

Universal Differential Equations for Scientific Machine Learning

Christopher Rackauckas (✉ crackauc@mit.edu)

Massachusetts Institute of Technology <https://orcid.org/0000-0001-5850-0663>

Yingbo Ma

Julia Computing

Julius Martensen

University of Bremen <https://orcid.org/0000-0003-4143-3040>

Collin Warner

Massachusetts Institute of Technology

Kirill Zubov

Saint Petersburg State University <https://orcid.org/0000-0003-0441-449X>

Rohit Supekar

Massachusetts Institute of Technology

Dominic Skinner

Massachusetts Institute of Technology <https://orcid.org/0000-0002-2698-041X>

Ali Ramadhan

Massachusetts Institute of Technology

Alan Edelman

Massachusetts Institute of Technology

Article

Keywords: machine learning, generalized approaches, modeling methodology

Posted Date: August 31st, 2020

DOI: <https://doi.org/10.21203/rs.3.rs-55125/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Universal Differential Equations for Scientific Machine Learning

Christopher Rackauckas^{a,b}, Yingbo Ma^c, Julius Martensen^d, Collin Warner^a, Kirill Zubov^e, Rohit Supekar^a, Dominic Skinner^a, Ali Ramadhan^a, and Alan Edelman^a

^aMassachusetts Institute of Technology

^bUniversity of Maryland, Baltimore

^cJulia Computing

^dUniversity of Bremen

^eSaint Petersburg State University

August 26, 2020

Abstract

1
2 In the context of science, the well-known adage “a picture is worth a
3 thousand words” might well be “a model is worth a thousand datasets.”
4 Scientific models, such as Newtonian physics or biological gene regula-
5 tory networks, are human-driven simplifications of complex phenomena
6 that serve as surrogates for the countless experiments that validated the
7 models. Recently, machine learning has been able to overcome the in-
8 accuracies of approximate modeling by directly learning the entire set of
9 nonlinear interactions from data. However, without any predetermined
10 structure from the scientific basis behind the problem, machine learning
11 approaches are flexible but data-expensive, requiring large databases of
12 homogeneous labeled training data. A central challenge is reconciling data
13 that is at odds with simplified models without requiring “big data”.

14 In this work demonstrate how a mathematical object, which we de-
15 note universal differential equations (UDEs), can be utilized as a theo-
16 retical underpinning to a diverse array of problems in scientific machine
17 learning to yield efficient algorithms and generalized approaches. The
18 UDE model augments scientific models with machine-learnable structures
19 for scientifically-based learning. We show how UDEs can be utilized to
20 discover previously unknown governing equations, accurately extrapolate
21 beyond the original data, and accelerate model simulation, all in a time
22 and data-efficient manner. This advance is coupled with open-source soft-
23 ware that allows for training UDEs which incorporate physical constraints,
24 delayed interactions, implicitly-defined events, and intrinsic stochasticity
25 in the model. Our examples show how a diverse set of computationally-
26 difficult modeling issues across scientific disciplines, from automatically

27 discovering biological mechanisms to accelerating the training of physics-
28 informed neural networks and large-eddy simulations, can all be trans-
29 formed into UDE training problems that are efficiently solved by a single
30 software methodology.

31 Recent advances in machine learning have been dominated by deep learning
32 which utilizes readily available “big data” to solve previously difficult problems
33 such as image recognition [1, 2, 3] and natural language processing [4, 5, 6].
34 While some areas of science have begun to generate the large amounts of data
35 required to train deep learning models, notably bioinformatics [7, 8, 9, 10, 11],
36 in many areas the expense of scientific experiments has prohibited the effective-
37 ness of these ground breaking techniques. In these domains, such as aerospace
38 engineering, quantitative systems pharmacology, and macroeconomics, mecha-
39 nistic models which synthesize the knowledge of the scientific literature are still
40 predominantly deployed due to the inaccuracy of deep learning techniques with
41 small training datasets. While these mechanistic models are constrained to be
42 predictive by utilizing prior structural knowledge conglomerated throughout the
43 scientific literature, the data-driven approach of machine learning can be more
44 flexible and allows one to drop the simplifying assumptions required to derive
45 theoretical models. The purpose of this work is to bridge the gap by merging
46 the best of both methodologies while mitigating the deficiencies.

47 It has recently been shown to be advantageous to merge differential equa-
48 tions with machine learning. Physics-Informed Neural Networks (PINNs) utilize
49 partial differential equations in the cost functions of neural networks to incor-
50 porate prior scientific knowledge [12]. While this has been shown to be a form
51 of data-efficient machine learning for some scientific applications, the resulting
52 model does not have the interpretability of mechanistic models. On the other
53 end of the spectrum, machine learning practitioners have begun to make use
54 of scientific structures as a modeling basis for machine learning. For exam-
55 ple, neural ordinary differential equations are initial value problems of the form
56 [13, 14, 15, 16]:

$$u' = \text{NN}_\theta(u, t), \tag{1}$$

57 defined by a neural network NN_θ where θ are the weights. As an example, a
58 neural network with two hidden layers can be written as

$$\text{NN}_\theta(u, t) = W_3\sigma_2(W_2\sigma_1(W_1[u; t] + b_1) + b_2) + b_3, \tag{2}$$

59 where $\theta = (W_1, W_2, W_3, b_1, b_2, b_3)$ where W_i are matrices and b_i are vectors of
60 weights, and (σ_1, σ_2) are the choices of activation functions. Because the embed-
61 ded function is a universal approximator (UA), it follows that NN_θ can learn to
62 approximate any sufficiently regular differential equation. However, the result-
63 ing model is defined without direct incorporation of known mechanisms. The
64 Universal Approximation Theorem (UAT) demonstrates that sufficiently large
65 neural networks can approximate any nonlinear function with a finite set of
66 parameters [17, 18, 19]. Our approach extends the previous data-driven neural
67 ODE approaches to directly utilize mechanistic modeling simultaneously with
68 UAs in order to allow for arbitrary data-driven model extensions. The objects

69 of this semi-mechanistic approach, which we denote as Universal Differential
70 Equations (UDEs) for universal approximators in differential equations, are dif-
71 ferential equation models where part of the differential equation contains an
72 embedded UA, such as a neural network, Chebyshev expansion, or a random
73 forest.

