and diffusion corrections. Model performance was evaluated based on fitting accuracy, forecasting ability, parameter recovery, and the analysis of learned neural components. Results show both UDEs and Neural SDEs effectively captured the complex oscillatory dynamics in both datasets. The UDE achieved higher deterministic accuracy on the synthetic data, while the Neural SDE demonstrated more robust forecasting performance on the challenging empirical data. Parameter recovery proved difficult for both methods, with neural networks learning significant dynamics corrections. This research offers practical insights into the relative strengths and trade-offs of applying Neural SDEs and UDEs in ecological modeling.*

1. Introduction

1.1 Background of the Study

In an era defined by rapid environmental change and escalating biodiversity concerns, the need for accurate and robust models of ecological systems is paramount [1]. Under-standing core ecological processes, such as predator-prey dynamics, is essential for predicting population trajectories, assessing ecosystem stability, and informing effective conservation strategies [2]. While foundational models like the Lotka-Volterra equations provided initial mathematical frameworks based on Ordinary Differential Equations (ODEs) [3, 4], their inherent simplicity often falls short. These classic deterministic models typically neglect complex functional responses, environmental drivers, spatial heterogeneity, and the pervasive stochasticity that characterize real-world ecological interactions [5]. As ecological data becomes increasingly available through advanced monitoring, developing models that can effectively utilize this data while respecting biological principles is a key challenge.

Despite the theoretical appeal of these advanced methods, their practical performance and nuances in ecological applications warrant careful examination.

1.2 Research Questions

This study aims to systematically evaluate the performance of Neural SDEs and Neural UDEs on both controlled synthetic data generated from the Lotka-Volterra equations with added noise, and on real-world empirical data from an algae-rotifer chemostat experiment. Specifically, this research addresses the following questions:

1. How well do the trained Neural UDE and Neural SDE models replicate dynamics on both synthetic Lotka-Volterra and empirical algae-rotifer data considering both model fit and forecasting ability?
2. How accurate can the models recover known parameters of the underlying Lotka-Volterra system when trained on noisy synthetic data and empirical data?
3. What is the structure and magnitude of the learned neural network components relative to the known Lotka-Volterra components?
4. What are the practical challenges and trade-off associated with implementing and training Neural UDEs versus Neural SDEs for each dataset?

2. Methodology

This section presents the development and comparison of four distinct predator-prey models: a Neural Universal Differential Equation (Neural UDE) and a Neural Stochastic Differential Equation (Neural SDE) applied to both synthetic Lotka-Volterra data and empirical algae-rotifer data.

3.1. Data Collection

3.1.1. Synthetic Data

Synthetic predator-prey time-series data was generated using the classical Lotka-Volterra equations from Eq. (1):

$$\frac{dx}{dt} = \alpha x - \beta x y \tag{1}$$

$$\frac{dy}{dt} = \delta x y - \gamma y$$

where x represents prey population and y represents predator population. The parameters were set to $\alpha = 1.5$, $\beta = 1.0$, $\gamma = 2.0$ and $\delta = 0.5$. The system was integrated numerically over the time interval $t \in [0,10]$ with initial conditions $x(0) = 8.0$ and $y(0) = 3.0$ with Gaussian Noise $\sigma = 0.1$. A total of $N = 200$ data points were sampled uniformly across this time span.

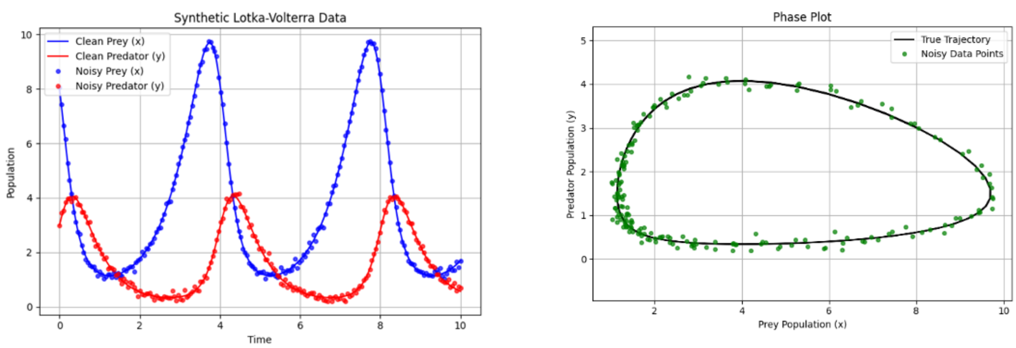

Figure 1. Generated Lotka-Volterra Synthetic Data: (**a**) Population-Time Graph; (**b**) Phase Plane

3.1.2. Empirical Data

Time series data is employed capturing the population densities of green algae (*Chlorella vulgaris*) and rotifers (*Brachionus calyciflorus*) interacting within a laboratory chemostat, based on experiments conducted by Blasius et al. and Fussmann et al. The specific dataset was set for 16 days and contains 200 interpolated data points for temporal density.

