

# UC Santa Barbara

## UC Santa Barbara Previously Published Works

### Title

Interpretable polynomial neural ordinary differential equations

### Permalink

<https://escholarship.org/uc/item/063153hd>

### Journal

Chaos An Interdisciplinary Journal of Nonlinear Science, 33(4)

### ISSN

1054-1500

### Authors

Fronk, Colby  
Petzold, Linda

### Publication Date

2023-04-01

### DOI

10.1063/5.0130803

Peer reviewed

---

# Interpretable Polynomial Neural Ordinary Differential Equations

---

**Colby Fronk**

University of California, Santa Barbara  
Santa Barbara, CA 93106  
colbyfronk@ucsb.edu

**Linda Petzold**

University of California, Santa Barbara  
Santa Barbara, CA 93106  
petzold@ucsb.edu

## Abstract

Neural networks have the ability to serve as universal function approximators, but they are not interpretable and don't generalize well outside of their training region. Both of these issues are problematic when trying to apply standard neural ordinary differential equations (ODEs) to dynamical systems. We introduce the polynomial neural ODE, which is a deep polynomial neural network inside of the neural ODE framework. We demonstrate the capability of polynomial neural ODEs to predict outside of the training region, as well as to perform direct symbolic regression without using additional tools such as SINDy.

## 1 Introduction

Dynamical systems are mathematical equations written to describe the interactions of quantities that change in time in many science and engineering applications. Traditionally, mechanistic models describing these systems were obtained from the iterative process of deriving equations from first principles and testing the models with physical experimentation. The power of mechanistic models comes from the ability to directly explain the system at hand with known physical principles such as the thermodynamics of the system, heat and mass transfer processes, chemical kinetics, and the system's forces. Reliable predictions can be made about the system behavior in different regimes from mechanistic models due to the scientist's knowledge of the assumptions governing the various first principles and physical laws. Mechanistic models are preferred by scientists and engineers; however, it can take years to develop a complete and accurate description of a system and we do not know the underlying first principles describing all systems. Brunton et al. [2016] sped up the identification of mechanistic models for nonlinear dynamical systems with the method Sparse Identification of Nonlinear Dynamics (SINDy), which is a sparse regression problem of time derivatives obtained by finite difference methods with a library of candidate terms appearing in the dynamical system. SINDy has had great success with system identification for various applications such as plasma physics [Alves and Fiuza, 2020], nonlinear optical communication [Sorokina et al., 2016], biological chemical reaction networks [Mangan et al., 2016, Hoffmann et al., 2019], and fluid dynamics [Rudy et al., 2016].

With the emergence of increased computational power from GPUs and CPUs, along with the exponential growth of the amount of data gathered such as via automated experimentation in chemistry and biology [Mayr and Bojanic, 2009, Szymański et al., 2011], sensor data from factories [Kalsoom et al., 2020], satellite and in situ earth observations [Balsamo et al., 2018], and large-scale simulations such as computational fluid dynamics (CFD) simulations [Calzolari and Liu, 2021] and climate models [Randall et al., 2007, Maher et al., 2021], data-driven models such as deep learning have emerged as a way to process and understand this large amount of data quickly.

Neural ordinary differential equations (ODEs) are a recent approach to data-driven modeling of time-series data and dynamical systems in which a neural network is used to learn an approximation

to an equation governing the dynamics of the system. Neural ODEs were first introduced in Chen et al. [2018]’s seminal NeurIPS best paper. In the last few years, several other types of neural ODEs have emerged such as latent ODEs [Rubanova et al., 2019], Bayesian neural ODEs [Dandekar et al., 2020], and neural stochastic ODEs [Li et al., 2020]. Models describing the system can be created solely from observed data without the need for expert domain knowledge, which Rackauckas et al. [2020] term as the universal differential equation. These models can then be used to make predictions of what the system will do for unobserved conditions. Traditional data-driven neural ODE models can be obtained in a few hours as opposed to the many years required to develop complete and accurate mechanistic models for dynamical systems. Neural ODEs integrate the neural ODE in time to obtain predictions for the observed data, whereas SINDy uses finite difference methods with the data to obtain numerical approximations for the time derivatives, which gives neural ODEs the advantage of having less stringent requirements on the frequency of the observations [Champion et al., 2019]. However, neural ODEs have the well-known major problems that they are not directly interpretable, and they do not make reliable predictions outside of the domain of their training region. [Lee et al., 2021] were able to make the neural ODE interpretable by successfully recovering symbolic equations from conventional neural ODEs by using SINDy with predictions for the time derivatives obtained from the trained neural ODE.

Our work is motivated by the need to make neural ODEs directly interpretable without using additional methods like SINDy after training. We address this problem by making the case to use directly interpretable neural network architectures inside of the neural ODE framework rather than the conventional neural network based on the standard multilayer perceptron (MLP) [Goodfellow et al., 2016]. Nonlinear dynamical systems can be complex symbolic expressions. We chose to address the class of dynamical systems described by polynomials first before tackling more advanced expressions. Dynamical systems involving polynomials arise in a number of physical systems such as gene regulatory networks [Sanguinetti and Huynh-Thu, 2018] and cell signaling networks [Gutkind, 2000] in systems biology, population models in ecology [Royle and Dorazio, 2008] and epidemiology [Singh and Dhar, 2018], and atmospheric chemical kinetics [Brasseur and Jacob, 2017]. In the process of this work, we developed a few deep polynomial neural network architectures, but had the most success with [Chrysos et al., 2022]’s  $\pi$ -net. We are the first to use any deep polynomial neural network for the purpose of direct symbolic regression, which we demonstrate on a fourth order univariate polynomial. We are also the first to put a deep polynomial neural network into the neural ODE framework, which we term the polynomial neural ODE, and use the polynomial neural ODE to perform direct symbolic regression on nonstiff dynamical systems described by a polynomials such as the Lotka-Volterra model, Damped Oscillatory System, and Van der Pol model. In addition, we test the polynomial neural ODE on a model that does not involve polynomials, and demonstrate its effectiveness as another form of a deep polynomial function approximator. This work, along with future work, will allow the neural ODE to serve as a tool complementary to SINDy for system identification.

## 2 Methods

### 2.1 Neural ODEs

Neural ordinary differential equations (ODEs) learn an approximate ODE, given data for the solution,  $y(t)$  [Chen et al., 2018]. The ODE that we seek to approximate is given by

$$\frac{dy(t)}{dt} = f(t, y(t), \theta), \tag{1}$$

where  $t$  is the time,  $y(t)$  is the vector of state variables,  $\theta$  is the vector of parameters, and  $f$  is the ODE model. For many scientific problems, it can take years to discover the functional form of the ODE described by  $f$ . Neural ODEs solve this problem by learning an approximation to the dynamics described by  $f$  without learning the exact functional equation. The neural ode, which we denote by  $NN$ , is a neural network that approximates the model  $f$ :

$$\frac{dy(t)}{dt} \approx NN(t, y(t), \theta). \tag{2}$$

Once the neural ODE has been trained, it is treated exactly the same way as an ODE. To obtain predictions for  $y(t)$ , the neural ODE is integrated as an initial value problem (IVP) with an ODE solver. Traditionally, neural ODEs have used the same architecture as a standard multilayer perceptron (MLP)

$$NN(x) = (L_1 \circ L_2 \cdots \circ L_{l-1} \circ L_l)(x), \tag{3}$$

that is the composition of several neural network layers,  $L_i$

$$L_i(x) = \sigma(x * w_i + b_i), \tag{4}$$

with nonlinear activation function  $\sigma$ , weights  $w_i$ , and bias  $b_i$  [Goodfellow et al., 2016].

## 2.2 Polynomial Neural ODEs

Mathematical models in numerous application areas including gene regulatory networks [Sanguinetti and Huynh-Thu, 2018] and cell signaling networks [Gutkind, 2000] in systems biology, population models in ecology [Royle and Dorazio, 2008] and epidemiology [Singh and Dhar, 2018], and atmospheric chemical kinetics [Brasseur and Jacob, 2017] are expressed as differential equations where the right hand side functions  $f$  are polynomials. For this class of problems we present the polynomial neural ODE. While Andoni et al. [2014] have theoretically and experimentally shown that conventional feedforward MLPs work as universal approximators for polynomials, we will make the case for using polynomial neural ODEs instead of conventional neural ODEs for this application space.

Polynomial neural networks are function approximators in which the output layer is expressed as a polynomial of the input layer. There are several types of polynomial neural networks. Designing polynomial neural networks that can be trained easily without an explosion of parameters is still an active area of research [Chrysos et al., 2022, Fan et al., 2020, Du and Lee, 2018, Liang and Srikant, 2016]. In this paper, we use the  $\pi$ -net architecture from Chrysos et al. [2022].  $\pi$ -nets were specifically designed to use skip connections to avoid the combinatorial explosion in the number of parameters of polynomial activation functions, which make the network harder to train. Essentially, these networks learn tensor decompositions of the polynomials. For our work we use  $\pi$ -net V1, for which the architecture is shown in Fig. 2. The architecture is centered around Hadamard Products [Horn et al., 1994] of linear layers without activation functions to form higher order polynomials.