74 As a motivating example, the universal ordinary differential equation (UODE):

$$u' = f(u, t, U_\theta(u, t)), \quad (3)$$

75 denotes a known mechanistic model form f with missing terms defined by some
76 UA U_θ . While the utility of this object has been seen before in optimal control
77 [20] and model augmentation [21], here we demonstrate that this object can be
78 used to solve a greatly expanded scope of problems and is thus able to be the
79 basis of much research in scientific machine learning, from accelerating models
80 to being a stepping stone in symbolic algorithms. We demonstrate general-
81 izations to incorporate process noise, delayed interactions, and physics-based
82 constraints are given by embedding UAs into stochastic, delay, and differential-
83 algebraic equations respectively. As a fundamental object underlying so many
84 algorithms, it then becomes essential to be able to efficiently train UDEs in
85 any context in which they arise. In Section 1 we describe our methodology and
86 software implementation for efficiently training UDEs of any of these forms in a
87 way that covers stiffness, nonlinear algebraic constraints, stochasticity, delays,
88 parallelism, and more. We then demonstrate that the following abilities within
89 the UDE framework:

- 90 • In Section 2.1 we recover governing equations from much lesser data than
91 prior methods and demonstrate the ability to accurately extrapolate from
92 a short time series.
- 93 • In Section 2.2 we demonstrate the ability to utilize arbitrary conservation
94 laws as prior knowledge in the discovery of dynamical systems.
- 95 • In Section 2.3 we discover the differential operator and nonlinear reaction
96 term of a biological partial differential equation (PDE) from spatiotempo-
97 ral data, demonstrating the interpretability of trained UDEs.
- 98 • In Section 3 We derive an adaptive method for automated solving of a 100-
99 dimensional nonlinear Hamilton-Jacobi-Bellman PDE, the first adaptive
100 method for this class of problems that the authors are aware of.
- 101 • In Section 4.1 we automate the discovery of fast, accurate, and physically-
102 consistent surrogates to accelerate a large-eddy simulation commonly used
103 in the voxels of a climate simulation.
- 104 • In Section 4.2 we approximate closure relations in viscoelastic fluids to
105 accelerate the simulation of a system of 6 differential-algebraic equations
106 by 2x, showing that this methodology is also applicable to small-scale
107 problems.

- In Section 4.3 we demonstrate that discrete physics-informed neural networks fall into a subclass of universal ODEs and extend previous methods directly through this formalism.

1 Efficient Training of Universal Differential Equations via Differentiable Programming

Training a UDE amounts to minimizing a cost function $C(\theta)$ defined on the current solution $u_\theta(t)$, the current solution to the differential equation with respect to the choice of parameters θ . One choice is the Euclidean distance $C(\theta) = \sum_i \|u_\theta(t_i) - d_i\|$ at discrete data points (t_i, d_i) . When optimized with local derivative-based methods, such as stochastic gradient decent, ADAM [22], or L-BFGS [23], this requires the calculation of $\frac{dC}{d\theta}$ which by the chain rule amounts to calculating $\frac{du}{d\theta}$. Thus the problem of efficiently training a UDE reduces to calculating gradients of the differential equation solution with respect to parameters.

In certain special cases there exist efficient methods for calculating these gradients called adjoints [24, 25, 26, 27, 28]. The asymptotic computational cost of these methods does not grow multiplicatively with the number of state variables and parameters like numerical or forward sensitivity approaches, and thus it has been shown empirically that adjoint methods are more efficient on large parameter models [29, 30]. However, given the number of different families of UDE models we wish to train, we generalize to a differentiable programming framework with reverse-mode accumulation in order to allow for deriving on-the-fly approximations for the wide range of possible differential equation types.

Given a function $f(x) = y$, the pullback at x is the function:

$$B_f^x(y) = y^T f'(x), \quad (4)$$

where $f'(x)$ is the Jacobian J . We note that $B_f^x(1) = (\nabla f)^T$ for a function f producing a scalar output, meaning the pullback of a cost function computes the gradient. A general computer program can be written as the composition of discrete steps:

$$f = f^L \circ f^{L-1} \circ \dots \circ f^1, \quad (5)$$

and thus the vector-Jacobian product can be decomposed:

$$v^T J = (\dots ((v^T J_L) J_{L-1}) \dots) J_1, \quad (6)$$

which allows for recursively decomposing a the pullback to a primitively known set of $\mathcal{B}_{f^i}^x$:

$$\mathcal{B}_f^x(A) = \mathcal{B}_{f^1}^x \left(\dots \left(\mathcal{B}_{f^{L-2}}^{x_{L-2}} \left(\mathcal{B}_{f^{L-1}}^{x_{L-1}}(A) \right) \right) \dots \right), \quad (7)$$

139 where $x_i = (f^i \circ f^{i-1} \circ \dots \circ f^1)(x)$. Implementations of code generation for the
140 backwards pass of an arbitrary program in a dynamic programming language can
141 vary. For example, building a list of function compositions (a tape) is provided
142 by libraries such as Tracker.jl [31] and PyTorch [32], while other libraries perform
143 direct generation of backward pass source code such as Zygote.jl [33], TAF [34],
144 and Tapenade [35].

145 The open-source differential equation solvers of DifferentialEquations.jl [36]
146 were developed in a manner such that all steps of the programs have a well-
147 defined pullback when using a Julia-based backwards pass generation system.
148 Our software allows for automatic differentiation to be utilized over differen-
149 tial equation solves without any modification to the user code. This enables
150 the simulation software already written with DifferentialEquations.jl, including
151 large software infrastructures such as the MIT-CalTech CLiMA climate mod-
152 eling system [37] and the QuantumOptics.jl simulation framework [38], to be
153 compatible with all of the techniques mentioned in the rest of the paper. Thus
154 while we detail our results in isolation from these larger simulation frameworks,
155 the UDE methodology can be readily used in full-scale simulation packages
156 which are already built on top of the Julia SciML ecosystem.

157 The full set of adjoint options, which includes continuous adjoint methods
158 and pure reverse-mode AD approaches, is described in Supplement S1. Meth-
159 ods via solving ODEs in reverse [16] are the common adjoint utilized in neural
160 ODE software such as torchdiffeq and are $O(1)$ in memory, but are known to
161 be unstable under certain conditions such as on stiff equations [39]. Check-
162 pointed interpolation adjoints [27] and continuous quadrature approaches are
163 available which do not require stable reversibility of the ODEs while retain-
164 ing a relatively low-memory implementation via checkpointing (in particular
165 Section 2.2 and 4.1 are noted as a cases which are not stable under the re-
166 versed adjoint but stable under the checkpointing adjoint approach). These
167 adjoint methods fall under the continuous optimize-then-discretize approach.
168 Through the differentiable programming integration, discrete adjoint sensitivity
169 analysis [40, 41] is implemented through both tape-based reverse-mode [42] and
170 source-to-source translation [33], with computational trade-offs between the two
171 approaches. The former can be faster on scalarized heterogeneous differential
172 equations while the latter is more optimized for homogeneous vectorized func-
173 tions calls like are demonstrated in neural networks and discretizations of partial
174 differential equations. Full discretize-then-optimize is implemented using this
175 package by utilizing the step-wise integrator interface in conjunction with these
176 discrete adjoints of the steps. Continuous and discrete forward mode sensitivity
177 analysis approaches are also provided and optimized for equations with smaller
178 numbers of parameters.

