The Journal of Open Source Software

Interfere: Studying Intervention Response Prediction in Complex Dynamic Models

D. J. Passey [1], **Alice C. Schwarze** [2], **Zachary M. Boyd** [3], and **Peter J. Mucha** [2]

1 University of North Carolina at Chapel Hill, United States **2** Dartmouth College, United States **3** Brigham Young University, United States

DOI: 10.xxxxxx/draft

Software

- Review
- Repository
- Archive

Editor: Open Journals

Reviewers:

- @openjournals

Submitted: 01 January 1970
Published: unpublished

Summary

The vision of Interfere is simple: What if we used high-quality scientific models to create causal dynamic benchmark scenarios? Randomized experimental data and intervention response time series are essential for testing methods that attempt to infer dynamic relationships from data, but obtaining such datasets can be expensive and difficult. Mechanistic models are commonly developed to simulate scenarios and predict the response of systems to interventions across economics, neuroscience, ecology, systems biology and other areas (Baker et al., 2018; Banks et al., 2017; Brayton et al., 2014; Izhikevich & Edelman, 2008). Because these models are calibrated to the real world, they have the ability to generate diverse, complex, synthetic intervention responses that are characteristic of the real processes they emulate. Interfere offers the first steps towards this vision by combining (1) a general interface for simulating the effect of interventions on dynamic models, (2) a suite of predictive methods and cross validated hyper parameter optimization tools, and (3) the first known extensible benchmark data set of dynamic intervention response scenarios see Figure 1.

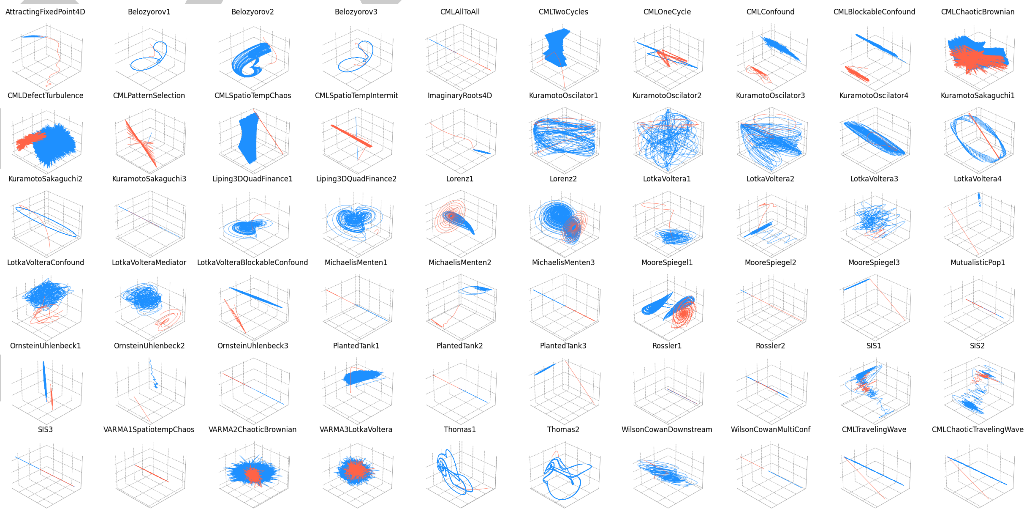

Figure 1: Three-dimensional trajectories of sixty scenarios simulated with the Interfere package. The models simulated here are either differential equations or discrete time difference equations. For each system, the trajectory in blue represents the natural behavior of the system and the red depicts how the system responds to a specified intervention. Many of the models pictured have more than three dimensions (in such cases, only the three dimensions of the trajectory with the highest variance are shown). These sixty scenarios make up the Interfere Benchmark 1.1.1 for intervention response prediction which is available online for download.