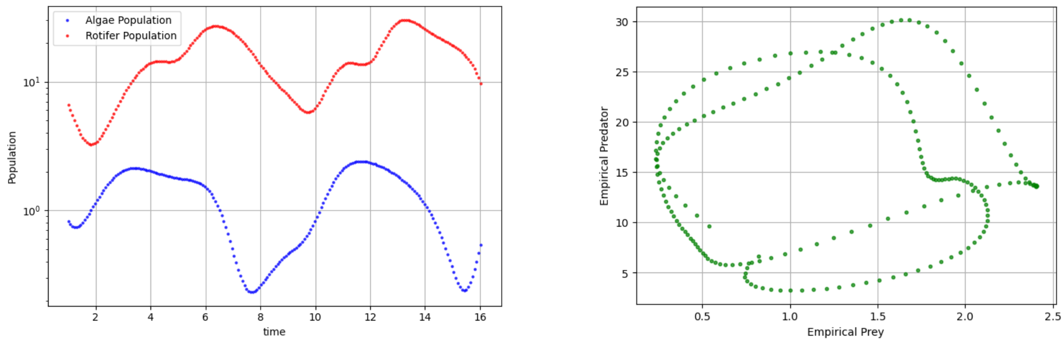

Figure 2. Algae-Rotifer Empirical Data: (**a**) Population-Time Graph; (**b**) Phase Plane

3.2 Data Preparation and Model Training

For both synthetic and empirical datasets, the time series data $(\boldsymbol{t}, \boldsymbol{u_{data}})$ was split chronologically into a training set (first 80% of points) and a test set (remaining 20%). Let **N** be the total number of points, the split index was $\boldsymbol{N_{train}} = \lfloor \mathbf{0.8}\boldsymbol{N} \rfloor$. The training set $\left(\boldsymbol{t_{train}}, \boldsymbol{u_{train}^{orig}}\right)$ was used to train the models, while the test set $\left(\boldsymbol{t_{test}}, \boldsymbol{u_{test}^{orig}}\right)$ was reserved for evaluating forecasting performance.

All four models were trained using their respective training datasets ($\boldsymbol{t_{train}}\ \boldsymbol{u_{train}^{scaled}}$ for UDEs, or $\boldsymbol{t_{train}}, \boldsymbol{u_{train}^{scaled}}$ target with SDE solving on $\boldsymbol{t_{train}}$ from $\boldsymbol{u_0^{scaled}}$ for SDEs). The objective was to minimize the Mean Squared Error (MSE) between the model's predicted trajectory (at the training time points) and the observed scaled training data $\boldsymbol{u_{train}^{scaled}}$. The Adam optimizer was used for training all models. Specific hyperparameters and solver settings varied slightly per model, primarily based on the provided code snippets, and are summarized in Table 1.

Table 1. Training Hyperparameters and Solver Settings per Model

Parameter/Setting	**SDE Synthetic Data**	**UDE Synthetic Data**	**SDE Empirical**	**UDE Empirical**
Learning Rate	1×10^{-3}	1×10^{-2}	1×10^{-3}	1×10^{-3}
Epochs	5000	3000	5000	5000
Weight Decay	1×10^{-6}	1×10^{-6}	1×10^{-6}	1×10^{-6}
Hidden Dim (NN)	32	32	32	32
Optimizer	Adam	Adam	Adam	Adam
Solver Library	torchsde	torchdiffeq	torchsde	torchdiffeq
Solver Method	euler	Dopri5	euler	Dopri5
Parameter/Setting	SDE Synthetic Data	UDE Synthetic Data	SDE Empirical	UDE Empirical

3.3 Model Evaluation

Model performance was evaluated using the saved best model state for each of the four cases. The evaluation involved generating predictions over the entire time span $\boldsymbol{t_{data}}$ (covering both training and test periods) starting from the initial condition $\boldsymbol{u_0^{scaled}}$. All predicted trajectories, initially in the scaled space, were inverse-transformed back to the original population scale using the scaler fitted on the training data. These original-scale predictions $\left(\boldsymbol{u_{pred}^{orig}}\right)$ were then compared against the original-scale true data $\left(\boldsymbol{u_{data}^{orig}}\right)$.