Conventional neural networks are not interpretable, due to the complex arrangement of nonlinear activation functions. Lee et al. [2021] have had success using Brunton et al. [2016]’s Sparse Identification of Nonlinear Dynamics (SINDy) to recover symbolic forms of neural ODEs following training. However, the polynomial neural ODE architecture has a unique advantage. Since the output layer is a direct mapping of the input in terms of tensor and Hadamard products, symbolic tensor math can be used to obtain a direct polynomial form of the polynomial neural network without using additional tools such as SINDy. We use SymPy, the Python library for symbolic mathematics, to obtain a symbolic form of the polynomial neural ODE following training [Meurer et al., 2017]. Since we expect to see a plethora of new neural network architectures for symbolic regression beyond polynomials, we anticipate the need for compatibility of symbolic computation with software such as PyTorch, TensorFlow, and JAX, which would avoid the need to extract the network architecture, weights, and biases into SymPy following training to obtain a symbolic equation.

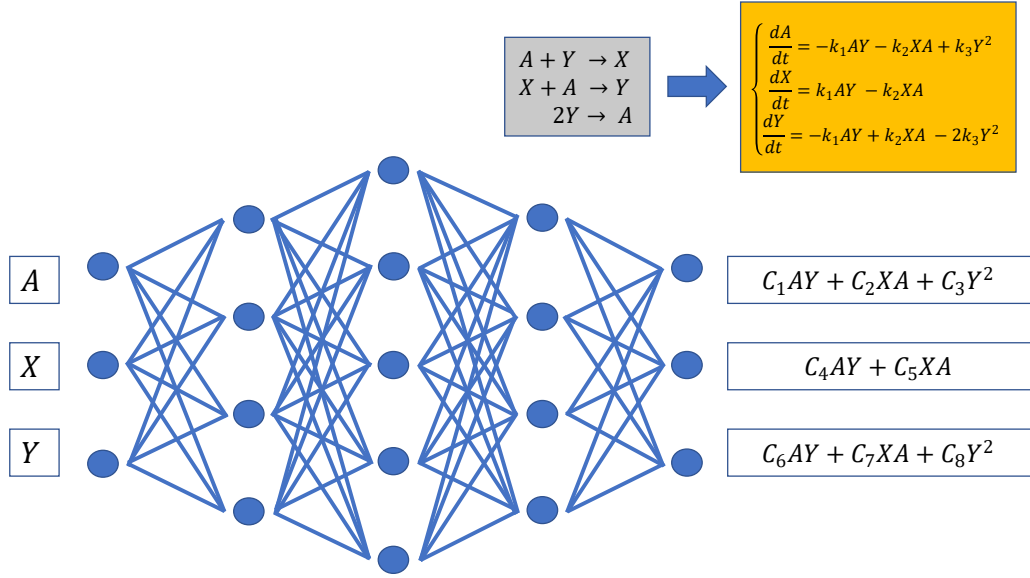


Figure 1: Example polynomial neural ODE for a chemical reaction system with molecules A, X, and Y. The neural network outputs a polynomial transformation of the input, which are the concentrations of the chemical species.

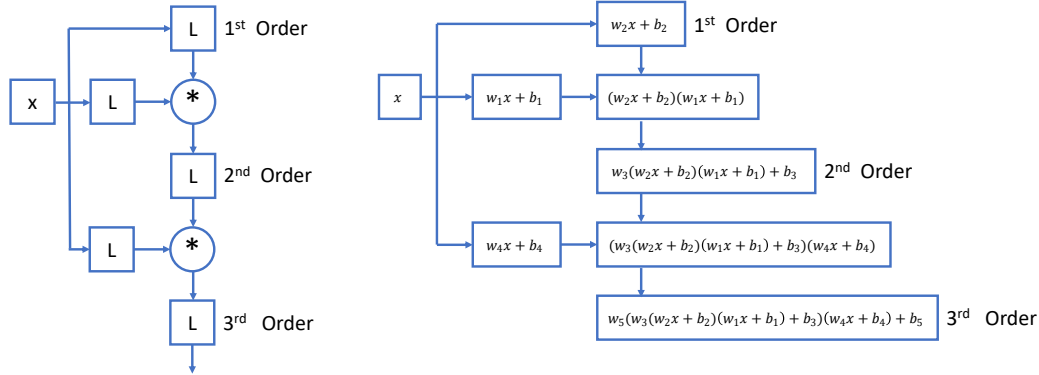


Figure 2: The network architecture of Chrysos et al. [2022]’s  $\pi$ -net V1 is shown (left) alongside a worked example of a one layer input with variable name  $x$  (right). The boxes labeled L represent standard linear layers, while the circles with the \* symbol represent a layer that is the Hadamard product of the layer’s inputs. There are no standard activation functions such as tanh or ReLU in this network architecture, which makes the architecture directly interpretable.

### 2.3 Training Neural ODEs

Prior to training the neural ODE, the architecture is defined and the parameters in the network are initialized. For all of our neural ODEs, we initialized our weights and biases with a normal distribution with mean of 0. We used a standard deviation of 0.00005 and 0.01 for the conventional and  $\pi$ -net V1 neural ODEs respectively.

When training a neural ODE, the goal is to fit the neural ODE to observed data for the state variables,  $y_{obs}$ , as a function of time. The neural ODE is integrated with an ODE solver to obtain predictions

for  $y_{obs}$ , which we will call  $y_{pred}$ . We used gradient descent to minimize the normalized MSE loss  $L(y_{obs}, y_{pred})$  between  $y_{obs}$  and  $y_{pred}$ :

$$L(y_{obs}, y_{pred}) = \frac{1}{N} \sum_i^N \left( \frac{y_{pred,i} - y_{obs,i}}{y_{scale}} \right)^2, \quad (5)$$

where  $y_{scale} = |y_{max} - y_{min}|$  is the scale factor used to normalize the values in the loss function. We batched our data into  $N - 1$  samples consisting of IVPs between 2 adjacent known data points. For each iteration of gradient descent, we simultaneously solve the  $N - 1$  initial value problems forward in time to the next observed data point explicitly using the fourth order explicit Runge–Kutta–Fehlberg method [Fehlberg, 1968]. Since we are only working with nonstiff ODEs, we are able to use an explicit discretization method to solve the neural ODEs. The advantage of using this method to solve the neural ODE is efficient direct backpropagation through the explicit ODE scheme, which the popular continuous time sensitivity adjoint method from Chen et al. [2018] lacks.

### 3 Results

#### 3.1 Univariate 4th Order Polynomial

Prior to looking at any dynamical systems with neural ODEs, the ability of  $\pi$ -net V1 to learn basis polynomials was first tested. For the test case, we used the following fourth order univariate function:

$$f(x) = 3x^4 + 16x^3 + 5x^2 \quad (6)$$

The training data for the  $x$ -values consisted of 20 uniformly spaced data points in the range -5.3 to 2.2. The values of  $f(x)$  corresponding to the values of  $x$  were obtained by directly substituting the  $x$ -values into the function.

Two neural networks were trained using the training data: (1) a conventional neural network with 5 layers consisting of 1x100x100x100x1 neurons in each layer with tanh activation functions and (2) a  $\pi$ -net V1 polynomial neural network that outputs fourth order polynomials. Each neural network was trained a total of 5 times and the best network was chosen based on the normalized MSE test loss,  $L(y_{obs}, y_{pred})$ , as well as the visual fit of the neural network’s predictions against the known data.

Results from the neural networks with the best fit are shown in Fig 3. While not shown, conventional neural networks with ReLU activation functions produced similar results. The fourth order  $\pi$ -net V1 polynomial neural network reliably reproduced the same model, due to the polynomial constraint defined by its architecture, whereas there was some variability in the final conventional neural network model. The conventional neural network exhibits Gibbs phenomenon, an oscillatory behavior around the observed data usually caused by discontinuities, which is typical for approximation functions such as Fourier series, orthogonal polynomials, splines, and wavelets [Wilbraham, 1848, Jerri, 2013]. Since a neural network is another type of approximation function, this behavior is not surprising. Gibbs phenomenon was found every time the conventional neural network was fit to the data. The model in Fig 3 was chosen because it exhibited the least amount of Gibbs phenomenon. Additionally, we observed that ReLU activation functions produce Gibbs phenomenon to a lesser extent than tanh activation functions. We suspect that this is because tanh functions bound the output between -1 and 1, whereas ReLU functions bound the output between 0 and  $\infty$ , which leads to fewer discontinuities.

Neural networks are known to be poor at making predictions outside of their training region, which Fig 3 also shows. When the prediction range is slightly extended, the conventional neural network is completely unable to make an accurate prediction, whereas the polynomial neural network makes accurate predictions in this range because it learned the functional form of the data.