179 Previous research has shown that the discrete adjoint approach is more stable
180 than continuous adjoints in some cases [43, 39, 44, 45, 46, 47] while continuous
181 adjoints have been demonstrated to be more stable in others [48, 45] and can
182 reduce spurious oscillations [49, 50, 51]. This trade-off between discrete and
183 continuous adjoint approaches has been demonstrated on some equations as a
184 trade-off between stability and computational efficiency [52, 53, 54, 55, 56, 57,

185 58, 59, 60]. Care has to be taken as the stability of an adjoint approach can
186 be dependent on the chosen discretization method [61, 62, 63, 64, 65], and our
187 software contribution helps researchers switch between all of these optimization
188 approaches in combination with hundreds of differential equation solver methods
189 with a single line of code change.

190 As described in Supplement S1.1, these adjoints utilize reverse-mode auto-
191 matic differentiation for vector transposed Jacobian products within the adjoint
192 definitions to reduce the computational complexity while supporting advanced
193 features like constraint and conservation equations. In addition, the module
194 DiffEqFlux.jl handles compatibility with the Flux.jl neural network library so
195 that these vector Jacobian products are automatically replaced with efficient
196 pullback implementations for embedded deep neural networks (also known as
197 backpropagation) wherever neural networks are encountered in the right hand
198 side of any differential equation definitions. This allows for common deep archi-
199 tectures, such as convolutional neural networks and recurrent neural networks,
200 to be efficiently used as the basis for a UDE without any Jacobians being cal-
201 culated in the full adjoint and without requiring any intervention from users.

202 Using this approach, the solvers are capable of building efficient gradient
203 calculations for training ML-embedded UDEs of the classes:

- 204 • Universal Ordinary Differential Equations (UODEs)
- 205 • Universal Stochastic Differential Equations (USDEs), or universal differ-
206 ential equations with continuous process noise
- 207 • Universal Delay Differential Equations (UDDEs), or universal differential
208 equations with delayed interactions
- 209 • Universal Differential-Algebraic Equations (UDAEs), or universal differ-
210 ential equations with constraint equations and conservation laws
- 211 • Universal Boundary Value Problems (UBVPs), or universal differential
212 equations with final time point constraints
- 213 • Universal Partial Differential Equations (UPDEs)
- 214 • Universal Hybrid (Event-Driven) Differential Equations

215 as well as the combinations, such as stochastic delay differential equations, jump
216 diffusions, and stochastic partial differential equations. A combination of over
217 300 solver methods cover the efficient training of stiff and non-stiff versions
218 of each of these equations, with support for adaptivity, high-order, automatic
219 stiffness detection, sparse differentiation with automatic sparsity detection and
220 coloring [66], Newton-Krylov implicit handling, GPU compatibility, and multi-
221 node parallelism via MPI compatibility. Thus together, semi-mechanistic UDEs
222 of any form can embed machine learning models and be trained using this open-
223 source library with the most effective differential equation solvers for that class
224 of equations.

Feature	Stiff	DAEs	SDEs	DDEs	Stabilized	DtO	GPU	Dist	MT	Sparse
SciML	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
torchdiffeq	0	0	0	0	0	✓	✓	0	0	0
torchsde	0	0	✓	0	0	0	✓	0	0	0
tfdiffeq	0	0	0	0	0	0	✓	0	0	0

Table 1: Feature comparison of ML-augmented differential equation libraries. First first column corresponds to support for stiff ODEs, then DAEs, SDEs, DDEs, stabilized non-reversing adjoints, discretize-then-optimize methods, distributed computing, and multithreading. Sparse refers to automated sparsity handling in Jacobian calculations of implicit methods.

1.1 Features and Performance

225

226 We assessed the viability of alternative differential equation libraries for univer-
227 sal differential equation workflows by comparing the features and performance
228 of the given libraries. Table 1 demonstrates that the Julia SciML ecosystem is
229 the only differential equation solver library with deep learning integration that
230 supports stiff ODEs, DAEs, DDEs, stabilized adjoints, distributed and multi-
231 threaded computation. We note the importance of the stabilized adjoints in
232 Section 4.1 as many PDE discretizations with upwinding exhibit unconditional
233 instability when reversed, and thus this is a crucial feature when training em-
234 bedded neural networks in many PDE applications. Table 2 demonstrates that
235 the SciML ecosystem exhibits more than an order of magnitude performance
236 when solving ODEs against torchdiffeq of up to systems of 1 million equations.
237 Because the adjoint calculation itself is a differential equation, this also corre-
238 sponds to increased training times on scientific models. To reinforce this result,
239 Supplement S2 demonstrates a 100x performance difference over torchdiffeq
240 when training the spiral neural ODE from [16, 43]. We note that the author
241 of the tfdiffeq library has previously concluded “speed is almost the same as the
242 PyTorch (torchdiffeq) codebase ($\pm 2\%$)”. Additionally, Supplement S2 demon-
243 strates a 1,600x performance advantage for the SciML ecosystem over torchsde
244 using the geometric Brownian motion example from the torchsde documentation
245 [67]. Given the computational burden, the mix of stiffness, and non-reversibility
246 of the examples which follow in this paper, these results demonstrate that the
247 SciML ecosystem is the first deep learning integrated differential equation soft-
248 ware ecosystem that can train all of the equations necessary for the results of
249 this paper. Note that this does not infer that our solvers will demonstrate
250 more than an order of magnitude performance difference on all equations, for
251 example very non-stiff ODEs dominated by large dense matrix multiplications
252 like in image classification neural ODEs, but it does demonstrate that on the
253 equations generally derived from scientific models (ODEs derived from PDE dis-
254 cretizations, heterogeneous differential equation systems, and neural networks
255 in sufficiently small systems) that an order of magnitude or more performance
256 difference can exist.

# of ODEs	3	28	768	3,072	12,288	49,152	196,608	786,432
SciML	1.0x	1.0x	1.0x	1.0x	1.0x	1.0x	1.0x	1.0x
SciML DP5	1.0x	1.9x	2.9x	2.9x	2.9x	2.9x	3.4x	2.6x
torchdiffeq dopri5	5,850x	1700x	420x	280x	120x	31x	41x	38x
torchdiffeq adams	7,600x	1100x	710x	490x	170x	44x	47x	43x

Table 2: Relative time to solve for ML-augmented differential equation libraries (smaller is better). Standard non-stiff solver benchmarks from representative scientific systems were taken from [68] as described in Supplement S2. SciML stands for the optimal method choice out of the 300+ from the SciML, which for the first is DP5, for the second is VCABM, and for the rest is ROCK4.