Goodness-of-fit and forecasting accuracy were quantified using several standard metrics, calculated separately for the training period (fit performance) and the test period (forecast performance). Let $\boldsymbol{u_i^{orig}} = \left[\boldsymbol{x_i^{orig}}, \boldsymbol{y_i^{orig}}\right]^{\boldsymbol{T}}$ be the true data point at time $\boldsymbol{t_i}$ and $\widehat{\boldsymbol{u_i^{orig}}} = \left[\widehat{\boldsymbol{x_i^{orig}}}, \widehat{\boldsymbol{y_i^{orig}}}\right]^{\boldsymbol{T}}$ be the corresponding model prediction (mean path for SDEs) in the original scale. The metrics calculated were Mean Squared Error (MSE), Mean Absolute Error (MAE), and Coefficient of Determination (R^2).

3. Results and Discussion

This section presents the results obtained from training and evaluating the Neural UDE and Neural SDE models on both synthetic and empirical predator-prey datasets. We analyze the model performance based on quantitative metrics, visual inspection of predicted trajectories, learned parameters, and the structure of the learned dynamic components.

3.1 Modeling Synthetic Lotka-Volterra Data

3.1.1 Neural SDE Performance on Synthetic Data

Figure 3 shows both the population-time graph and phase plane with both time-series fit and forecast, excellently capturing the behavior of the graph.

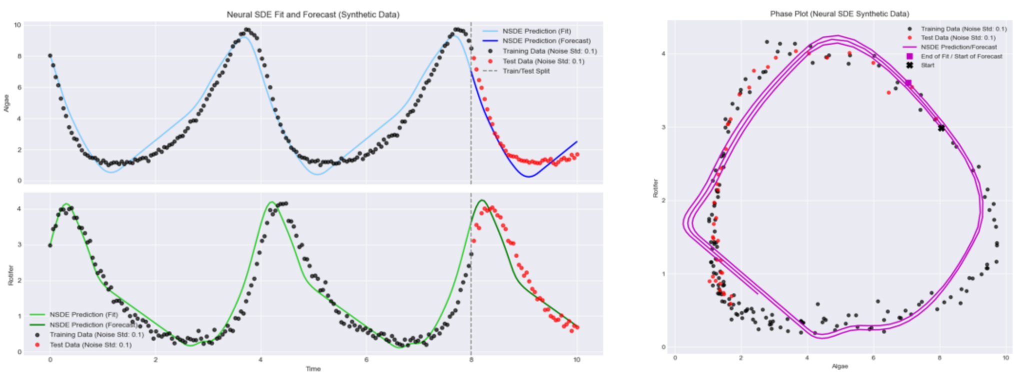

(a) **(b)**

Figure 3. Neural SDE Fit and Forecast on Synthetic Data (**a**) Population-Time Graph; (**b**) Phase Plane

Performance metrics calculated on the original data scale are presented in Table 2. The model achieved a good fit on the training data (e.g., $\boldsymbol{R^2_{total} = 0.9464}$) and reasonable forecasting performance on the test set $\left(\boldsymbol{R^2_{total} = 0.9198}\right)$.

Table 2. Performance Metrics for Neural SDE on Synthetic Data

Metric (Train Set – Fit)	**Prey/Algae**	**Predator/Rotifer**	**Total**
MSE	0.4242	0.0867	0.2555
RMSE	0.6513	0.2945	0.5054
R^2	0.9486	0.9441	0.9464
Metric (Test Set - Forecast)	**Prey/Algae**	**Predator/Rotifer**	**Total**
MSE	0.3072	0.0970	0.2021
RMSE	0.5543	0.3114	0.4496
R^2	0.9057	0.9340	0.9198

Analysis of the learned dynamics components in the scaled space (Fig. 4) reveals that the $\boldsymbol{NN_{drift}}$ term provides substantial corrections to the $\boldsymbol{f_{known}}$ (parameterized by the inaccurately learned parameters). The learned diffusion term $\left(\boldsymbol{g = NN_{diffusion}}\right)$ is non-zero, indicating the model captured some stochasticity.