The main advantage of polynomial neural networks is the ability to directly obtain a symbolic representation of the network. As shown in Fig. 3, coefficients of the polynomial accurate to 7 significant figures were directly predicted from 20 observations. After performing equation recovery, the extra terms can be systematically removed to simplify the model. In this case, the red terms are small enough to be dropped after further analysis by the modeler.

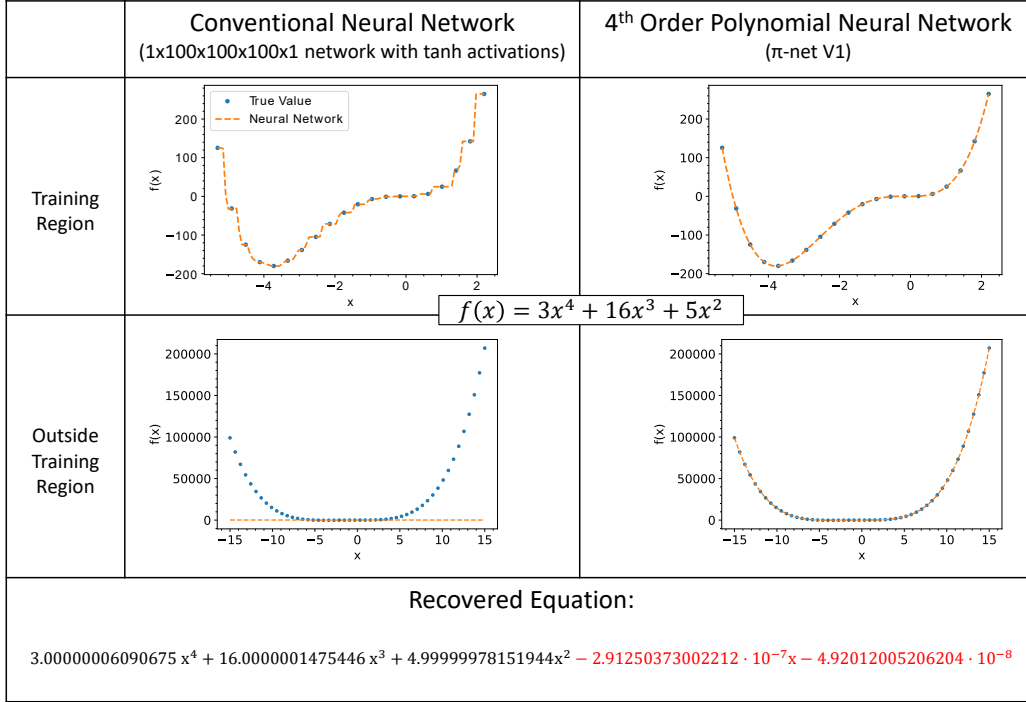


Figure 3: Predictions inside and outside of the training region for a conventional neural network with tanh activation functions, and a fourth order polynomial neural network after learning the fourth order polynomial. The recovered symbolic equation from the polynomial neural network is shown below the figures. The red terms indicate additional terms.

### 3.2 Lotka-Volterra Deterministic Oscillator

Our first demonstration of a polynomial neural ODE is on the deterministic Lotka-Volterra ODE model, which describes the predator-prey population dynamics of a biological system [Lotka, 1925, Volterra, 1926]. When written as a set of first order nonlinear ODEs, the model is given by

$$\frac{dx}{dt} = 1.5x - xy, \quad (7)$$

$$\frac{dy}{dt} = -3y + xy, \quad (8)$$

with initial conditions  $x = 1$  and  $y = 1$ . Since the problem is nonstiff, we generate our training data by integrating the IVP with SciPy and torchdiffeq using DOPRI5, a fourth order embedded method in the Runge–Kutta family of ODE solvers, with the default settings at 200 points uniformly spaced in time between 0 and 10 [Virtanen et al., 2020, Chen, 2021, Dormand and Prince, 1980]. As discussed in the methods section, we batch our data into 199 training samples consisting of IVPs between 2 adjacent known data points, and simultaneously solve the 199 IVPs during each epoch using our own code for the fourth order explicit Runge–Kutta–Fehlberg method, which allows us to directly perform backpropagation through the ODE discretization scheme [Fehlberg, 1968].

Four neural networks were trained using the training data: (1) a conventional neural network with 5 layers consisting of 2x50x50x50x2 neurons in each layer with tanh activation functions and (2) three separate  $\pi$ -net V1 polynomial neural networks of degree two, three, and four.

Following training, the performance of the neural networks was tested by integrating the neural ODE from the initial conditions up to a time value of 40. As shown in Fig 4, both the conventional neural ODE and the fourth order polynomial neural ODE are able to accurately predict the trajectory of the dynamical system beyond the training region for the same initial values as the training data.

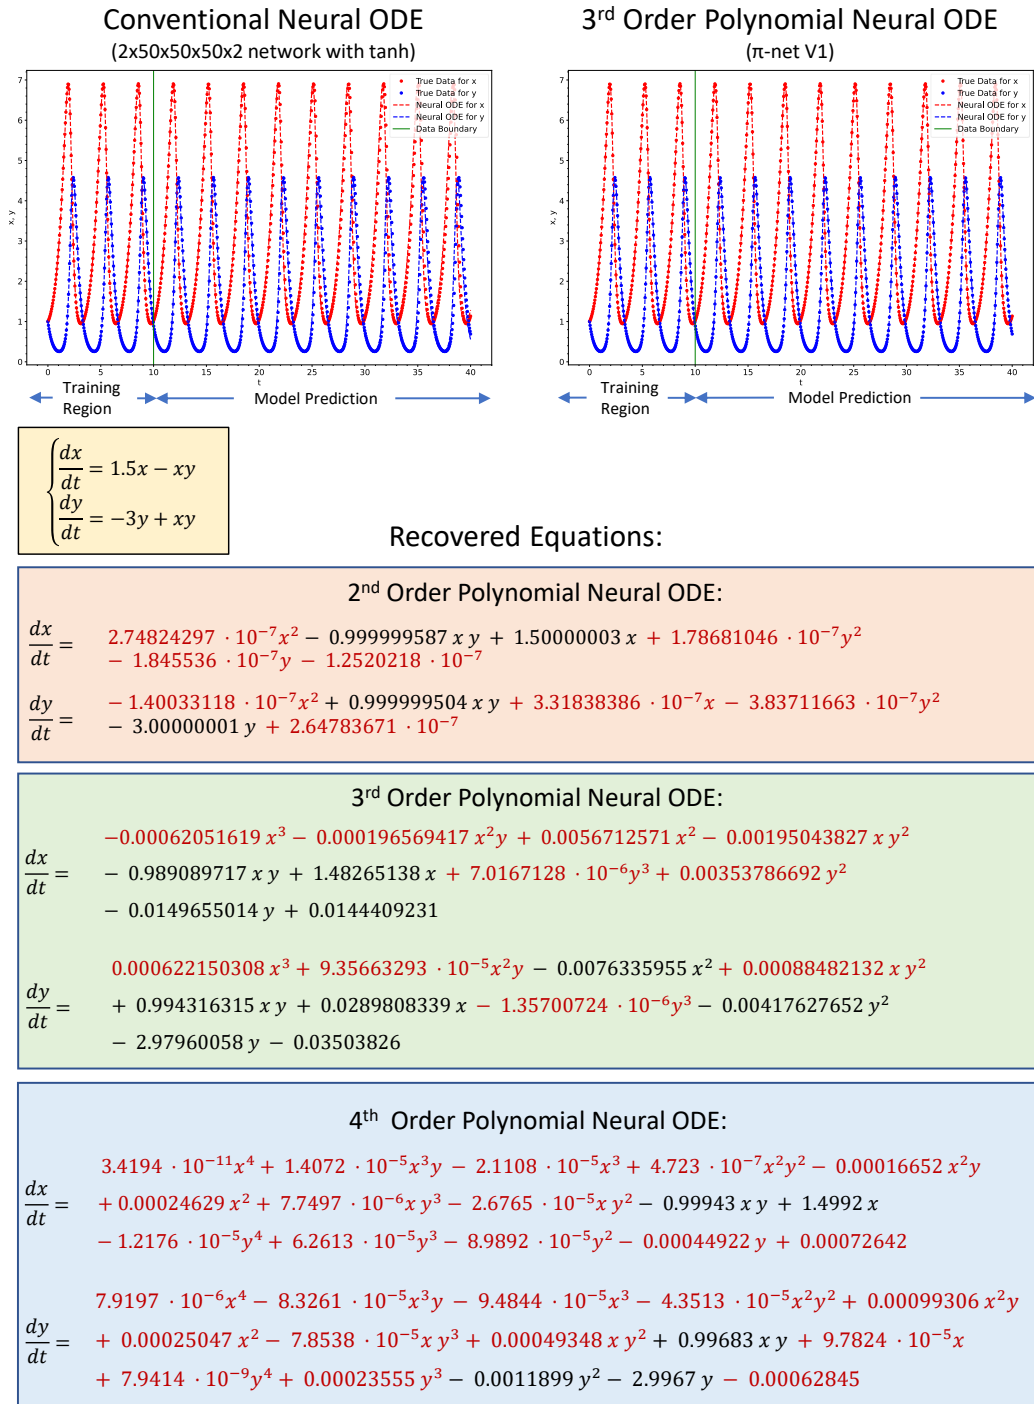


Figure 4: The conventional neural ODE and the  $\pi$ -net V1 polynomial neural ODE were integrated from the initial conditions to a time of 40, to test the performance of the neural ODEs. The vertical green line indicates where the training data ends and the neural ODE begins making a prediction. The symbolic form of the  $\pi$ -net V1 polynomial neural network was obtained for each of the degrees shown. The correct equation is shown above the recovered equations. The red terms indicate terms that can be dropped.