2 Knowledge-Enhanced Model Reconstruction of Biological Models

Automatic reconstruction of models from observable data has been extensively studied. Many methods produce non-symbolic representations by learning functional representations [69, 70] or through dynamic mode decomposition (DMD, eDMD) [71, 72, 73, 74]. Symbolic reconstruction of equations has utilized symbolic regressions which require a prechosen basis [75, 76], or evolutionary methods to grow a basis [77, 78]. However, a common thread throughout much of the literature is that added domain knowledge constrains the problem to allow for more data-efficient reconstruction [79, 80]. Here we detail how a UA embedded workflow can augment existing symbolic regression frameworks to allow for reconstruction from partially known models in a more data-efficient manner.

2.1 Automated Identification of Nonlinear Interactions with Universal Ordinary Differential Equations

The SInDy algorithm [81, 82, 83] finds a sparse basis Ξ over a given candidate library Θ minimizing the objective function $\|\dot{\mathbf{X}} - \Theta\Xi\|_2 + \lambda\|\Xi\|_1$ using data for $\dot{\mathbf{X}}$ generated by interpolating the trajectory data \mathbf{X} . Here we describe a UDE approach to extend SInDy in a way that embeds prior structural knowledge.

As a motivating example, take the Lotka-Volterra system:

$$\begin{aligned} \dot{x} &= \alpha x - \beta xy, \\ \dot{y} &= \gamma xy - \delta y. \end{aligned} \tag{8}$$

Assume that a scientist has a short time series from this system but knows the birth rate of the prey x and the death rate of the predator y . With only this information, a scientist can propose the knowledge-based UODE as:

$$\begin{aligned} \dot{x} &= \alpha x + U_1(x, y), \\ \dot{y} &= -\delta y + U_2(x, y), \end{aligned} \tag{9}$$

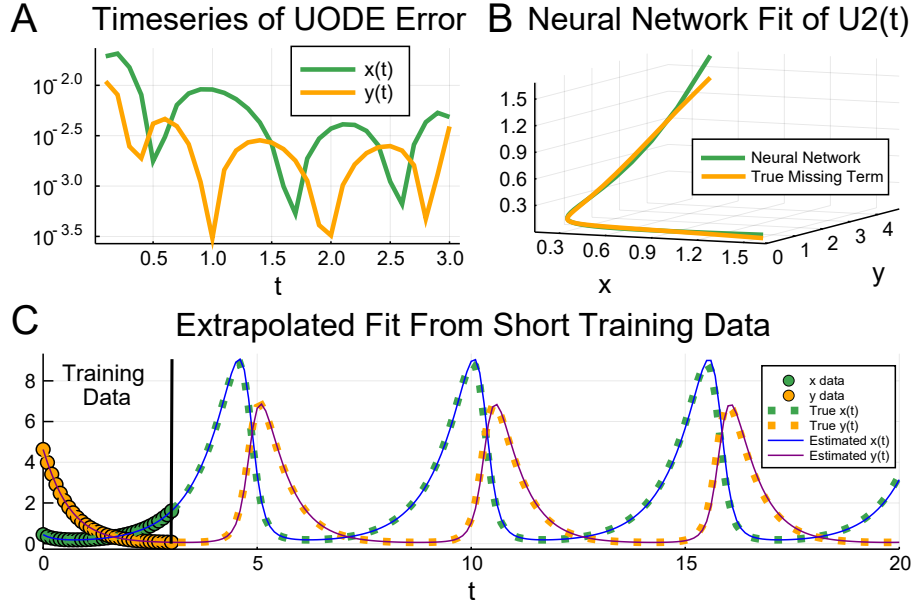


Figure 1: Automated Lotka-Volterra equation discovery with UODE-enhanced SInDy. (A) The error in the trained UODE against $x(t)$ and $y(t)$ in green and yellow respectively. (B) The measured values of the missing term $U_2(x, y)$ throughout the time series, with the neural network approximate in green and the true value γxy in yellow. (C) The extrapolation of the knowledge-enhanced SInDy fit series. The green and yellow dots show the data that was used to fit the UODE, and the dots show the true solution of the Lotka-Volterra Equations 8 beyond the training data. The blue and purple lines show the extrapolated solution how the UODE-enhanced SInDy recovered equations.

279 which is a system of ordinary differential equations that incorporates the known
 280 structure but leaves room for learning unknown interactions between the the
 281 predator and prey populations. Learning the unknown interactions corresponds
 282 training the UA $U : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ in this UODE.

283 While the SInDy method normally approximates derivatives using a spline
 284 over the data points or similar numerical techniques, here we have $U_\theta(x, y)$ as an
 285 estimator of the derivative for only the missing terms of the model and we can
 286 perform a sparse regression on samples from the trained $U_\theta(x, y)$ to reconstruct
 287 only the unknown interaction equations. As described in Supplement S3.1, we
 288 trained $U_\theta(x, y)$ as a neural network against the simulated data for $t \in [0, 3]$ and
 289 utilized a sparse regression techniques [81, 84, 85] on the neural network out-
 290 puts to reconstruct the missing dynamical equations. Using a 10-dimensional
 291 polynomial basis extended with trigonometric functions, the sparse regression
 292 yields 0 for all terms except for the missing quadratic terms, directly learning
 293 the original equations in an interpretable form. Even though the original data

294 did not contain a full period of the cyclic solution, the resulting fit is then able
 295 to accurately extrapolate from the short time series data as shown in Figure 1.
 296 Supplement S3.1 further demonstrates the robustness of the discovery approach
 297 to noise in the data. Likewise, when attempting to learn full ODE with the
 298 original SInDy approach on the same trained data with the analytical deriva-
 299 tive values, we were unable to recover the exact original equations from the
 300 sparse regression, indicating that the knowledge-enhanced approach increases
 301 the robustness equation discovery.

302 We note that collaborators using the preprint of this manuscript have suc-
 303 cessfully demonstrated the ability to construct UODE models which improve
 304 the prediction of Li-ion battery performance [86] and for automated discovery
 305 of droplet physics directly from imaging data, effectively replicating the theo-
 306 retical results of a one and a half year study with a UODE discovery process
 307 which trains in less than an hour [87].

308 2.2 Incorporating Prior Knowledge of Conservation Equa- 309 tions

The extra features of the SciML ecosystem can be utilized to encode more infor-
 mation into the model. For example, when attempting to discover a biological
 chemical reaction network or a chemical combustion network, one may only have
 prior knowledge of the conservation laws between the constituent substrates. As
 a demonstration, in the Robertson equation

$$\frac{dy_1}{dt} = -0.04y_1 + 10^4 y_2 y_3 \quad (10)$$

$$\frac{dy_2}{dt} = 0.04y_1 - 10^4 y_2 y_3 - 3 * 10^7 y_2^2 \quad (11)$$

$$1 = y_1 + y_2 + y_3 \quad (12)$$

one might only have prior knowledge of the conservation equation $1 = y_1 + y_2 + y_3$. In this case, a universal DAE of the form:

$$\frac{d[y_1, y_2]}{dt} = U_\theta(y_1, y_2, y_3) \quad (13)$$

$$1 = y_1 + y_2 + y_3 \quad (14)$$

310 can be utilized to encode this prior knowledge. This can then be trained by
 311 utilizing a singular mass matrix in the form $Mu' = f(u, p, t)$. Supplement
 312 S1’s derivation of the adjoint method describes a new initialization scheme for
 313 index-1 DAEs in mass matrix form which directly solves a linear system for new
 314 consistent algebraic variables in the adjoint pass without requiring the approx-
 315 imate nonlinear iterations of [88], thus further demonstrating the efficiency and
 316 accuracy of the SciML software’s methods for UDE workflows. Supplement S3.2
 317 demonstrates the ability to learn this system utilizing the SciML tools through
 318 this universal DAE approach.