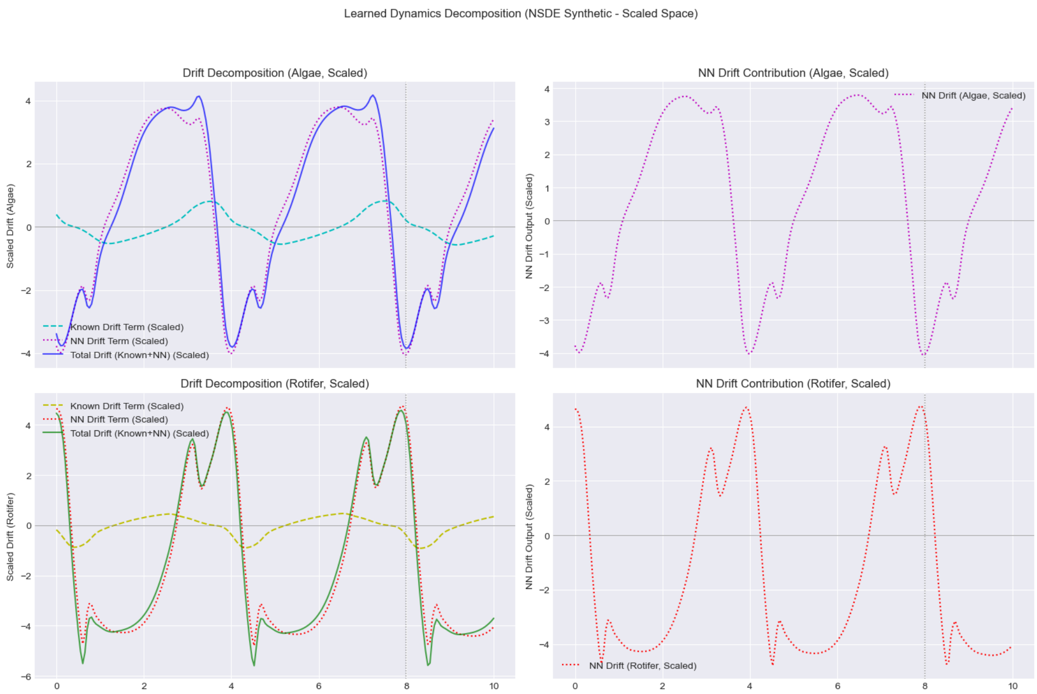

Figure 4. Learned Dynamics Decomposition for Neural SDE on Synthetic Data

4.1.2. Neural UDE Performance on Synthetic Data

Figure 5 shows the UDE prediction closely tracking the underlying structure within the noisy data points.

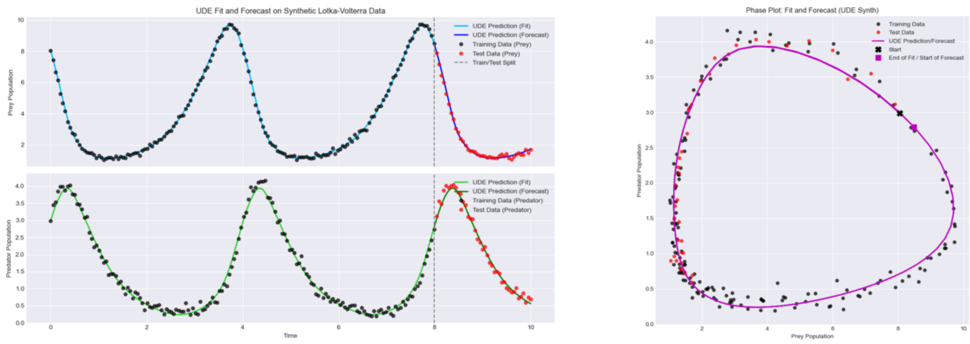

(a) **(b)**

Figure 5. Neural UDE Fit and Forecast on Synthetic Data (**a**) Population-Time Graph; (**b**) Phase Plane

Metrics (Table 3) indicate an excellent fit to the training data $\left(\boldsymbol{R^2_{total}} = \mathbf{0.9940}\right)$ and strong forecasting performance $\left(\boldsymbol{R^2_{total}} = \mathbf{0.9928}\right)$ surpassing the SDE model on this noisy synthetic dataset in terms of deterministic accuracy.