Statement of Need

Over the past twenty years, the scientific community has experienced the emergence of multiple frameworks for identifying causal relationships in observational data (Imbens & Rubin, 2015; Pearl, 2009; Wieczorek & Roth, 2019). The most influential frameworks are probabilistic and, while it is not a necessary condition for identifying causality, historically a static, linear relationship has often been assumed. However, when attempting to anticipate the response of complex dynamic systems in the medium and long term, a linear approximation of the dynamics can be insufficient. Therefore, researchers have increasingly begun to employ non-linear, dynamic techniques for causal discovery and forecasting (e.g. Runge, 2022). Still, there are relatively few techniques that are able to fit causal dynamic nonlinear models to data. Because of this, we see an opportunity to bring together the insights from recent advancements in causal inference with historical work in dynamic modeling and simulation.

In order to facilitate this cross pollination, we focus on a key problem — predicting how a complex system responds to a previously unobserved intervention — and designed the Interfere package for benchmarking tools aimed at intervention response prediction. The dynamic models contained in Interfere present challenges for computational methods that can likely only be addressed with the incorporation of mechanistic assumptions alongside probabilistic frameworks for causality. The Interfere package is a toolbox that allows researcher to validate predictive dynamic methods against simulated intervention scenarios. As such, the Interfere package encourages an opportunity for cross pollination between the probabilistic causal inference community and the modeling and simulation community.

Usage

The Interfere package is designed around four tasks: (1) simulation, (2) intervention, (3) forecasting method optimization and (4) intervention response prediction. The following section will describe each task in detail alongside example code.

1. Simulation.

The models implemented in the Interfere package are mainly stochastic differential equations (e.g. $d\mathbf{X} = A\mathbf{X} + d\mathbf{W}$) simulated with the users choice of Ito's method or SR1, (a strong Runge-Kutta method, see (Rößler, 2010)) or stochastic difference equations (e.g. $x[n+1] = 0.25x[n] - 0.5x[n-1]$), simulated via initial conditions and stepping forward in time. Each dynamic model class included in the Interfere package has a simulate method. To run a simulation, the package requires an array of equally spaced time values and and array containing initial conditions or historic observations. For example, the following code block produces the trajectories similar to those plotted in Figure 2:

```python
import numpy as np
import interfere
import optuna

# Set up simulation parameters
initial_cond = np.random.rand(3)
t_train = np.arange(0, 10, 0.05)
dynamics = interfere.dynamics.Belozyorov3DQuad(sigma=0.5)

# Generate trajectory
sim_states = dynamics.simulate(t_train, initial_cond)
```

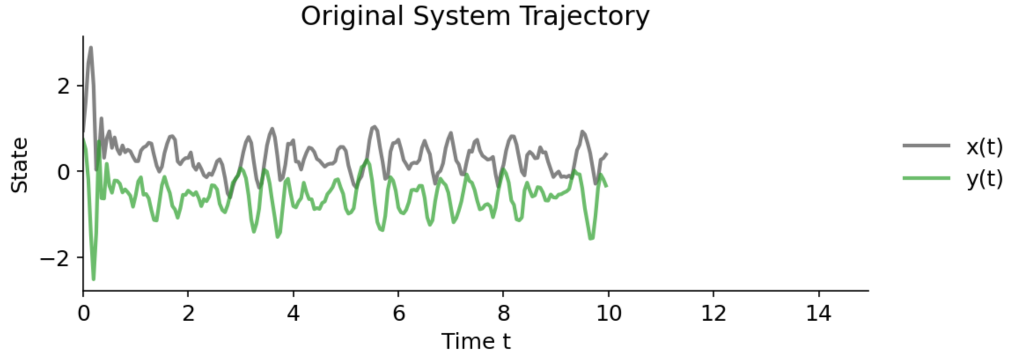

Figure 2: Simulation of System: The natural, uninterupted evolution of the quaratic Belozyorov system (Belozyorov, 2015) with the addition of a small amount of stochastic noise simulated using the Interfere package.