319 **2.3 Reconstruction of Spatial Dynamics with Universal** 320 **Partial Differential Equations**

321 To demonstrate discovery of spatiotemporal equations directly from data, we
322 consider data generated from the one-dimensional Fisher-KPP
323 (Kolmogorov–Petrovsky–Piskunov) PDE [89]:

$$\frac{\partial \rho}{\partial t} = r\rho(1 - \rho) + D\frac{\partial^2 \rho}{\partial x^2}, \quad (15)$$

324 with $x \in [0, 1]$, $t \in [0, T]$, and periodic boundary condition $\rho(0, t) = \rho(1, t)$. Here
325 ρ represents population density of a species, r is the local growth rate and D
326 is the diffusion coefficient. Such reaction-diffusion equations appear in diverse
327 physical, chemical and biological problems [90]. To learn the generated data,
328 we define the UPDE:

$$\rho_t = \text{NN}_\theta(\rho) + \hat{D} \text{CNN}(\rho), \quad (16)$$

329 where NN_θ is a neural network representing the local growth term. The deriva-
330 tive operator is approximated as a convolutional neural network CNN, a learn-
331 able arbitrary representation of a stencil while treating the coefficient \hat{D} as an
332 unknown. We encode in the loss function extra constraints to ensure the learned
333 equation is physically realizable, i.e. the derivative stencil must be conservative
334 (the coefficients sum to zero), as described in Supplement S4. Figure 2 shows
335 the result of training the UPDE against the simulated data, which recovers the
336 canonical $[1, -2, 1]$ stencil of the one-dimensional Laplacian and the diffusion
337 constant while simultaneously finding a neural representation of the unknown
338 quadratic growth term. We note that the differentiable programming integration
339 in conjunction with the Flux.jl deep learning framework allows for the adjoints
340 to automatically utilize efficient backpropagation of the embedded convolutional
341 neural networks and automatically utilizes the fast kernels provided by cudnn
342 when trained using GPUs.

343 **3 Computationally-Efficient Solving of** 344 **High-Dimensional Partial Differential** 345 **Equations**

346 It is impractical to solve high dimensional PDEs with mesh-based techniques
347 since the number of mesh points scales exponentially with the number of dimen-
348 sions. Given this difficulty, mesh-free methods based on universal approximators
349 such as neural networks have been constructed to allow for direct solving of high
350 dimensional PDEs [91, 92]. Recently, methods based on transforming partial
351 differential equations into alternative forms, such as backwards stochastic differ-
352 ential equations (BSDEs), which are then approximated by neural networks have
353 been shown to be highly efficient on important equations such as the nonlinear
354 Black-Scholes and Hamilton-Jacobi-Bellman (HJB) equations [93, 94, 95, 96].

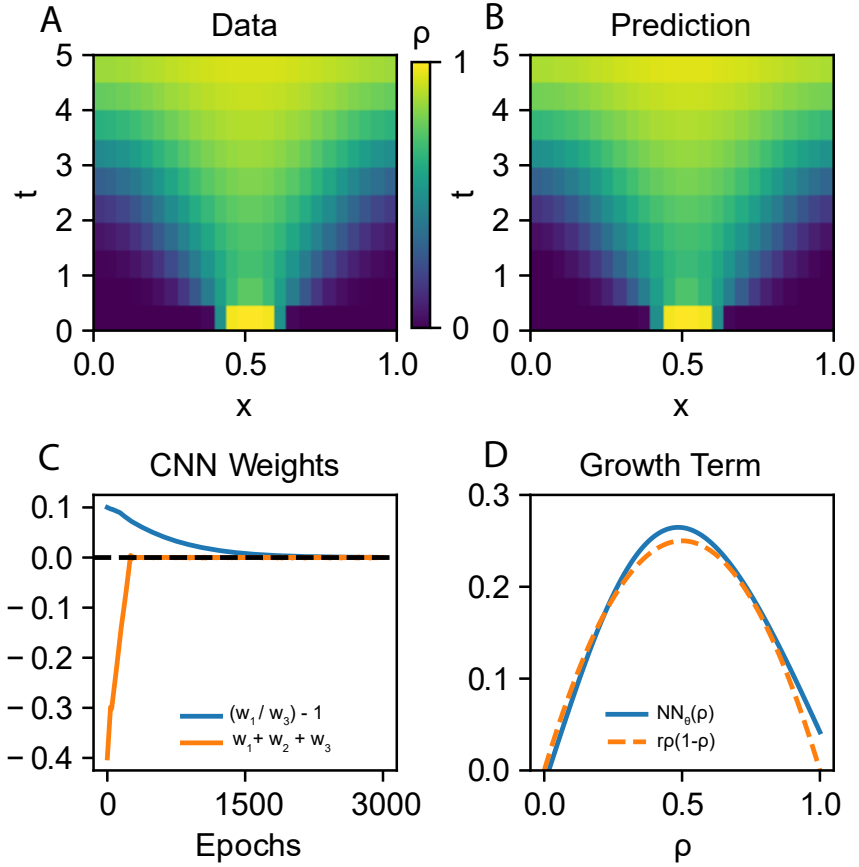


Figure 2: Recovery of the UPDE for the Fisher-KPP equation. (A) Training data and (B) prediction of the UPDE for $\rho(x, t)$. (C) Curves for the weights of the CNN filter $[w_1, w_2, w_3]$ indicate the recovery of the $[1, -2, 1]$ stencil for the 1-dimensional Laplacian. (D) Comparison of the learned (blue) and the true growth term (orange) showcases the learned parabolic form of the missing nonlinear equation.

355 Here we will showcase how one of these methods, a deep BSDE method for semi-
 356 linear parabolic equations [94], can be reinterpreted as a universal stochastic dif-
 357 ferential equation (USDE) to generalize the method and allow for enhancements
 358 like adaptivity, higher order integration for increased efficiency, and handling of
 359 stiff driving equations through the SciML software.