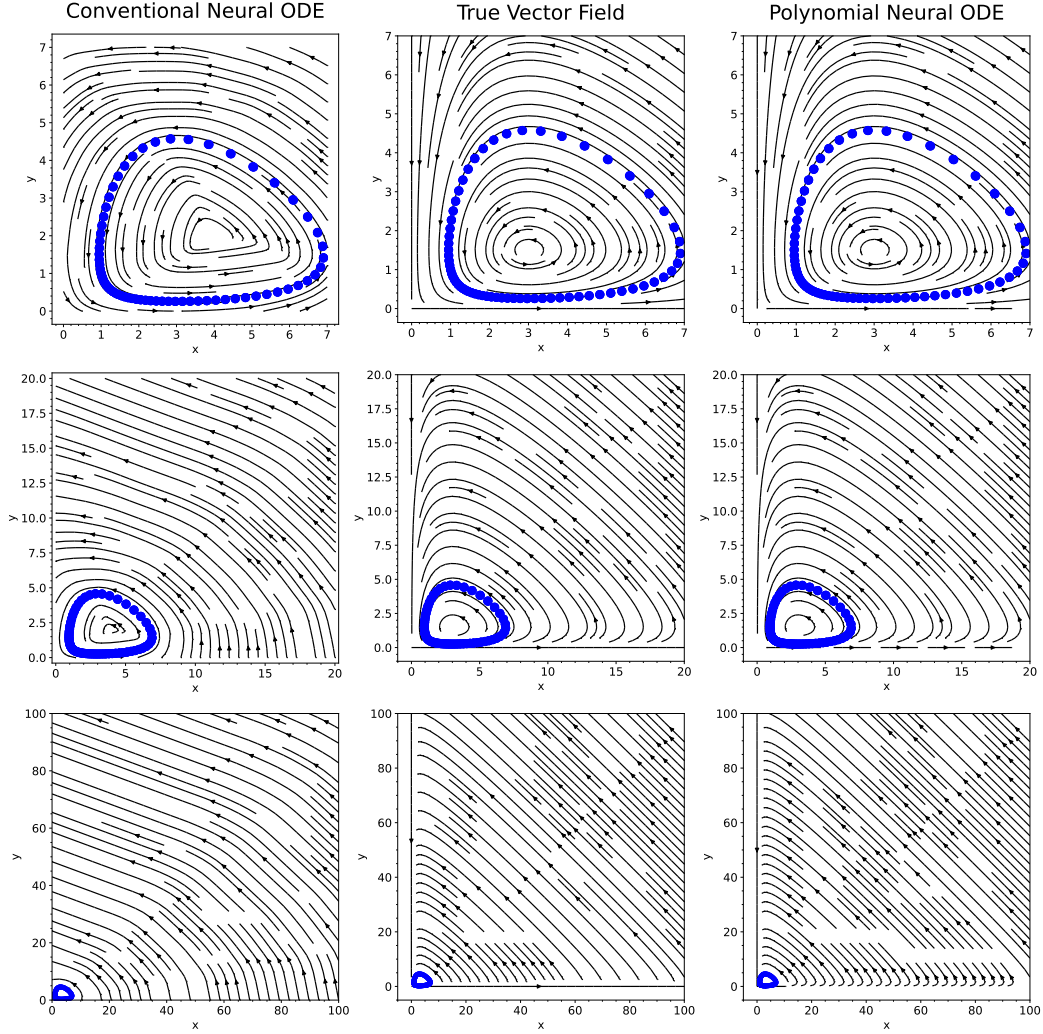


Figure 5: The vector fields of the time derivatives of variables  $x$  and  $y$  are plotted for the conventional neural ODE with tanh activations (left), the true ODE equation (middle), and the second order  $\pi$ -net V1 polynomial neural ode (right). The top row shows the training region, while the remaining rows show the region outside of the training region. The blue circles mark the training data.

Fig 4 doesn't offer much insight about how the conventional neural ODE differs from the polynomial neural ODE, so the vector fields of the time derivatives of  $x$  and  $y$  were also plotted. As shown in Fig 5, the second order  $\pi$ -net V1 neural ODE accurately learns the true vector field for the training region, as well as outside of the training region. The conventional neural ODE learns a close approximation to the training region's vector field, which is why it was able to accurately predict the trajectory of the dynamical system beyond the training region for the same initial values as the training data, as shown in Fig 4. However, it doesn't learn an accurate enough approximation to make predictions for observations outside of its training region. This is clear just by looking at the difference in vector fields for the region of space outside of the training region. On the other hand, the polynomial neural ODE has no problem with the same task.

To show that the  $\pi$ -net V1 polynomial neural ODE learns the same equation regardless of the degree of the polynomial output, three separate polynomial neural ODEs were trained with degrees of two, three, and four. Following training, the symbolic form of each of the polynomial neural ODEs was obtained, as shown in Fig 4. The second, third, and fourth order polynomial neural ODEs were able to recover 6, 2, and 3 significant digits for the coefficients respectively. The values of the coefficients

belonging to the terms not found in the original ODE differ each time, which serves as a clue to the modeler that these terms can be dropped from the final equation.

### 3.3 Damped Oscillatory System

Our next demonstration of a polynomial neural ODE is on the deterministic damped oscillatory system, a popular toy problem for Neural ODEs [Chen et al., 2018, Roesch et al., 2021]. Damped oscillations are common in engineering, physics, and biology [Janson, 2012, Karnopp et al., 1990]. For our work, we refer to the following model as the damped oscillator:

$$\begin{aligned}\frac{dx}{dt} &= -0.1x^3 - 2y^3 \\ \frac{dy}{dt} &= 2x^3 - 0.1y^3,\end{aligned}\tag{9}$$

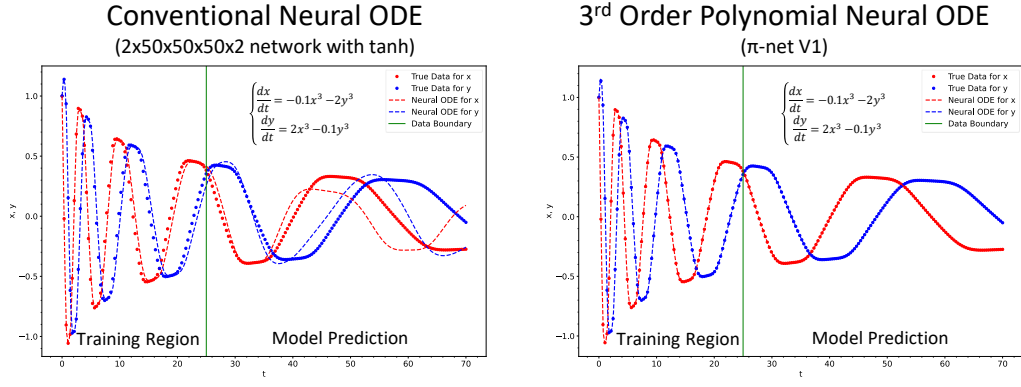
with initial conditions  $x = 1$  and  $y = 1$ . Since the problem is nonstiff, we generate our training data by integrating the IVP with SciPy and torchdiffeq using DOPRI5, a fourth order embedded method in the Runge–Kutta family of ODE solvers, with the default settings at 100 points uniformly spaced in time between 0 and 25 [Virtanen et al., 2020, Chen, 2021, Dormand and Prince, 1980]. As discussed in the methods section, we batch our data into 99 training samples consisting of IVPs between 2 adjacent known data points, and simultaneously solve the 99 IVPs during each epoch using our own code for the fourth order explicit Runge–Kutta–Fehlberg method, which allows us to directly perform backpropagation through the ODE discretization scheme [Fehlberg, 1968].

Three neural networks were trained using the training data: (1) a conventional neural network with 5 layers consisting of  $2 \times 50 \times 50 \times 50 \times 2$  neurons in each layer, with tanh activation functions, and (2) two separate  $\pi$ -net V1 polynomial neural networks of degree three and four. Previous work has demonstrated that this damped oscillatory system is challenging for traditional neural ODEs to learn [Chen et al., 2018, Roesch et al., 2021]. Several authors have made the model easier to learn with neural ODEs by making the first layer of the neural network a cubic function; however, this approach requires a priori knowledge that the functional form of the equation is cubic [Chen, 2021, Roesch et al., 2021]. Since the traditional neural network is supposed to learn the dynamics on its own, we do not use this approach for any of our neural ODEs.