2. Intervention

Next, we can take exogenous control of x by pinning it to $\sin(t)$ and simulate the response of y. The resulting simulation reveals how the behavior of the system is altered by this particular intervention. See Figure 3 to see how the quadratic Belozyorov system (Belozyorov, 2015) responds to this intervention.

```
# Time points for the intervention simulation
test_t = np.arange(t_train[-1], 15, 0.05)

# Intervention initialization
intervention = interfere.SignalIntervention(iv_idxs=1, signals=np.sin)

# Simulate intervention
interv_states = dynamics.simulate(
    test_t,
    prior_states=sim_states,
    intervention=intervention,
)
```

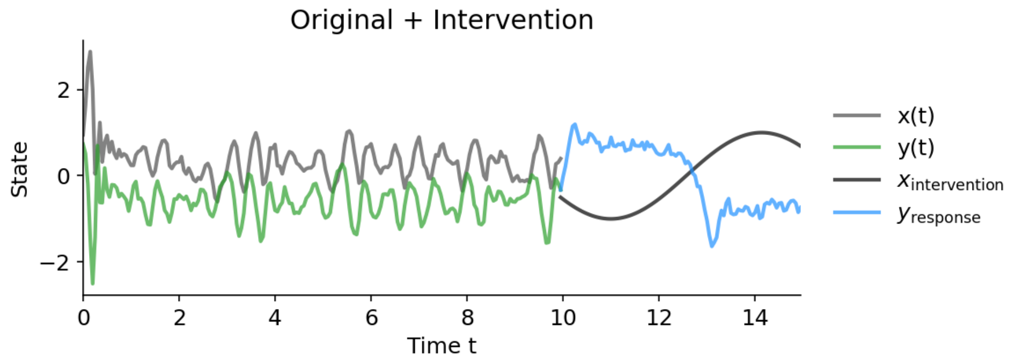

Figure 3: System trajectory with intervention: The figure above demonstrates the effect that taking exogenous control of $x(t)$ by via $\text{do}(x(t) = \sin(t))$ has on y. The intervention (black) and response (blue), depict a clear departure from the natural evolution behavior of the system.

3. Optimization

Interfere offers tools to optimize forecasting methods for time series prediction. By using Interfere's cross validation objective function along with a hyperparameter optimizer such as Optuna (Akiba et al., 2019), it is possible to compare hyperparameter settings on multiple folds of time series data. To simplify this process, every Interfere forecasting method comes with sensible preset hyperparameter ranges for the optimizer to explore. Using the cross validation objective function for hyperparameter optimization is demonstrated in the following code block.

```
# Select the SINDy method for hyperparameter optimization.
method_type = interfere.SINDy

# Create an objective function that aims to minimize cross validation error
# over different hyper parameter configurations for SINDy
cv_obj = interfere.CrossValObjective(
    method_type=method_type,
    data=sim_states,
    times=t_train,
    train_window_percent=0.3,
    num_folds=5,
    exog_idxs=intervention.iv_idxs,
)

# Run the study using optuna.
study = optuna.create_study()
study.optimize(cv_obj, n_trials=25)

# Collect the best hyperparameters into a dictionary.
best_param_dict = study.best_params
```

4. Intervention Response Prediction

Using the best parameters from the hyperparameter optimization run, we can fit the forecasting method to all the data that occurred prior to the intervention, treating the states we plan to manipulate as exogenous. This way, the method expects to be given exogenous data about the intervention variable(s). After fitting to the unperturbed system, we forecast the intervention response by treating the desired intervention as an exogenous input signal applied during forecasting. Figure 4 demonstrates the optimized model's ability to forecast the intervention response.

```
# Initialize SINDy with the best perfoming parameters.
method = interfere.SINDy(**study.best_params)

# Use an intervention helper function to split the pre-intervention data
# into endogenous and exogenous columns.
Y_endog, Y_exog = intervention.split_exog(sim_states)

# Fit SINDy to the pre-intervention data.
method.fit(t_train, Y_endog, Y_exog)

# Use the inherited interfere.ForecastingMethod.simulate() method
# To simulate intervention response using SINDy
pred_traj = method.simulate(
    test_t, prior_states=sim_states, intervention=intervention
)
```