360 Consider the class of semilinear parabolic PDEs with a finite time span
 361 $t \in [0, T]$ and d -dimensional space $x \in \mathbb{R}^d$ that have the form:

$$\begin{aligned} \frac{\partial u}{\partial t}(t, x) + \frac{1}{2} \text{Tr}(\sigma \sigma^T(t, x) (\text{Hess}_x u)(t, x)) \\ + \nabla u(t, x) \cdot \mu(t, x) \\ + f(t, x, u(t, x), \sigma^T(t, x) \nabla u(t, x)) = 0, \end{aligned} \quad (17)$$

362 with a terminal condition $u(T, x) = g(x)$. Supplement S5 describes how this
 363 PDE can be solved by approximating by approximating the FBSDE:

$$\begin{aligned} dX_t &= \mu(t, X_t)dt + \sigma(t, X_t)dW_t, \\ dU_t &= f(t, X_t, U_t, U_{\theta_1}^1(t, X_t))dt + [U_{\theta_1}^1(t, X_t)]^T dW_t, \end{aligned} \quad (18)$$

364 where $U_{\theta_1}^1$ and $U_{\theta_2}^2$ are UAs and the loss function is given by the requiring that
 365 the terminating condition $g(X_T) = u(X_T, W_T)$ is satisfied.

366 3.1 Adaptive Solution of High-Dimensional Hamilton-Jacobi- 367 Bellman Equations

368 A fixed time step Euler-Maryumana discretization of this USDE gives rise to
 369 the deep BSDE method [94]. However, this form as a USDE generalizes the
 370 approach in a way that makes all of the methodologies of our USDE training
 371 library readily available, such as higher order methods, adaptivity, and implicit
 372 methods for stiff SDEs. As a motivating example, consider the classical linear-
 373 quadratic Gaussian (LQG) control problem in 100 dimensions:

$$dX_t = 2\sqrt{\lambda}c_t dt + \sqrt{2}dW_t, \quad (19)$$

374 with $t \in [0, T]$, $X_0 = x$, and with a cost function $C(c_t) = \mathbb{E} \left[\int_0^T \|c_t\|^2 dt + g(X_t) \right]$
 375 where X_t is the state we wish to control, λ is the strength of the control, and
 376 c_t is the control process. Minimizing the control corresponds to solving the
 377 100-dimensional HJB equation:

$$\frac{\partial u}{\partial t} + \nabla^2 u - \lambda \|\nabla u\|^2 = 0 \quad (20)$$

378 We solve the PDE by training the USDE using an adaptive Euler-Maruyama
 379 method [97] as described in Supplement S5. Supplementary Figure 2 showcases
 380 that this methodology accurately solves the equations, effectively extending re-
 381 cent algorithmic advancements to adaptive forms simply by reinterpreting the

382 equation as a USDE. While classical methods would require an amount of mem-
 383 ory that is exponential in the number of dimensions making classical adaptively
 384 approaches infeasible, this approach is the first the authors are aware of to
 385 generalize the high order, adaptive, highly stable software tooling to the high-
 386 dimensional PDE setting.

387 4 Accelerated Scientific Simulation with Auto- 388 matically Constructed Closure Relations

389 4.1 Automated Discovery of Large-Eddy Model Parame- 390 terizations

391 As an example of directly accelerating existing scientific workflows, we focus
 392 on the Boussinesq equations [98]. The Boussinesq equations are a system of
 393 3+1-dimensional partial differential equations acquired through simplifying as-
 394 sumptions on the incompressible Navier-Stokes equations, represented by the
 395 system:

$$\begin{aligned} \nabla \cdot \mathbf{u} &= 0, \\ \frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} &= -\nabla p + \nu \nabla^2 \mathbf{u} + b \hat{z}, \\ \frac{\partial T}{\partial t} + \mathbf{u} \cdot \nabla T &= \kappa \nabla^2 T, \end{aligned} \quad (21)$$

396 where $\mathbf{u} = (u, v, w)$ is the fluid velocity, p is the kinematic pressure, ν is the
 397 kinematic viscosity, κ is the thermal diffusivity, T is the temperature, and b is
 398 the fluid buoyancy. We assume that density and temperature are related by a
 399 linear equation of state so that the buoyancy b is only a function $b = \alpha g T$ where
 400 α is the thermal expansion coefficient and g is the acceleration due to gravity.

401 This system is commonly used in climate modeling, especially as the vox-
 402 els for modeling the ocean [99, 100, 101, 98] in a multi-scale model that ap-
 403 proximates these equations by averaging out the horizontal dynamics $\bar{T}(z, t) =$
 404 $\iint T(x, y, z, t) dx dy$ in individual boxes. The resulting approximation is a lo-
 405 cal advection-diffusion equation describing the evolution of the horizontally-
 406 averaged temperature \bar{T} :

$$\frac{\partial \bar{T}}{\partial t} + \frac{\partial \overline{wT}}{\partial z} = \kappa \frac{\partial^2 \bar{T}}{\partial z^2}. \quad (22)$$

407 This one-dimensional approximating system is not closed since \overline{wT} is unknown.
 408 Common practice closes the system by manually determining an approximating
 409 \overline{wT} from ad-hoc models, physical reasoning, and scaling laws. However, we can
 410 utilize a UDE-automated approach to learn such an approximation from data.
 411 Let

$$\overline{wT} = U_\theta \left(P, \bar{T}, \frac{\partial \bar{T}}{\partial z} \right) \quad (23)$$

412 where P are the physical parameters of the Boussinesq equation at different
413 regimes of the ocean, such as the amount of surface heating or the strength
414 of the surface winds [102]. We can accurately capture the non-locality of the
415 convection in this term by making the UDE a high-dimensional neural network.
416 Using data from horizontal average temperatures \bar{T} with known physical param-
417 eters P , we can directly reconstruct a nonlinear P -dependent parameterization
418 by training a universal diffusion-advection partial differential equation. Sup-
419plementary Figure 3 demonstrates the accuracy of the approach using a deep
420 UPDE with high order stabilized-explicit Runge-Kutta (ROCK) methods where
421 the fitting is described in Supplement S6. To contrast the trained UPDE, we di-
422rectly simulated the 3D Boussinesq equations under similar physical conditions
423 and demonstrated that the neural parameterization results in around a 15,000x
424 acceleration. This demonstrates that physical-dependent parameterizations for
425 acceleration can be directly learned from data utilizing the previous knowl-
426 edge of the averaging approximation and mixed with a data-driven discovery
427 approach.