Table 3. Performance Metrics for Neural UDE on Synthetic Data

Metric (Train Set – Fit)	**Prey/Algae**	**Predator/Rotifer**	**Total**
MSE	0.0127	0.0161	0.0144
RMSE	0.1126	0.1269	0.1200
R^2	0.9985	0.9896	0.9940
Metric (Test Set - Forecast)	**Prey/Algae**	**Predator/Rotifer**	**Total**
MSE	0.0170	0.0134	0.0152
RMSE	0.1303	0.1158	0.1233
R^2	0.9948	0.9909	0.9928

The dynamics decomposition plot (Fig. 5) clearly shows the $NN_{unknown}$ term making significant, structured contributions to correct the dynamics generated by the f_{known} term. The UDE dynamics closely match the target behavior.

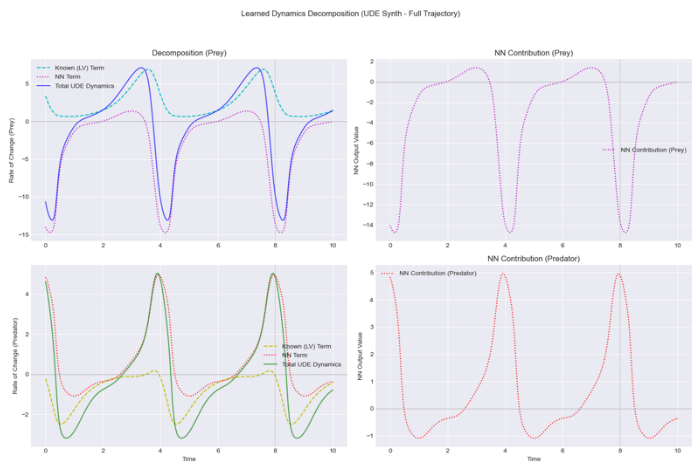

Figure 5: Learned Dynamics Decomposition for Neural UDE on Synthetic Data

4.2. Modeling Empirical Algae-Rotifer Data

4.2.1. Neural SDE Performance on Empirical Data

The population-time graph and phase plane on Fig. 6 shows the mean prediction path capturing the main oscillations.

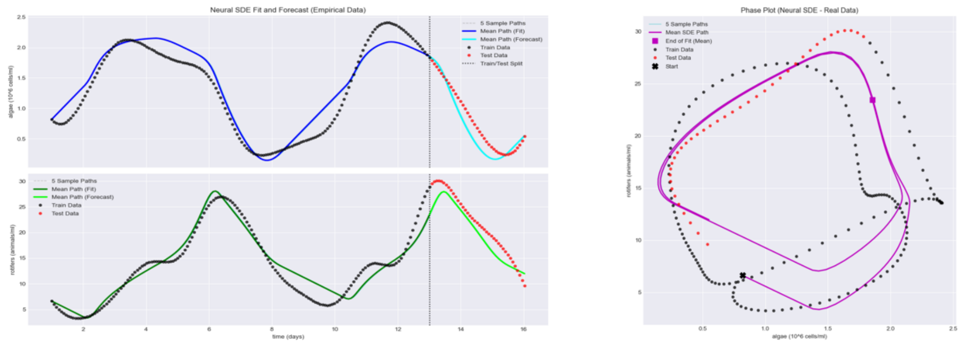

(a) **(b)**

Figure 6. Neural SDE Fit and Forecast on Empirical Data (**a**) Population-Time Graph; (**b**) Phase Plane

Metrics (Table 4) show a reasonably good fit to the training data $\left(\boldsymbol{R^2_{total}} = \mathbf{0.9206}\right)$ and respectable forecasting performance $\left(\boldsymbol{R^2_{total}} = \mathbf{0.8753}\right)$.