The Journal of Open Source Software

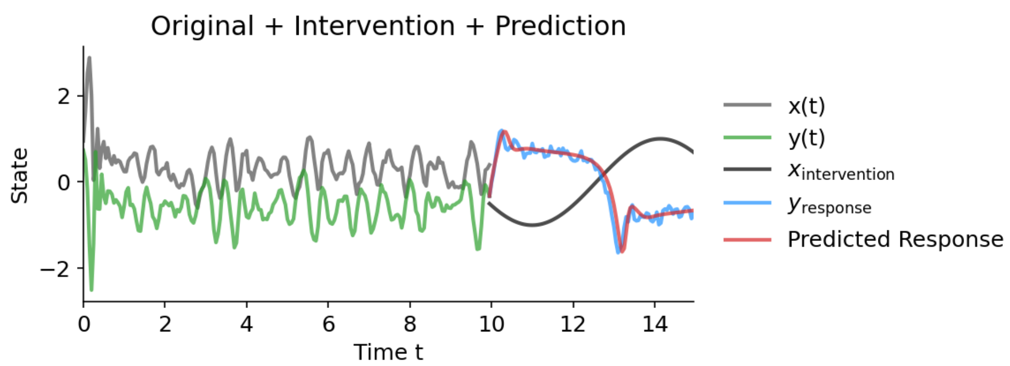

Figure 4: Forecasting Intervention Response: Example of forecasting the response of the Belozyorov system to a sinusoidal intervention. Here, the intervention consists of taking exogenous control of $x(t)$ (black). The ground truth response, $y(t)$ for $t > 10$ is plotted in blue. Here, an equation discovery algorithm, SINDy, (Brunton et al., 2016) is fit to the data that occurs prior to the intervention and makes an attempt to predict the intervention response (red curve). However, not all intervention responses can be so accurately predicted. See the bottom left of Figure 5 for a failed forecast scenario.

Primary Contributions

The Interfere package provides three primary contributions. (1) Dynamically diverse counterfactuals at scale, (2) cross diciplinary forecast methods, and (3) comprehensive and extensible benchmarking.

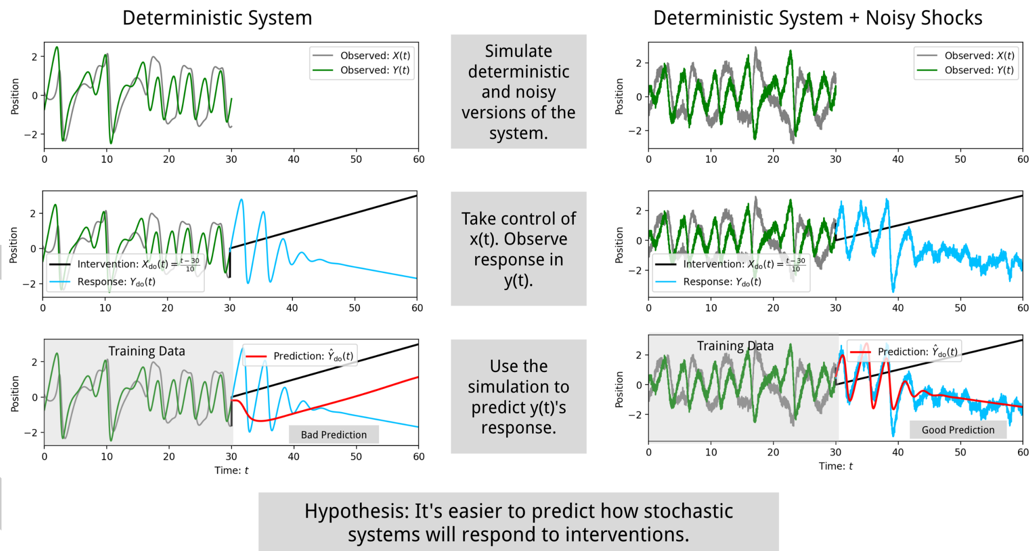

Figure 5: Example experimental setup possible with Interfere: Can stochasticity help reveal associations between variables? Interfere can be used to comparing intervention response prediction for deterministic and stochastic versions of the same system.