Following training, the performance of the neural networks was tested by integrating the neural ODE from the initial conditions up to a time value of 70. As shown in Fig 6, the third order polynomial neural ODE is able to accurately predict the trajectory of the dynamical system beyond the training region, for the same initial values as the training data. The identical match of vector fields between the third order polynomial and that of the true ODE model, as shown in Fig 7, indicates why the polynomial neural network is able to accurately predict the trajectory.

In contrast, the conventional neural ODE is able to make an accurate prediction only up to about 15 time units past the training region. The conventional neural network’s predictions do not preserve the general shape of the solution to the ODE, whereas the polynomial constraint on the polynomial neural ODE ensures that the functional shape is preserved for the prediction task. Looking at the vector field for the conventional neural ode between -0.5 and 0.5, which is the zoomed in portion of the training region, since the training data does not have many values in this range, the neural ode was unable to learn the center of the vector field’s spiral well. We observe and hypothesize that this is the reason why the conventional neural ODE is known to struggle with learning the damped oscillator model: the errors from the center of the spiral increase over the trajectory as it dampens. Additionally, the zoomed out vector field in the range -1000 to 1000 shows that the conventional neural ODE’s vector field undergoes major distortion outside of its training region. This effect is not observed in the polynomial neural network.

Fig 6 shows the symbolic equations recovered from the third and fourth order polynomial neural ODEs. Five significant figures were recovered from both of the neural odes, and the remaining coefficients are small enough to be dropped.



$$\begin{cases} \frac{dx}{dt} = -0.1x^3 - 2y^3 \\ \frac{dy}{dt} = 2x^3 - 0.1y^3 \end{cases}$$

Recovered Equations:

3<sup>rd</sup> Order Polynomial Neural ODE:

$$\begin{aligned} \frac{dx}{dt} &= -0.10002x^3 - 5.84862 \cdot 10^{-5}x^2y - 1.17748 \cdot 10^{-5}x^2 - 0.000606360xy^2 - 3.08590 \cdot 10^{-5}xy \\ &\quad + 5.28689 \cdot 10^{-6}x - 1.99985y^3 - 1.58209 \cdot 10^{-5}y^2 + 1.43468 \cdot 10^{-5}y + 4.49479 \cdot 10^{-11} \\ \frac{dy}{dt} &= 2.00005x^3 + 0.00107560x^2y - 4.00052 \cdot 10^{-5}x^2 - 1.14539 \cdot 10^{-5}xy^2 - 1.50010 \cdot 10^{-5}xy \\ &\quad + 2.30207 \cdot 10^{-6}x - 0.0999697y^3 - 7.32182 \cdot 10^{-7}y^2 - 8.47415 \cdot 10^{-7}y - 3.30187 \cdot 10^{-11} \end{aligned}$$

4<sup>th</sup> Order Polynomial Neural ODE:

$$\begin{aligned} \frac{dx}{dt} &= -9.55747 \cdot 10^{-7}x^4 + 9.47718 \cdot 10^{-7}x^3y - 0.099997x^3 + 1.22464 \cdot 10^{-5}x^2y^2 + 1.49426 \cdot 10^{-5}x^2y \\ &\quad + 1.31981 \cdot 10^{-6}x^2 + 9.02115 \cdot 10^{-6}xy^3 + 2.69115 \cdot 10^{-5}xy^2 + 2.70065 \cdot 10^{-5}xy + 1.69208 \cdot 10^{-5}x \\ &\quad + 5.09480 \cdot 10^{-6}y^4 - 1.99996y^3 + 7.07770 \cdot 10^{-5}y^2 + 4.59301 \cdot 10^{-5}y - 5.50757 \cdot 10^{-11} \\ \frac{dy}{dt} &= 3.70538 \cdot 10^{-6}x^4 - 4.10939 \cdot 10^{-6}x^3y + 1.99997x^3 - 9.08214 \cdot 10^{-7}x^2y^2 + 6.62073 \cdot 10^{-7}x^2y \\ &\quad + 7.36576 \cdot 10^{-6}x^2 + 7.88325 \cdot 10^{-8}xy^3 - 1.61711 \cdot 10^{-6}xy^2 - 1.62498 \cdot 10^{-5}xy - 3.02663 \cdot 10^{-5}x \\ &\quad + 2.31439 \cdot 10^{-7}y^4 - 0.0999993y^3 + 4.73148 \cdot 10^{-6}y^2 + 1.11502 \cdot 10^{-5}y + 9.48903 \cdot 10^{-12} \end{aligned}$$

Figure 6: The conventional neural ODE and the  $\pi$ -net V1 polynomial neural ODE were integrated from the initial conditions to a time of 70, to test the performance of the neural ODEs. The vertical green line indicates where the training data ends and the neural ODE begins making a prediction. The symbolic forms of the polynomial neural ODEs were also obtained for each of the degrees shown. The red terms indicate additional terms that can be dropped.

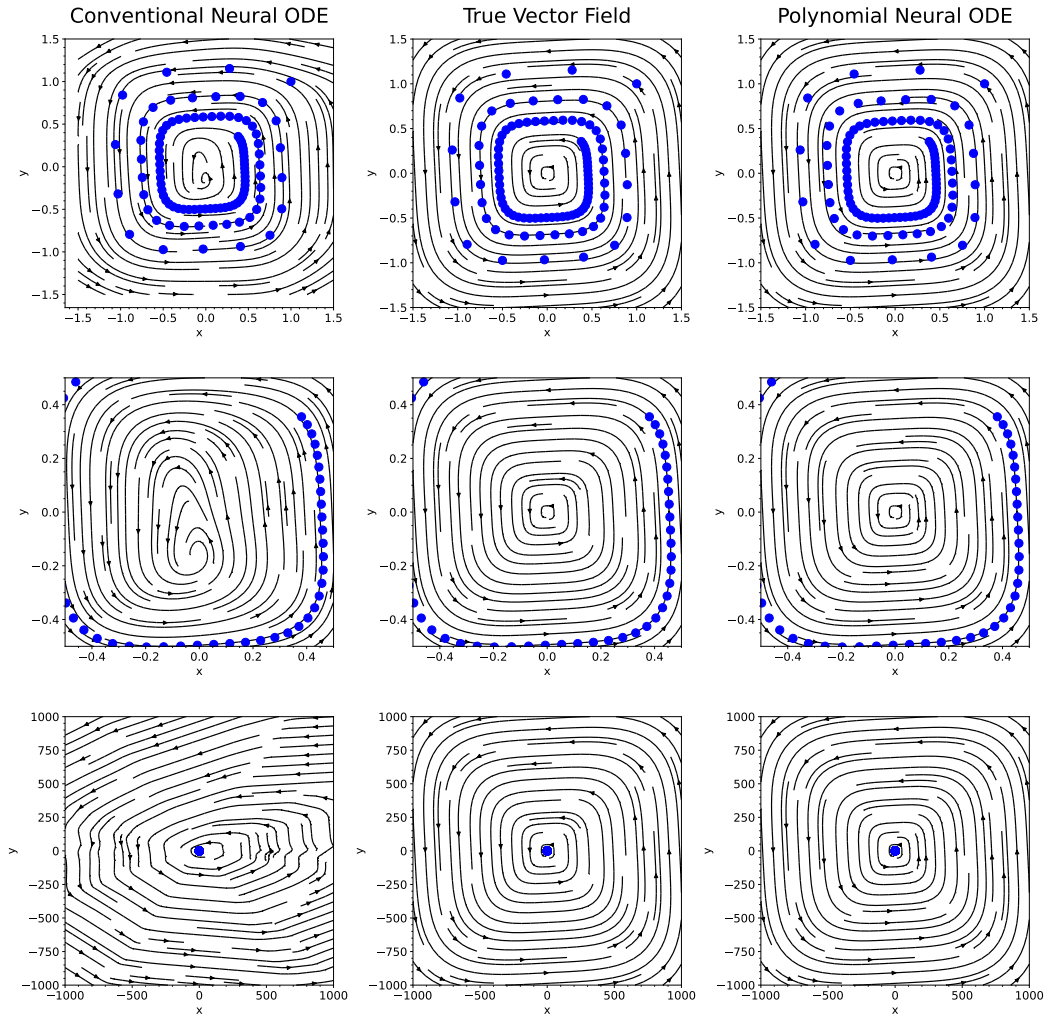


Figure 7: The vector fields of the time derivatives of variables  $x$  and  $y$  are plotted for the conventional neural ODE with tanh activations (left), the true ODE equation (middle), and the third order  $\pi$ -net V1 polynomial neural ode (right). The top row shows the entire training region, the middle row shows a smaller subset of the training region where oscillations dampen, and the bottom row shows the region outside of the training region. The blue circles mark the training data.