Table 4. Performance Metrics for Neural SDE on Empirical Data

Metric (Train Set – Fit)	**Prey/Algae**	**Predator/Rotifer**	**Total**
MSE	0.0397	3.9030	1.9714
RMSE	0.1992	1.9756	1.4041
R^2	0.9209	0.9204	0.9206
Metric (Test Set - Forecast)	**Prey/Algae**	**Predator/Rotifer**	**Total**
MSE	0.0293	4.9487	2.4890
RMSE	0.1711	2.2246	1.5777
R^2	0.8870	0.8635	0.8753

The dynamics decomposition (Fig. 7) shows significant contributions from both $\boldsymbol{NN_{drift}}$ and $\boldsymbol{NN_{diffusion}}$. The average magnitudes calculated indicate that the learned NN drift correction is considerably larger, on average, than the drift from the known LV part (with its learned parameters), highlighting the inadequacy of the simple LV model. The diffusion term, while smaller on average than the $\boldsymbol{NN_{drift}}$, is clearly non-zero and state-dependent.

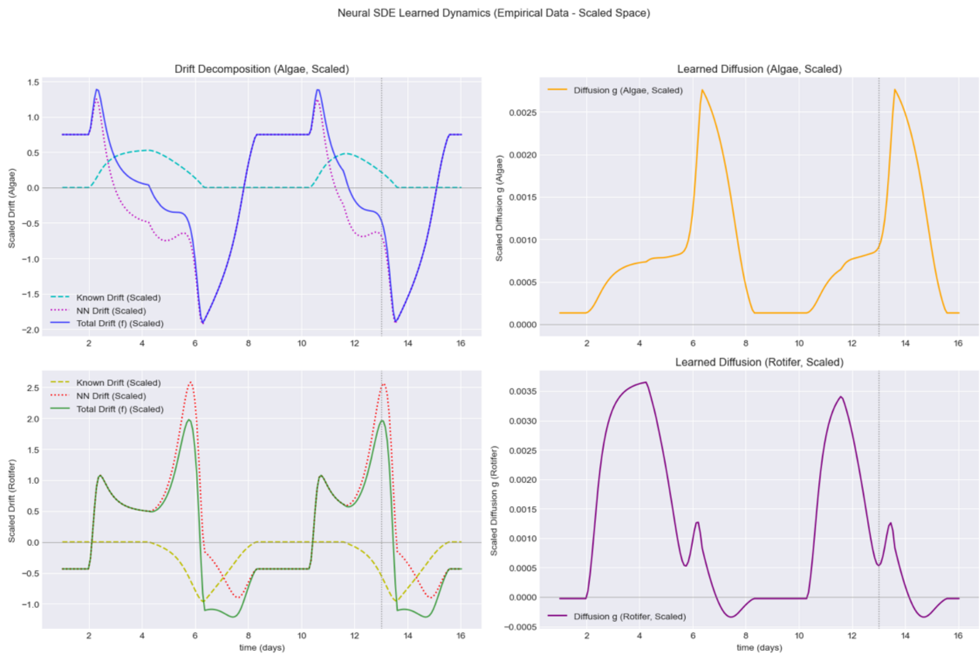

Figure 7. Learned Dynamics Decomposition for Neural SDE on Empirical Data

4.2.2. Neural UDE Performance on Empirical Data

The population-time graph and phase plane show the UDE capturing the training dynamics well, but diverging more significantly in the test phase compared to that of the SDE.

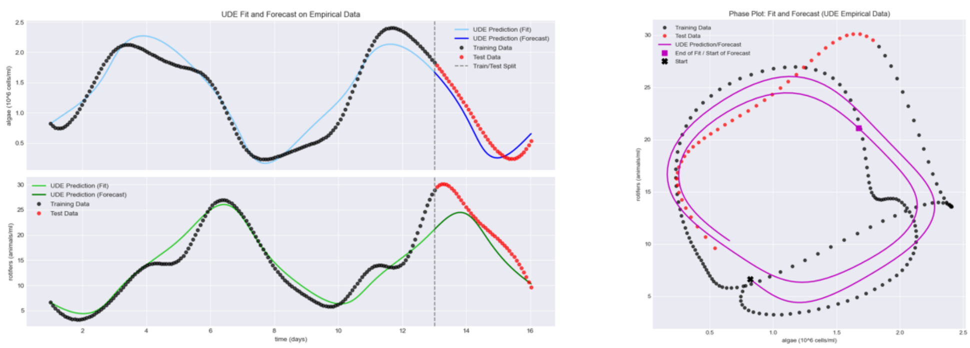

(a) **(b)**

Figure 8. Neural UDE Fit and Forecast on Empirical Data (**a**) Population-Time Graph; (**b**) Phase Plane

Metrics (Table 5) show a good fit on the training data show a good fit on the training data $\left(R^2_{total} = \mathbf{0.9250}\right)$, comparable to the SDE fit. However, the forecasting performance on the test set $\left(R^2_{total} = \mathbf{0.7690}\right)$ is noticeably worse than the SDE's, particularly for the rotifer component $(R^2 = \mathbf{0.6485})$.