1. Dynamically Diverse Counterfactuals at Scale

The "dynamics" submodule in the Interfere package contains over fifty dynamic models. It contains a mix of linear, nonlinear, chaotic, continuous time, discrete time, stochastic, and deterministic models. The models come from a variety of disciplines including finance, ecology, biology, neuroscience and public health. Each model inherits the from the Interfere

BaseDynamics type and gains the ability to take exogenous control of any observed state and to add measurement noise. Most models also gain the ability to make any observed state stochastic where magnitude of stochasticity can be controlled by a simple scalar parameter or fine tuned with a covariance matrix.

Because of the difficulty of building models of complex systems, predictive methods for complex dynamics are typically benchmarked on less than ten dynamical systems (Brunton et al., 2016; Challu et al., 2023; Pathak et al., 2018; Prasse & Van Mieghem, 2022; Vlachas et al., 2020). As such, Interfere offers a clear improvement over current benchmarking methods for prediction in complex dynamics.

Most importantly, Interfere is built around interventions: the ability to take exogenous control of one or several state variables in a complex system and observe the response. Imbuing a suite of scientific models with general exogenous control is no small feat because models can be complex and are implemented in a variety of ways. Interfere offers the ability to produce complex dynamic intervention response and standard forecasting scenarios at scale. This unique feature enables large scale evaluation of dynamic causal prediction methods—tested against systems with properties of interest to scientists. For example, we can simulate the change in concentration of ammonia based on the nitrogen cycle and an exogenous fertilizing schedule.

2. Cross Disciplinary Forecast Methods

A second contribution of Interfere is the integration of dynamic *forecasting* methodologies from deep learning (LSTM, NHITS), applied mathematics (SINDy, Reservoir Computers) and social science (VAR). The Interfere "ForecastingMethod" class is expressive enough to describe, fit and predict with multivariate dynamic models and apply interventions to the states of the models during prediction. This cross disciplinary mix of techniques has the potential to produce new insights into the problem of intervention response prediction among others. For example, experiments using this package have revealed that cross validation error does not correlate with well with prediction error when LSTM and NHITS attempt to predict intervention response.

3. Comprehensive and Extensible Benchmarking

The third major contribution of Interfere is the collection of dynamic scenarios organized into the Interfere Benchmark. The Interfere Benchmark is a comprehensive and extensible set of dynamic scenarios that are conveniently available for testing methods that predict the effects of interventions. The benchmark set contains 60 intervention response scenarios for testing, each simulated with different levels of stochastic noise. Each scenario is housed in a JSON file, complete with full metadata annotation, documentation, versioning and commit hashes marking the commit of Interfere that was used to generate the data. The scenarios was revied by hand with some systems exposed to exogenous input to ensure that none of the key variables settle into a steady state. Additionally, all interventions were chosen in a manner such that the response of the target variable is a significant departure from its previous behavior.

The Interfere package enables researchers from various backgrounds to systematically study the problem of predicting intervention response on simulated data from a wide range of disciplines. It thereby facilitates future progress towards correctly anticipating how complex systems will respond in new, never before seen scenarios.

Related Software and Mathematical Foundations

Predictive Methods

The Interfere package draws from the Nixtla open source ecosystem for time series forecasting. We implemented intervention support for LSTM and NHITS from the NeuralForecast package, and for ARIMA from the StatsForecast package (Azul Garza, 2022; Olivares et al., 2022). We followed Nixtla's example for cross validation and hyperparameter optimization approaches. We integrated predictive methods from the PySINDy (Kaptanoglu et al., 2022) and StatsModels (Seabold & Perktold, 2010) packages. We also include ResComp, a reservoir computing method for global forecasts from (Harding et al., 2024). Hyperparameter optimization is designed around the Optuna framework (Akiba et al., 2019).

While other forecasting methods exist, integrating a method with Interfere requires that the method is capable of (1) multivariate endogenous dynamic forecasting, (2) support for exogenous variables, and (3) support for flexible length forecast windows or recursive predictions. Few forecasting methods meet these criteria, and it is our hope that this package can encourage the development of additional methods.