### 3.4 Van der Pol Deterministic Oscillator

Our final demonstration of the performance of polynomial neural ODEs is on the Van der Pol oscillator, a nonconservative oscillator with nonlinear damping [van der Pol, 1926]. The Van der Pol equation describes several processes of relaxation-oscillations in the physical and biological sciences. For example, it has been used to model action potentials of neurons, tectonic plates in a geological fault, and oscillations of the left and right vocal cords during speech [FitzHugh, 1961, Nagumo et al., 1962, Cartwright et al., 1999, Lucero and Schoentgen, 2013]. The Van der Pol oscillator is described by the following second order ordinary differential equation:

$$\frac{d^2x}{dt^2} - \mu(1 - x^2)\frac{dx}{dt} + x = 0 \tag{10}$$

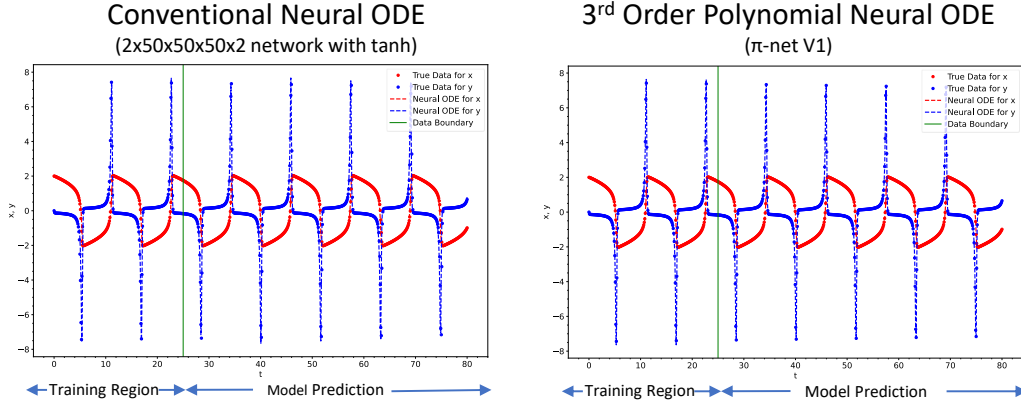
and can be rewritten as a system of first order ODEs:

$$\begin{aligned} \frac{dx}{dt} &= y \\ \frac{dy}{dt} &= \mu y - \mu x^2 y - x. \end{aligned} \tag{11}$$

We chose to assign  $\mu = 5$  and use initial conditions  $x = 2$  and  $\frac{dx}{dt} = y = 0$ . Since the problem is nonstiff, we generate our training data by integrating the IVP with SciPy and torchdiffeq using DOPRI5, a fourth order embedded method in the Runge–Kutta family of ODE solvers, with the default settings at 200 points uniformly spaced in time between 0 and 25 [Virtanen et al., 2020, Chen, 2021, Dormand and Prince, 1980]. As discussed in the methods section, we batch our data into 199 training samples consisting of IVPs between 2 adjacent known data points, and simultaneously solve the 199 IVPs during each epoch using our own code for the fourth order explicit Runge–Kutta–Fehlberg method, which allows us to directly perform backpropagation through the ODE discretization scheme [Fehlberg, 1968]. Three neural networks were trained using the training data: (1) a conventional neural network with 5 layers consisting of  $2 \times 50 \times 50 \times 50 \times 2$  neurons in each layer with tanh activation functions, and (2) two separate  $\pi$ -net V1 polynomial neural networks of degree three and four.

Following training, the performance of the neural networks were tested by: (1) integrating the neural ODEs from the training data’s initial conditions up to a time value of 80 and (2) plotting the vector fields of the time derivatives of the variables  $x$  and  $y$ . As shown in Fig 8 and 10, the third order  $\pi$ -net polynomial neural ODE was able to successfully learn the dynamics of the system. When the equations were recovered from the polynomial neural ODEs, as shown in Fig 8, the coefficients of the original Van der Pol ODE model were successfully recovered to 7 significant digits.

The conventional neural ODE with tanh activation functions was able to correctly predict the dynamical system’s trajectory starting at the training data’s initial conditions, as shown in Fig 8; however, the vector fields in Fig 10 demonstrate that it learned slightly different dynamics from the original Van der Pol ODE the data was generated from. The vector field of the traditional neural ODE does not show the limit cycle found in the Van der Pol ODE. We can show this by integrating the neural ODEs and Van der Pol ODE forward in time starting at a different set of initial conditions. Fig 9 integrates the ODEs forward in time starting at  $x = 4$  and  $y = \frac{dx}{dt} = 4$ . It can be seen that the conventional neural ODE fails to learn the limit cycle, whereas the polynomial neural ODE does not have the same issue.



$$\begin{cases} \frac{dx}{dt} = y \\ \frac{dy}{dt} = 5y - 5x^2y - x \end{cases}$$

Recovered Equations:

**3<sup>rd</sup> Order Polynomial Neural ODE:**

$$\frac{dx}{dt} = 7.2012814 \cdot 10^{-9}x^3 - 4.4566047 \cdot 10^{-8}x^2y - 6.7257892 \cdot 10^{-9}x^2 + 1.9197007 \cdot 10^{-8}xy^2 + 3.8386451 \cdot 10^{-8}xy + 1.9028246 \cdot 10^{-8}x - 1.0600239 \cdot 10^{-9}y^3 + 2.5864009 \cdot 10^{-8}y^2 + 1.0000002y - 9.6208355 \cdot 10^{-9}$$

$$\frac{dy}{dt} = -2.2990156 \cdot 10^{-7}x^3 - 5.0000006x^2y - 1.3328585 \cdot 10^{-7}x^2 - 1.6160624 \cdot 10^{-7}xy^2 - 4.4998796 \cdot 10^{-8}xy - 0.99999996x - 5.6293068 \cdot 10^{-9}y^3 + 7.2424686 \cdot 10^{-8}y^2 + 5.0000002y + 2.261526 \cdot 10^{-7}$$

**4<sup>th</sup> Order Polynomial Neural ODE:**

$$\frac{dx}{dt} = -5.7890282 \cdot 10^{-5}x^4 + 8.2630814 \cdot 10^{-5}x^3y - 2.7629706 \cdot 10^{-8}x^3 + 7.2781275 \cdot 10^{-6}x^2y^2 - 4.1394224 \cdot 10^{-7}x^2y + 0.00036728121x^2 + 2.8932511 \cdot 10^{-8}xy^3 + 1.1002618 \cdot 10^{-8}xy^2 - 0.00038938068xy - 5.490018 \cdot 10^{-7}x - 1.5044485 \cdot 10^{-9}y^4 - 3.5699984 \cdot 10^{-10}y^3 + 3.8001216 \cdot 10^{-5}y^2 + 1.0000001y - 0.00059056922$$

$$\frac{dy}{dt} = 2.3800507 \cdot 10^{-5}x^4 - 2.5142894 \cdot 10^{-5}x^3y - 6.0889057 \cdot 10^{-7}x^3 - 4.8105897 \cdot 10^{-6}x^2y^2 - 4.9999998x^2y - 0.00013980035x^2 - 3.5792187 \cdot 10^{-8}xy^3 - 3.0740606 \cdot 10^{-7}xy^2 + 0.00013637068xy - 0.99999866x + 1.967098 \cdot 10^{-9}y^4 + 1.0625662 \cdot 10^{-8}y^3 - 1.0153587 \cdot 10^{-5}y^2 + 5.0000002y + 0.00021566005$$

Figure 8: The conventional neural ODE and the  $\pi$ -net V1 polynomial neural ODE were integrated from the initial conditions to a time of 80 to test the performance of the neural ODEs. The vertical green line indicates where the training data ends and the neural ODE begins making a prediction. The symbolic forms of the polynomial neural ODEs were also obtained for each of the degrees shown. The red terms indicate additional terms that can be dropped.




Figure 9: To test the ability of the neural ODEs to learn the limit cycle of the Van der Pol oscillator, the various models were integrated forward in time from initial conditions  $x = 4$  and  $y = \frac{dx}{dt} = 4$ . The conventional neural ODE with tanh activation functions did not learn the limit cycle from the data, whereas the third order polynomial neural ODE successfully did.