Table 5: Performance Metrics for Neural UDE on Empirical Data

Metric (Train Set – Fit)	Prey/Algae	Predator/Rotifer	Total
MSE	0.0352	3.9203	1.9777
RMSE	0.1875	1.9800	1.4063
R^2	0.9299	0.9201	0.9250
Metric (Test Set - Forecast)	**Prey/Algae**	**Predator/Rotifer**	**Total**
MSE	0.0286	12.7469	6.3878
RMSE	0.1692	3.5703	2.5274
R^2	0.8896	0.6485	0.7690

The dynamics decomposition (Fig. 9) again shows a large contribution from the $\boldsymbol{NN_{unknown}}$ term, which learns complex, non-linear corrections to the basic Lotka-Volterra dynamics assumed in $\boldsymbol{f_{known}}$.

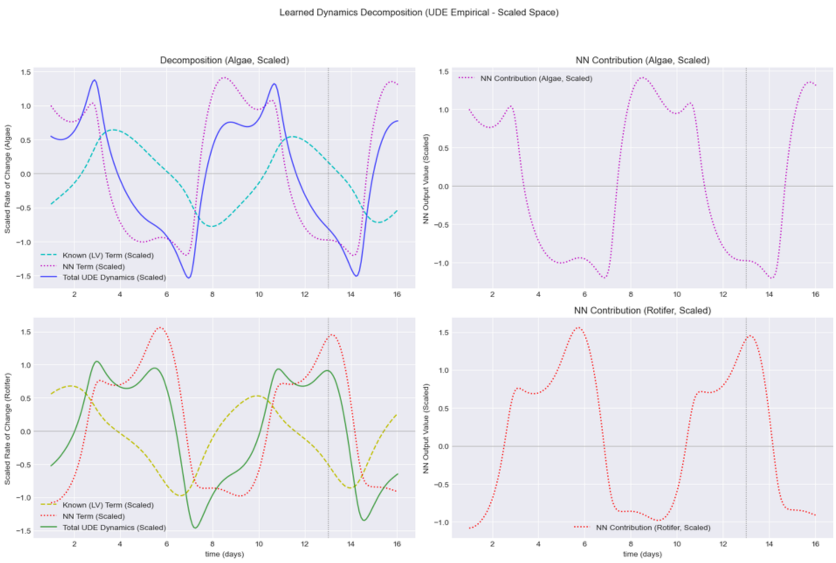

Figure 9. Learned Dynamics Decomposition for Neural UDE on Empirical Data

In both frameworks, the learned Lotka-Volterra parameters were likely non-representative of the true underlying biological rates, as the simple LV model is known to be inadequate for this system (Fussman, 2000). The neural network components in both models learned substantial, complex corrections, effectively dominating the dynamics. This highlights the power of SciML to adapt a potentially incorrect prior model structure to match observations, but also underscores the challenge of interpreting the learned parameters of the 'known' part when the NN correction is large. The SDE framework additionally provided an estimate of the state-dependent noise structure which could be valuable for understanding system variability.

4.3. Summary of Findings

The summary metrics reveal contrasting performance. On synthetic data, the Neural UDE significantly outperformed the Neural SDE, achieving lower RMSE and near-perfect R^2 scores

(>0.99) for both training fit and test forecast. Both models fit the empirical training data well ($R^2 \approx$ 0.92).

Table 6. Summary of Metrics

		Training Set (Fit)		Test Set (Forecast)	
Model Type	**Dataset**	**Total RMSE**	**Total R^2**	**Total RMSE**	**Total R^2**
Neural SDE	Synthetic	0.5054	0.9464	0.4496	0.9198
Neural UDE	Synthetic	0.1200	0.9940	0.1233	0.9928
Neural SDE	Empirical	1.4041	0.9206	1.5777	0.8753
Neural UDE	Empirical	1.4063	0.9250	2.5274	0.7690

However, on the empirical test set, the Neural SDE demonstrated better forecasting ability ($R^2 \approx 0.88$), whereas the Neural UDE's forecast performance declined substantially ($R^2 \approx 0.77$).