Dynamic Models

The table below list the dynamic models that are currently implemented in the Interfere package, plus attributions. These dynamic models in were implemented directly from mathematical descriptions except for two, "Hodgkin Huxley Pyclustering" and "Stuart Landau Kuramoto" which adapt existing simulations from the PyClustering package (Novikov, 2019).

Dynamic Model Class	Description and Source	Properties
Arithmetic Brownian Motion	Brownian motion with linear drift and constant diffusion (Øksendal, 2005)	Stochastic, Linear
Coupled Logistic Map	Discrete-time logistic map with spatial coupling (Lloyd, 1995)	Nonlinear, Chaotic
Stochastic Coupled Map Lattice	Coupled map lattice with stochastic noise (Kaneko, 1991)	Nonlinear, Stochastic, Chaotic
Michaelis Menten	Model for enzyme kinetics and biochemical reaction networks (Srinivasan, 2022)	Nonlinear, Stochastic
Lotka Voltera SDE	Stochastic Lotka-Volterra predator-prey model (Hening & Nguyen, 2018)	Nonlinear, Stochastic
Kuramoto	Coupled oscillator model to study synchronization (Rodrigues et al., 2016)	Nonlinear, Stochastic
Kuramoto Sakaguchi	Kuramoto model variant with phase frustration (Sakaguchi & Kuramoto, 1986)	Nonlinear, Stochastic
Hodgkin Huxley Pyclustering	Neuron action-potential dynamics based on Hodgkin-Huxley equations (Hodgkin & Huxley, 1952)	Nonlinear
Stuart Landau Kuramoto	Coupled oscillators with amplitude-phase dynamics (Cliff et al., 2023)	Nonlinear, Stochastic
Mutualistic Population	Dynamics of interacting mutualistic species (Prasse & Van Mieghem, 2022)	Nonlinear
Ornstein Uhlenbeck	Mean-reverting stochastic differential equation (Gardiner, 2009)	Stochastic, Linear
Belozyorov 3D Quad	3-dimensional quadratic chaotic system (Belozyorov, 2015)	Nonlinear, Chaotic
Liping 3D Quad Finance	Chaotic dynamics applied in financial modeling (Liping et al., 2021)	Nonlinear, Chaotic
Lorenz	Classic chaotic system describing atmospheric convection (Lorenz, 2017)	Nonlinear, Chaotic
Rossler	Simplified 3D chaotic attractor system (Rössler, 1976)	Nonlinear, Chaotic
Thomas	Chaotic attractor with simple structure and rich dynamics (Thomas, 1999)	Nonlinear, Chaotic

Dynamic Model Class	Description and Source	Properties
Damped Oscillator	Harmonic oscillator with damping and noise (Classical linear model)	Linear, Stochastic
SIS	Epidemiological model (Susceptible-Infected-Susceptible) (Prasse & Van Mieghem, 2022)	Nonlinear, Stochastic
VARMA Dynamics	Vector AutoRegressive Moving Average for time series modeling (Hamilton, 2020)	Linear, Stochastic
Wilson Cowan	Neural mass model for neuronal population dynamics (Wilson & Cowan, 1972)	Nonlinear
Geometric Brownian Motion	Stochastic model widely used in financial mathematics (Black & Scholes, 1973)	Nonlinear, Stochastic
Planted Tank Nitrogen Cycle	Biochemical cycle modeling nitrogen transformation in aquatic systems (Fazio & Jannelli, 2006)	Nonlinear
Generative Forecaster	Predictive forecasting models trained on simulation, then used to generate data (Written for Interfere)	Stochastic
Standard Normal Noise	IID noise from standard normal distribution (Cliff et al., 2023)	Stochastic
Standard Cauchy Noise	IID noise from standard Cauchy distribution (Cliff et al., 2023)	Stochastic
Standard Exponential Noise	IID noise from standard exponential distribution (Cliff et al., 2023)	Stochastic
Standard Gamma Noise	IID noise from standard gamma distribution (Cliff et al., 2023)	Stochastic
Standard T Noise	IID noise from Student's t-distribution (Cliff et al., 2023)	Stochastic

Acknowledgements

The work described here was supported by an NSF Graduate Research Fellowship (DJP) and by award W911NF2510049 from the Army Research Office. The content is solely the responsibility of the authors and does not necessarily represent the official views of any agency supporting this research.