428 4.2 Data-Driven Nonlinear Closure Relations for Model 429 Reduction in Non-Newtonian Viscoelastic Fluids

430 All continuum materials satisfy conservation equations for mass and momentum.
431 The difference between an elastic solid and a viscous fluid comes down to the
432 constitutive law relating the stresses and strains. In a one-dimensional system,
433 an elastic solid satisfies $\sigma = G\gamma$, with stress σ , strain γ , and elastic modulus
434 G , whereas a viscous fluid satisfies $\sigma = \eta\dot{\gamma}$, with viscosity η and strain rate $\dot{\gamma}$.
435 Non-Newtonian fluids have more complex constitutive laws, for instance when
436 stress depends on the history of deformation,

$$\sigma(t) = \int_{-\infty}^t G(t-s)F(\dot{\gamma}(s)) ds, \quad (24)$$

437 alternatively expressed in the instantaneous form [103]:

$$\begin{aligned} \sigma(t) &= \phi_1(t), \\ \frac{d\phi_1}{dt} &= G(0)F(\dot{\gamma}) + \phi_2, \\ \frac{d\phi_2}{dt} &= \frac{dG(0)}{dt}F(\dot{\gamma}) + \phi_3, \\ &\vdots \end{aligned} \quad (25)$$

438 where the history is stored in ϕ_i . To become computationally feasible, the
439 expansion is truncated, often in an ad-hoc manner, e.g. $\phi_n = \phi_{n+1} = \dots = 0$,
440 for some n . Only with a simple choice of $G(t)$ does an exact closure condition
441 exist, e.g. the Oldroyd-B model. For a fully nonlinear approximation, we train
442 a UODE according to the details in Supplement S7 to learn a closure relation:

$$\sigma(t) = U_0(\dot{\gamma}, \phi_1, \dots, \phi_N), \quad (26)$$

$$\frac{d\phi_i}{dt} = U_i(\dot{\gamma}, \phi_1, \dots, \phi_N), \quad \text{for } i = 1 \text{ to } N \quad (27)$$

443 from the numerical solution of the FENE-P equations, a fully non-linear consti-
 444 tutive law requiring a truncation condition [104]. Figure 3 compares the neural
 445 network approach to a linear, Oldroyd-B like, model for σ and showcases that
 446 the nonlinear approximation improves the accuracy by more than 50x. We
 447 note that the neural network approximation accelerates the solution by 2x over
 448 the original 6-state DAE, demonstrating that the universal differential equation
 449 approach to model acceleration is not just applicable to large-scale dynamical
 450 systems like PDEs but also can be effectively employed to accelerate small scale
 451 systems.

4.3 Efficient Discrete Physics-Informed Neural Networks as Universal ODEs

452 To further demonstrate the breadth of computational problems covered by the
 453 UODE framework, we note that the discrete physics-informed neural networks
 can be cast into the framework of UODEs. A physics-informed neural network
 is the representation of a PDE’s solution via a neural network, allowing machine
 learning training techniques to solve the equation [12]. These works note that
 the continuous PDE can be discretized in a single dimension to give rise to the
 discrete physics-informed neural network, simplified as:

$$u^{n+c_i} = u^n - \Delta t \sum_{j=1}^q a_{ij} \mathcal{N}[u^{n+c_j}] \quad (28)$$

$$u^{n+1} = u^n - \Delta t \sum_{j=1}^q b_j \mathcal{N}[u^{n+c_j}] \quad (29)$$

454 These results have demonstrated that the discrete form can enhance the com-
 455 putational efficiency of training physics-informed neural networks. However, we
 456 note that this directly corresponds to training the universal ODE $u' = \mathcal{N}(u)$ us-
 457 ing an explicit or implicit Runge-Kutta method in the SciML ecosystem. This
 458 directly gives rise to the further work on multistep discrete physics-informed
 459 neural networks [70, 80] by training the UODE via a multistep method, but
 460 also immediately gives the generalization to Runge-Kutta-Chebyshev, Rosen-
 461 brock, exponential integrator, and more formalizations which all are available
 462 via the SciML tools.

5 Discussion

463 While many attribute the success of deep learning to its blackbox nature, the
 464 key advances in deep learning applications have come by developing new archi-
 465 tectures which directly model the structures that are attempting to be learned.
 466

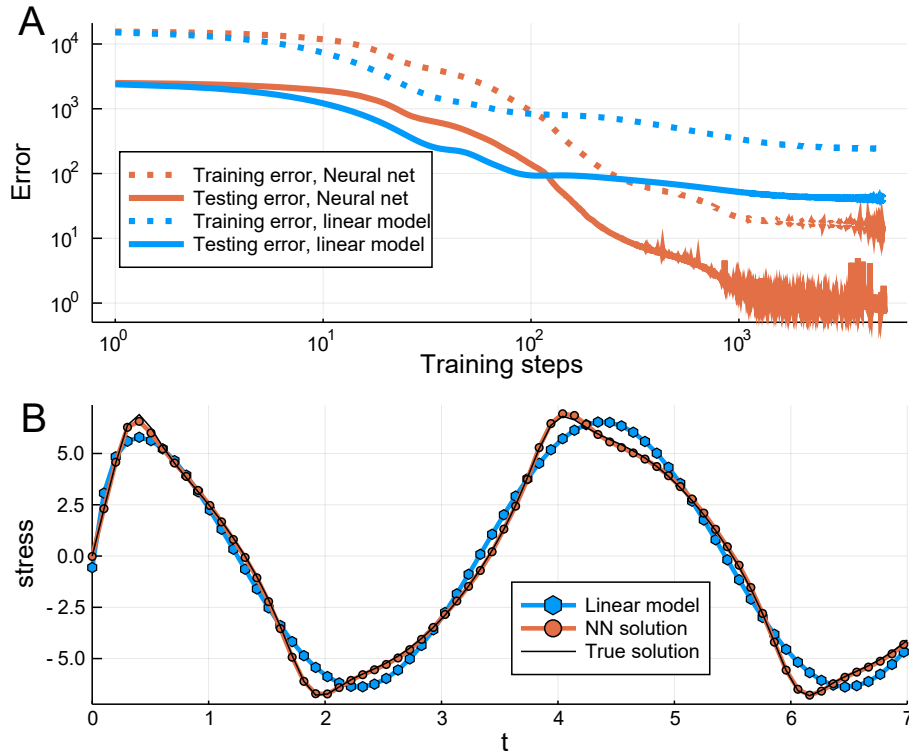


Figure 3: Convergence of neural closure relations for a non-Newtonian Fluid. (A) Error between the approximated σ using the linear approximation Equation 7 and the neural network closure relation Equation 26 against the full FENE-P solution. The error is measured for the strain rates $\dot{\gamma} = 12 \cos \omega t$ for $\omega = 1, 1.2, \dots, 2$ and tested with the strain rate $\dot{\gamma} = 12 \cos 1.5t$. (B) Predictions of stress for testing strain rate for the linear approximation and UODE solution against the exact FENE-P stress.

467 For example, deep convolutional neural networks for image processing directly
468 utilized the local spatial structure of images by modeling convolution stencil op-
469 erations. Similarly, recurrent neural networks encode a forward time progression
470 into a deep learning model and have excelled in natural language processing and
471 time series prediction. Here we present a software that allows for combining ex-
472 isting scientific simulation libraries with neural networks to train and augment
473 known models with data-driven components. Our results show that by build-
474 ing these hybrid mechanistic models with machine learning, we can arrive at
475 similar efficiency advancements by utilizing all known prior knowledge of the
476 underlying problem’s structure. While we demonstrate the utility of UDEs in
477 equation discovery, we have also demonstrated that these methods are capable
478 of solving many other problems, and many methods of recent interest, such as
479 discrete physics-informed neural networks, fall into the class of UDE methods
480 and can thus be analyzed and efficiently computed as part of this formalization.