5. Conclusion

This study provided a comparative evaluation of Neural Universal Differential Equations (UDEs) and Neural Stochastic Differential Equations (Neural SDEs) for modeling predator-prey dynamics, using both synthetic Lotka-Volterra data and empirical algae-rotifer time series. While both frameworks demonstrated the capacity to capture complex dynamics, a significant finding was the challenge of parameter identifiability. Neither method reliably recovered the true Lotka-Volterra parameters from noisy synthetic data when using the flexible configurations tested. The learned parameters likely represent 'effective' coefficients compensated by the neural network components, which readily adapt to model discrepancies or noise. This highlights an inherent difficulty in disentangling mechanism and data-driven correction in complex hybrid models and warrants caution in the direct biological interpretation of the learned mechanistic parameters without further analysis or constraints. The findings address the core research questions:

Accuracy: Both Neural UDEs and SDEs successfully replicated the complex oscillatory dynamics present in the training portions of both datasets. The Neural UDE generally achieved slightly higher deterministic accuracy (lower MSE/RMSE, higher R^2) on the synthetic data, while the Neural SDE exhibited better forecasting performance on the empirical test data, suggesting its explicit noise modeling aided generalization.

Parameter Learning: Neither method reliably recovered the true Lotka-Volterra parameters from noisy synthetic data when run with the flexible configurations used here. The learned parameters in both frameworks likely represent 'effective' parameters compensated by the neural network components. This highlights the inherent challenge of parameter identifiability in complex hybrid models.

Role of NN Components: The learned neural network components played a substantial role in both synthetic and empirical cases. They learned complex, state-dependent corrections that significantly modified the dynamics produced by the base Lotka-Volterra structure (using its learned parameters).

Analysis of these components revealed where the simple prior model was deficient. The learned diffusion term in the SDEs provided an estimate of the system's stochastic nature.

Practical Aspects: UDEs were generally simpler to implement and train using standard ODE solvers. SDEs required specialized solvers and careful hyperparameter tuning (e.g., step size dt, gradient clipping) to ensure stability, but offered the ability to model and simulate stochastic paths. These SciML frameworks offer valuable pedagogical tools for mathematical and ecological education. Students can explore the interplay between mechanistic knowledge and data-driven learning, gaining hands-on experience with hybrid modeling approaches. The visual dynamics decomposition particularly aids in understanding how neural networks complement traditional ecological models, bridging theoretical concepts with computational implementation.

In summary, both SciML approaches offer powerful tools for ecological modeling, capable of integrating prior knowledge with data-driven learning. Neural UDEs excel at finding deterministic corrections to model structure, while Neural SDEs provide a framework for explicitly learning and simulating stochastic dynamics. The choice between them may depend on the specific goals: prioritizing deterministic accuracy versus characterizing system noise and uncertainty.

Data Availability Statement: The original data presented in the study are openly available in FigShare at https://doi.org/10.6084/m9.figshare.10045976.v1.

Conflicts of Interest: The authors declare no conflicts of interest.

References

[1] Pecl, G. T., Araujo, M. B., Bell, J. D., Blanchard, J., Bonebrake, T. C., Chen, I.-C., Clark, T. D., Colwell, R. K., Danielsen, F., Evengard, B. (2017). Biodiversity redistribution under climate change: Impacts on ecosystems and human well-being. Science, 355(6332), eaai9214. https://doi.org/10.1126/science.aai9214

[2] Berryman, A. A. (1992). The origins and evolution of predator-prey theory. Ecology, 73(5), 1530–1535. https://doi.org/10.2307/1940017

[3] Lotka, A. J. (1925). Elements of physical biology. Williams & Wilkins Company.

[4] Volterra, V. (1926). Variazioni e fluttuazioni del numero d'individui in specie animali conviventi [Variations and fluctuations of the number of individuals in coexisting animal species]. Memorie della Reale Accademia Nazionale dei Lincei, Classe di Scienze Fisiche, Matematiche e Naturali, Serie VI, 2, 31–113.

[5] Abrams, P. A. (2000). The complexity of simplicity: the population consequences of prey behavioural responses to predators. In A. P. Kinclle & D. Tilman (Eds.), *Functional consequences of biodiversity: experimental progress and theoretical extensions* (pp. 205–234). Princeton University Press