References

Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. (2019). Optuna: A next-generation hyperparameter optimization framework. *The 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2623–2631.

Azul Garza, C. C., Max Mergenthaler Canseco. (2022). *StatsForecast: Lightning fast forecasting with statistical and econometric models*. PyCon Salt Lake City, Utah, US 2022. https://github.com/Nixtla/statsforecast

Baker, R. E., Peña, J.-M., Jayamohan, J., & Jérusalem, A. (2018). Mechanistic models versus machine learning, a fight worth fighting for the biological community? *Biology Letters*, *14*(5), 20170660. https://doi.org/10.1098/rsbl.2017.0660

Banks, H. T., Banks, J. E., Bommarco, R., Curtsdotter, A., Jonsson, T., & Laubmeier, A. N. (2017). Parameter estimation for an allometric food web model. *International Journal of Pure and Apllied Mathematics*, *114*(1). https://doi.org/10.12732/ijpam.v114i1.12

Belozyorov, V. Y. (2015). Exponential-algebraic maps and chaos in 3D autonomous quadratic systems. *International Journal of Bifurcation and Chaos*, *25*(04), 1550048.

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, *81*(3), 637–654.

Brayton, F., Laubach, T., & Reifschneider, D. (2014). The FRB/US model: A tool for macroeconomic policy analysis. *FEDS Notes*, *2014-04*, 03.

Brunton, S. L., Proctor, J. L., & Kutz, J. N. (2016). Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, *113*(15), 3932–3937. https://doi.org/10.1073/pnas.1517384113

Challu, C., Olivares, K. G., Oreshkin, B. N., Ramirez, F. G., Canseco, M. M., & Dubrawski, A. (2023). NHITS: Neural Hierarchical Interpolation for Time Series Forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, *37*(6), 6989–6997. https://doi.org/10.1609/aaai.v37i6.25854

Cliff, O. M., Bryant, A. G., Lizier, J. T., Tsuchiya, N., & Fulcher, B. D. (2023). Unifying pairwise interactions in complex dynamics. *Nature Computational Science*, *3*(10), 883–893. https://doi.org/10.1038/s43588-023-00519-x

Fazio, R., & Jannelli, A. (2006). Mathematical and numerical modeling for a bio-chemical aquarium. *Applied Mathematics and Computation*, *174*(2), 1370–1383.

Gardiner, C. (2009). *Stochastic methods* (Vol. 4). Springer Berlin Heidelberg.

Hamilton, J. D. (2020). *Time series analysis*. Princeton university press.

Harding, S., Leishman, Q., Lunceford, W., Passey, D. J., Pool, T., & Webb, B. (2024). Global forecasts in reservoir computers. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, *34*(2), 023136. https://doi.org/10.1063/5.0181694

Hening, A., & Nguyen, D. H. (2018). Stochastic lotka–volterra food chains. *Journal of Mathematical Biology*, *77*, 135–163.

Hodgkin, A. L., & Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, *117*(4), 500.

Imbens, G. W., & Rubin, D. B. (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction* (1st ed.). Cambridge University Press. https://doi.org/10.1017/CBO9781139025751

Izhikevich, E. M., & Edelman, G. M. (2008). Large-scale model of mammalian thalamocortical systems. *Proceedings of the National Academy of Sciences*, *105*(9), 3593–3598. https://doi.org/10.1073/pnas.0712231105

Kaneko, K. (1991). Coupled map lattice. In *Chaos, Order, and Patterns* (pp. 237–247). Springer.

Kaptanoglu, A. A., Silva, B. M. de, Fasel, U., Kaheman, K., Goldschmidt, A. J., Callaham, J., Delahunt, C. B., Nicolaou, Z. G., Champion, K., Loiseau, J.-C., Kutz, J. N., & Brunton, S. L. (2022). PySINDy: A comprehensive python package for robust sparse system identification. *Journal of Open Source Software*, *7*(69), 3994. https://doi.org/10.21105/joss.03994

Liping, C., Khan, M. A., Atangana, A., & Kumar, S. (2021). A new financial chaotic model in Atangana-Baleanu stochastic fractional differential equations. *Alexandria Engineering Journal*, *60*(6), 5193–5204.