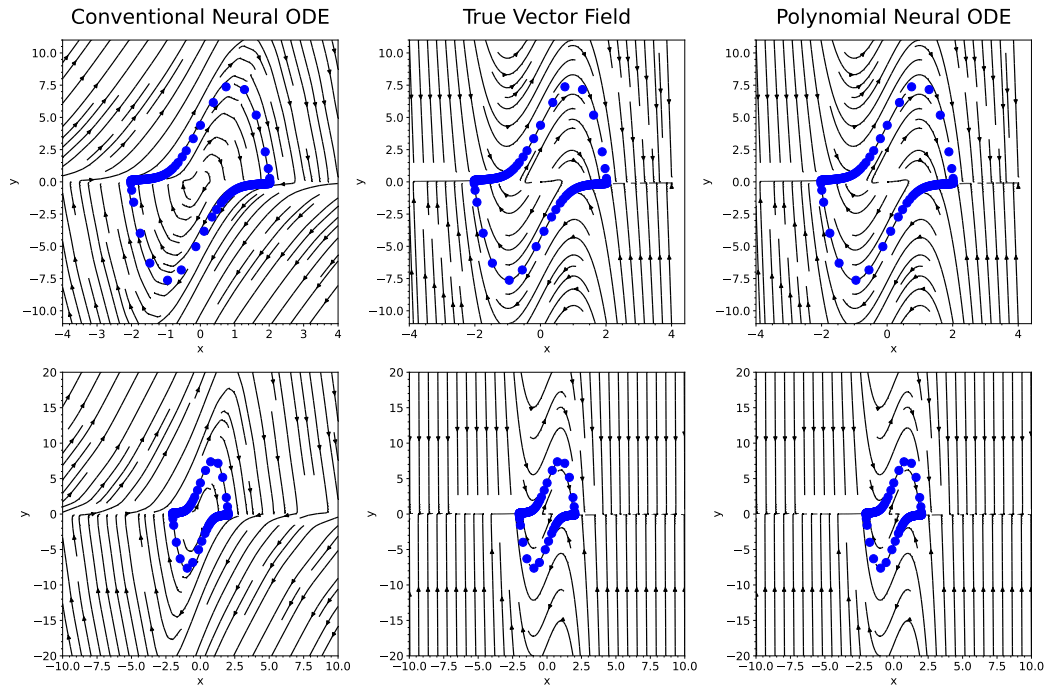


Figure 10: The vector fields of the time derivatives of variables  $x$  and  $y$  are plotted for the conventional neural ODE with tanh activations (left), the true ODE equation (middle), and the third order  $\pi$ -net V1 polynomial neural ode (right). The top row shows the entire training region, while the bottom row shows the region outside of the training region. The blue circles mark the training data.

### 3.5 Learning Other Types of Equations with Polynomial Neural ODEs

All of our previous demonstrations were on ODEs with equations defined by polynomials, so our last demonstration is on a dynamical system which is not a polynomial:

$$\begin{aligned} \frac{dx}{dt} &= \cos(y) \\ \frac{dy}{dt} &= -\sin(x). \end{aligned} \tag{12}$$

with initial conditions  $x = 0.5$  and  $y = 1$ . Since the problem is nonstiff, we generate our training data by integrating the IVP with SciPy and torchdiffeq using DOPRI5, a fourth order embedded method in the Runge–Kutta family of ODE solvers, with the default settings at 200 points uniformly spaced in time between 0 and 40 [Virtanen et al., 2020, Chen, 2021, Dormand and Prince, 1980]. As discussed in the methods section, we batch our data into 199 training samples consisting of IVPs between 2 adjacent known data points, and simultaneously solve the 199 IVPs during each epoch using the fourth order explicit Runge–Kutta–Fehlberg method [Fehlberg, 1968]. The following neural networks were trained using the training data: (1) a conventional neural network with 5 layers consisting of  $2 \times 50 \times 50 \times 50 \times 2$  neurons in each layer with tanh activation functions, and (2) separate  $\pi$ -net V1 polynomial neural networks of degrees 4, 5, 6, and 15.

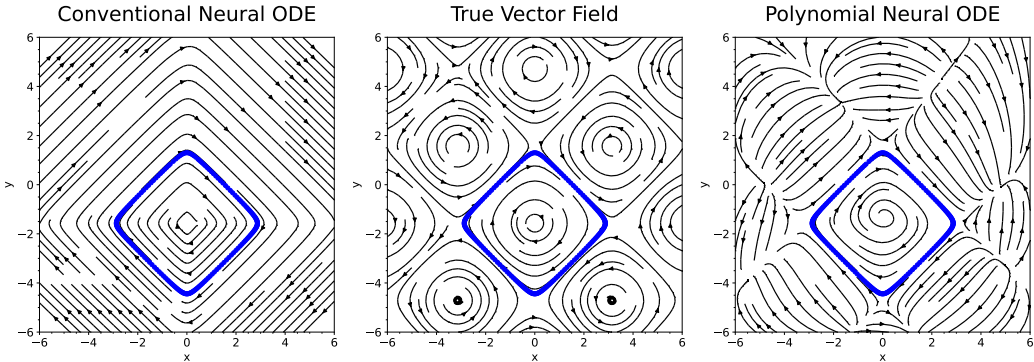


Figure 11: The vector fields of the time derivatives of variables  $x$  and  $y$  are plotted for the conventional neural ODE with hidden layers  $2 \times 50 \times 50 \times 50 \times 2$  and tanh activations (left), the true ODE equation (middle), and the sixth order  $\pi$ -net V1 polynomial neural ode (right). The blue circles mark the training data.

Following training, the performance of the neural networks were tested by plotting the vector fields of the time derivatives of the variables  $x$  and  $y$ . As shown in Fig 11, the polynomial neural ODE was able to learn a better approximation to the true vector field than the conventional neural ODE. While the polynomial neural ODE does not learn a perfect approximation to the field, it is better able to predict a rough approximation of the structure of the true vector field, including where additional spirals in the vector field would be, whereas the conventional neural ODE only learned one spiral.

## 4 Discussion/Conclusion

This work introduced (1) symbolic regression with deep polynomial neural networks and (2) the polynomial neural ODE, which is a deep polynomial neural network implemented into the neural ODE framework. Deep polynomial neural networks directly output a polynomial transformation of the input, which makes them directly interpretable. We are the first to show that symbolic computing can be used to directly recover a symbolic representation of deep polynomial neural networks. We demonstrated successful symbolic regression with a deep polynomial neural network on data generated from a univariate fourth order polynomial. We also successfully demonstrated symbolic



regression of dynamical systems governed by ODEs with the polynomial neural ODE on data from the Lotka-Volterra deterministic oscillator, damped oscillatory system, and Van der Pol deterministic oscillator.

We understand that real world experimental data will be noisy, but we have chosen to devote an analysis on experimental noise to a future study. We have a follow up paper on the data requirements for training conventional and polynomial neural ODEs such as noise, sampling frequency, and training size. Additionally, most chemical kinetics systems are stiff ODEs [Shampine and Gear, 1979], which arise from reaction rates that differ by many orders of magnitude, so neural ODEs need to be able to handle stiffness for this application space. This work only shows examples for nonstiff ODEs. Stiff ODEs require special treatment [Kim et al., 2021], which we will show in a follow up paper.

Traditional deep learning approaches from the field of computer science can be used on scientific problems, but this paper makes the case for developing deep learning techniques specifically tailored for scientific applications. Rather than using "black box" data-driven approaches to describe physical phenomenon, we should be creating a suite of mechanistic data-driven approaches. The polynomial neural ODE is one such approach; however, mathematical models usually have additional types of functions such as trigonometric functions and exponential functions. The scientific machine learning community will need to work on developing more interpretable neural network architectures including more complicated functions such as these.

## 5 Acknowledgements

The authors acknowledge research funding from National Institutes of Health (NIH) NIBIB Award No. 2-R01-EB014877-04A1. Use was made of computational facilities purchased with funds from the National Science Foundation (CNS-1725797) and administered by the Center for Scientific Computing (CSC). The CSC is supported by the California NanoSystems Institute and the Materials Research Science and Engineering Center (MRSEC; NSF DMR 1720256) at UC Santa Barbara.

This work was supported in part by NSF awards CNS-1730158, ACI-1540112, ACI-1541349, OAC-1826967, OAC-2112167, CNS-2120019, the University of California Office of the President, and the University of California San Diego's California Institute for Telecommunications and Information Technology/Qualcomm Institute. Thanks to CENIC for the 100Gbps networks. The content of the information does not necessarily reflect the position or the policy of the funding agencies, and no official endorsement should be inferred.

## References

- E. P. Alves and F. Fiuza. Robust data-driven discovery of reduced plasma physics models from fully kinetic simulations. In *APS Division of Plasma Physics Meeting Abstracts*, volume 2020 of *APS Meeting Abstracts*, page GO10.006, jan 2020.
- A. Andoni, R. Panigrahy, G. Valiant, and L. Zhang. Learning polynomials with neural networks. In E. P. Xing and T. Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1908–1916, Beijing, China, 22–24 Jun 2014. PMLR.
- G. Balsamo, A. Agusti-Panareda, C. Albergel, G. Arduini, A. Beljaars, J. Bidlot, N. Bousserez, S. Bousssetta, A. Brown, R. Buizza, C. Buontempo, F. Chevallier, M. Choulga, H. Cloke, M. Cronin, M. Dahoui, P. Rosnay, P. Dirmeyer, E. Dutra, and X. Zeng. Satellite and in situ observations for advancing global earth surface modelling: A review. *Remote Sensing*, 10:2038, 12 2018. doi: 10.3390/rs10122038.
- G. Brasseur and D. Jacob. *Modeling of Atmospheric Chemistry*. Cambridge University Press, 2017. ISBN 9781108210959.
- S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016. doi: 10.1073/pnas.1517384113.
- G. Calzolari and W. Liu. Deep learning to replace, improve, or aid cfd analysis in built environment applications: A review. *Building and Environment*, 206:108315, 2021.