Lloyd, A. L. (1995). The coupled logistic map: A simple model for the effects of spatial heterogeneity on population dynamics. *Journal of Theoretical Biology*, *173*(3), 217–230.

Lorenz, E. N. (2017). Deterministic Nonperiodic Flow 1. In *Universality in Chaos, 2nd edition* (pp. 367–378). Routledge.

Novikov, A. V. (2019). PyClustering: Data mining library. *Journal of Open Source Software*, *4*(36), 1230. https://doi.org/10.21105/joss.01230

Øksendal, B. K. (2005). *Stochastic differential equations: An introduction with applications* (6th ed., corrected third printing 2005). Springer. ISBN: 978-3-540-04758-2

Olivares, K. G., Challú, C., Garza, A., Canseco, M. M., & Dubrawski, A. (2022). *NeuralForecast: User friendly state-of-the-art neural forecasting models.* PyCon Salt Lake City, Utah, US 2022. https://github.com/Nixtla/neuralforecast

Pathak, J., Hunt, B., Girvan, M., Lu, Z., & Ott, E. (2018). Model-Free Prediction of Large Spatiotemporally Chaotic Systems from Data: A Reservoir Computing Approach. *Physical Review Letters*, *120*(2), 024102. https://doi.org/10.1103/PhysRevLett.120.024102

Pearl, J. (2009). *Causality* (2nd ed.). Cambridge University Press. https://doi.org/10.1017/CBO9780511803161

Prasse, B., & Van Mieghem, P. (2022). Predicting network dynamics without requiring the knowledge of the interaction graph. *Proceedings of the National Academy of Sciences*, *119*(44), e2205517119. https://doi.org/10.1073/pnas.2205517119

Rodrigues, F. A., Peron, T. K. D., Ji, P., & Kurths, J. (2016). The Kuramoto model in complex networks. *Physics Reports*, *610*, 1–98.

Rößler, A. (2010). Runge–Kutta Methods for the Strong Approximation of Solutions of Stochastic Differential Equations. *SIAM Journal on Numerical Analysis*, *48*(3), 922–952. https://doi.org/10.1137/09076636X

Rössler, O. E. (1976). An equation for continuous chaos. *Physics Letters A*, *57*(5), 397–398.

Runge, J. (2022). *Discovering contemporaneous and lagged causal relations in autocorrelated nonlinear time series datasets.* arXiv. https://doi.org/10.48550/arXiv.2003.03685

Sakaguchi, H., & Kuramoto, Y. (1986). A soluble active rotater model showing phase transitions via mutual entertainment. *Progress of Theoretical Physics*, *76*(3), 576–581.

Seabold, S., & Perktold, J. (2010). Statsmodels: Econometric and statistical modeling with python. *9th Python in Science Conference*.

Srinivasan, B. (2022). A guide to the Michaelis–Menten equation: Steady state and beyond. *The FEBS Journal*, *289*(20), 6086–6098.

Thomas, R. (1999). Deterministic chaos seen in terms of feedback circuits: Analysis, synthesis," labyrinth chaos". *International Journal of Bifurcation and Chaos*, *9*(10), 1889–1905.

Vlachas, P. R., Pathak, J., Hunt, B. R., Sapsis, T. P., Girvan, M., Ott, E., & Koumoutsakos, P. (2020). Backpropagation algorithms and Reservoir Computing in Recurrent Neural Networks for the forecasting of complex spatiotemporal dynamics. *Neural Networks*, *126*, 191–217. https://doi.org/10.1016/j.neunet.2020.02.016

Wieczorek, A., & Roth, V. (2019). Information Theoretic Causal Effect Quantification. *Entropy*, *21*(10), 975. https://doi.org/10.3390/e21100975

Wilson, H. R., & Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, *12*(1), 1–24.