- J. H. E. Cartwright, V. M. EGUÍLUZ, E. HERNÁNDEZ-GARCÍA, and O. PIRO. Dynamics of elastic excitable media. *International Journal of Bifurcation and Chaos*, 09(11):2197–2202, 1999. doi: 10.1142/S0218127499001620.
- K. Champion, S. Brunton, and J. Kutz. Discovery of nonlinear multiscale systems: Sampling strategies and embeddings. *SIAM Journal on Applied Dynamical Systems*, 18:312–333, 01 2019. doi: 10.1137/18M1188227.
- R. T. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- R. T. Q. Chen. torchdiffeq, 6 2021. URL <https://github.com/rtqichen/torchdiffeq>.
- G. G. Chrysos, S. Moschoglou, G. Bouritsas, J. Deng, Y. Panagakis, and S. Zafeiriou. Deep polynomial neural networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(8):4021–4034, 2022. doi: 10.1109/TPAMI.2021.3058891.
- R. Dandekar, V. Dixit, M. Tarek, A. Garcia-Valadez, and C. Rackauckas. Bayesian neural ordinary differential equations. *CoRR*, abs/2012.07244, 2020.
- J. Dormand and P. Prince. A family of embedded runge-kutta formulae. *Journal of Computational and Applied Mathematics*, 6(1):19–26, 1980.
- S. Du and J. Lee. On the power of over-parametrization in neural networks with quadratic activation. In *International conference on machine learning*, pages 1329–1338. PMLR, 2018.
- F. Fan, J. Xiong, and G. Wang. Universal approximation with quadratic deep networks. *Neural Networks*, 124:383–392, 2020. ISSN 0893-6080. doi: <https://doi.org/10.1016/j.neunet.2020.01.007>.
- E. Fehlberg. *Classical fifth-, sixth-, seventh-, and eighth-order Runge-Kutta formulas with stepsize control*. National Aeronautics and Space Administration, 1968.
- R. FitzHugh. Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, 1(6):445–466, 1961. ISSN 0006-3495. doi: [https://doi.org/10.1016/S0006-3495\(61\)86902-6](https://doi.org/10.1016/S0006-3495(61)86902-6).
- I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- J. Gutkind. *Signaling Networks and Cell Cycle Control: The Molecular Basis of Cancer and Other Diseases*. Cancer Drug Discovery and Development. Humana Press, 2000. ISBN 9781592592180.
- M. Hoffmann, C. Fröhner, and F. Noé. Reactive sindy: Discovering governing reactions from concentration data. *The Journal of Chemical Physics*, 150(2):025101, 2019. doi: 10.1063/1.5066099.
- R. Horn, R. Horn, and C. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, 1994. ISBN 9780521467131.
- N. B. Janson. Non-linear dynamics of biological systems. *Contemporary Physics*, 53(2):137–168, 2012. doi: 10.1080/00107514.2011.644441.
- A. Jerri. *The Gibbs Phenomenon in Fourier Analysis, Splines and Wavelet Approximations*. Mathematics and Its Applications. Springer US, 2013.
- T. Kalsoom, N. Ramzan, S. Ahmed, and M. Ur Rehman. Advances in sensor technologies in the era of smart factory and industry 4.0. *Sensors*, 20:6783, 11 2020. doi: 10.3390/s20236783.
- D. Karnopp, D. L. Margolis, and R. C. Rosenberg. *System dynamics*. Wiley New York, 1990.
- S. Kim, W. Ji, S. Deng, Y. Ma, and C. Rackauckas. Stiff neural ordinary differential equations. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 31(9):093122, 2021. doi: 10.1063/5.0060697.

- K. Lee, N. Trask, and P. Stinis. Structure-preserving sparse identification of nonlinear dynamics for data-driven modeling. *arXiv preprint arXiv:2109.05364*, 2021.
- X. Li, T.-K. L. Wong, R. T. Q. Chen, and D. Duvenaud. Scalable gradients for stochastic differential equations. In S. Chiappa and R. Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 3870–3882. PMLR, 26–28 Aug 2020.
- S. Liang and R. Srikant. Why deep neural networks for function approximation? *arXiv preprint arXiv:1610.04161*, 2016.
- A. Lotka. *Elements of physical biology*. Williams and Wilkins Company, 1925.
- J. C. Lucero and J. Schoentgen. Modeling vocal fold asymmetries with coupled van der pol oscillators. *Proceedings of Meetings on Acoustics*, 19(1):060165, 2013. doi: 10.1121/1.4798467.
- N. Maher, S. Milinski, and R. Ludwig. Large ensemble climate model simulations: introduction, overview, and future prospects for utilising multiple types of large ensemble. *Earth System Dynamics*, 12(2):401–418, 2021. doi: 10.5194/esd-12-401-2021.
- N. M. Mangan, S. L. Brunton, J. L. Proctor, and J. N. Kutz. Inferring biological networks by sparse identification of nonlinear dynamics. *IEEE Transactions on Molecular, Biological and Multi-Scale Communications*, 2(1):52–63, 2016. doi: 10.1109/TMBMC.2016.2633265.
- L. M. Mayr and D. Bojanic. Novel trends in high-throughput screening. *Current opinion in pharmacology*, 9(5):580–588, 2009.
- A. Meurer, C. P. Smith, M. Paprocki, O. Čertík, S. B. Kirpichev, M. Rocklin, A. Kumar, S. Ivanov, J. K. Moore, S. Singh, T. Rathnayake, S. Vig, B. E. Granger, R. P. Muller, F. Bonazzi, H. Gupta, S. Vats, F. Johansson, F. Pedregosa, M. J. Curry, A. R. Terrel, v. Roučka, A. Saboo, I. Fernando, S. Kulal, R. Cimman, and A. Scopatz. Sympy: symbolic computing in python. *PeerJ Computer Science*, 3:e103, Jan. 2017. ISSN 2376-5992. doi: 10.7717/peerj-cs.103.
- J. Nagumo, S. Arimoto, and S. Yoshizawa. An active pulse transmission line simulating nerve axon. *Proceedings of the IRE*, 50(10):2061–2070, 1962. doi: 10.1109/JRPROC.1962.288235.
- C. Rackauckas, Y. Ma, J. Martensen, C. Warner, K. Zubov, R. Supekar, D. Skinner, and A. J. Ramadhan. Universal differential equations for scientific machine learning. *CoRR*, abs/2001.04385, 2020.
- D. A. Randall, R. A. Wood, S. Bony, R. Colman, T. Fichet, J. Fyfe, V. Kattsov, A. Pitman, J. Shukla, J. Srinivasan, et al. Climate models and their evaluation. In *Climate change 2007: The physical science basis. Contribution of Working Group I to the Fourth Assessment Report of the IPCC (FAR)*, pages 589–662. Cambridge University Press, 2007.
- E. Roesch, C. Rackauckas, and M. Stumpf. Collocation based training of neural ordinary differential equations. *Statistical Applications in Genetics and Molecular Biology*, 20, 07 2021. doi: 10.1515/sagmb-2020-0025.
- J. Royle and R. Dorazio. *Hierarchical Modeling and Inference in Ecology: The Analysis of Data from Populations, Metapopulations and Communities*. Elsevier Science, 2008. ISBN 9780080559254.
- Y. Rubanova, R. T. Q. Chen, and D. K. Duvenaud. Latent ordinary differential equations for irregularly-sampled time series. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- S. Rudy, S. Brunton, J. Proctor, and J. Kutz. Data-driven discovery of partial differential equations. *Science Advances*, 3, 09 2016. doi: 10.1126/sciadv.1602614.
- G. Sanguinetti and V. Huynh-Thu. *Gene Regulatory Networks: Methods and Protocols*. Methods in Molecular Biology. Springer New York, 2018. ISBN 9781493988815.

- L. F. Shampine and C. W. Gear. A user's view of solving stiff ordinary differential equations. *SIAM Review*, 21(1):1–17, 1979. doi: 10.1137/1021001.
- H. Singh and J. Dhar. *Mathematical Population Dynamics and Epidemiology in Temporal and Spatio-Temporal Domains*. Apple Academic Press, 2018. ISBN 9781351251686.
- M. Sorokina, S. Sygletos, and S. Turitsyn. Sparse identification for nonlinear optical communication systems: Sino method. *Opt. Express*, 24(26):30433–30443, Dec 2016. doi: 10.1364/OE.24.030433.
- P. Szymański, M. Markowicz, and E. Mikiciuk-Olasik. Adaptation of high-throughput screening in drug discovery—toxicological screening tests. *International journal of molecular sciences*, 13(1): 427–452, 2011.
- B. van der Pol. On relaxation-oscillations. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):978–992, 1926. doi: 10.1080/14786442608564127.
- P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17: 261–272, 2020. doi: 10.1038/s41592-019-0686-2.
- V. Volterra. *Variazioni e fluttuazioni del numero d'individui in specie animali conviventi*. Società anonima tipografica" Leonardo da Vinci", 1926.
- H. Wilbraham. On a certain periodic function. *Cambridge and Dublin Mathematical Journal*, pages 198–201, 1848.