481 Our software implementation is the first deep learning integrated differen-
482 tial equation library to include the full spectrum of adjoint sensitivity analysis
483 methods that is required to both efficiently and accurately handle the range
484 of training problems that can arise from universal differential equations. We
485 have demonstrated orders of magnitude performance advantages over previous
486 machine learning enhanced adjoint sensitivity ODE software in a variety of sci-
487 entific models and demonstrated generalizations to stiff equations, DAEs, SDEs,
488 and more. While the results of this paper span many scientific disciplines and
489 incorporate many different modeling approaches, together all of the examples
490 shown in this manuscript can be implemented using the SciML software ecosys-
491 tem in just hundreds of lines of code each, with none of the examples taking
492 more than half an hour to train on a standard laptop. This both demonstrates
493 the efficiency of the software and its methodologies, along with the potential to
494 scale to much larger applications.

495 The code for reproducing the computational experiments can be found at:

496 https://github.com/ChrisRackauckas/universal_differential_equations

497 **6 Acknowledgements**

498 We thank Jesse Bettencourt, Mike Innes, and Lyndon White for being instru-
499 mental in the early development of the DiffEqFlux.jl library, Tim Besard and
500 Valentin Churavy for the help with the GPU tooling, and David Widmann and
501 Kanav Gupta for their fundamental work across DifferentialEquations.jl. Spe-
502 cial thanks to Viral Shah and Steven Johnson who have been helpful in the
503 refinement of these ideas. We thank Charlie Strauss, Steven Johnson, Nathan
504 Urban, and Adam Gerlach for enlightening discussions and remarks on our
505 manuscript and software. We thank Stuart Rogers for his careful read and cor-
506 rections. We thank David Duvenaud for extended discussions on this work. We
507 thank the author of the torchsde library, Xuechen Li, for optimizing the SDE
508 benchmark code.

509 References

- 510 [1] Boukaye Boubacar Traore, Bernard Kamsu-Foguem, and Fana Tangara.
511 Deep convolution neural network for image recognition. *Ecological informatics*, 48:257–268, 2018.
512
- 513 [2] M. T. Islam, B. M. N. Karim Siddique, S. Rahman, and T. Jabid. Image
514 recognition with deep learning. In *2018 International Conference on Intelligent Informatics and Biomedical Sciences (ICIIBMS)*, volume 3, pages
515 106–110, Oct 2018.
516
- 517 [3] Chaofan Chen, Oscar Li, Daniel Tao, Alina Barnett, Cynthia Rudin, and
518 Jonathan K Su. This looks like that: deep learning for interpretable image
519 recognition. In *Advances in Neural Information Processing Systems*, pages
520 8928–8939, 2019.
- 521 [4] Tom Young, Devamanyu Hazarika, Soujanya Poria, and Erik Cambria.
522 Recent trends in deep learning based natural language processing. *IEEE Computational Intelligence Magazine*, 13(3):55–75, 2018.
523
- 524 [5] Daniel W Otter, Julian R Medina, and Jugal K Kalita. A survey of the
525 usages of deep learning in natural language processing. *arXiv preprint*
526 *arXiv:1807.10854*, 2018.
- 527 [6] Y Tsuruoka. Deep learning and natural language processing. *Brain and*
528 *nerve= Shinkei kenkyu no shinpo*, 71(1):45, 2019.
- 529 [7] Yu Li, Chao Huang, Lizhong Ding, Zhongxiao Li, Yijie Pan, and Xin Gao.
530 Deep learning in bioinformatics: Introduction, application, and perspective
531 in the big data era. *Methods*, 2019.
- 532 [8] Binhua Tang, Zixiang Pan, Kang Yin, and Asif Khateeb. Recent advances
533 of deep learning in bioinformatics and computational biology. *Frontiers*
534 *in Genetics*, 10, 2019.
- 535 [9] James Zou, Mikael Huss, Abubakar Abid, Pejman Mohammadi, Ali Torkamani,
536 and Amalio Telenti. A primer on deep learning in genomics. *Nature*
537 *genetics*, 51(1):12–18, 2019.
- 538 [10] Christof Angermueller, Tanel Pärnamaa, Leopold Parts, and Oliver Stegle.
539 Deep learning for computational biology. *Molecular systems biology*, 12(7),
540 2016.
- 541 [11] Davide Bacciu, Paulo JG Lisboa, José D Martín, Ruxandra Stoean, and
542 Alfredo Vellido. Bioinformatics and medicine in the era of deep learning.
543 *arXiv preprint arXiv:1802.09791*, 2018.
- 544 [12] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-
545 informed neural networks: A deep learning framework for solving forward
546 and inverse problems involving nonlinear partial differential equations.
547 *Journal of Computational Physics*, 378:686–707, 2019.

- 548 [13] Mauricio Alvarez, David Luengo, and Neil D Lawrence. Latent force
549 models. In *Artificial Intelligence and Statistics*, pages 9–16, 2009.
- 550 [14] Yueqin Hu, Steve Boker, Michael Neale, and Kelly L Klump. Coupled
551 latent differential equation with moderators: Simulation and application.
552 *Psychological Methods*, 19(1):56, 2014.
- 553 [15] Mauricio Alvarez, Jan R Peters, Neil D Lawrence, and Bernhard
554 Schölkopf. Switched latent force models for movement segmentation. In
555 *Advances in neural information processing systems*, pages 55–63, 2010.
- 556 [16] Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David K Du-
557 venaud. Neural ordinary differential equations. In *Advances in neural*
558 *information processing systems*, pages 6571–6583, 2018.
- 559 [17] Hongzhou Lin and Stefanie Jegelka. Resnet with one-neuron hidden layers
560 is a universal approximator. In *Advances in Neural Information Processing*
561 *Systems*, pages 6169–6178, 2018.
- 562 [18] David A Winkler and Tu C Le. Performance of deep and shallow neural
563 networks, the universal approximation theorem, activity cliffs, and qsar.
564 *Molecular informatics*, 36(1-2):1600118, 2017.
- 565 [19] Alexander N Gorban and Donald C Wunsch. The general approximation
566 theorem. In *1998 IEEE International Joint Conference on Neural Net-*
567 *works Proceedings. IEEE World Congress on Computational Intelligence*
568 *(Cat. No. 98CH36227)*, volume 2, pages 1271–1274. IEEE, 1998.
- 569 [20] M.R. Arahall and E.F. Camacho. Neural network adaptive control of non-
570 linear plants. *IFAC Proceedings Volumes*, 28(13):239 – 244, 1995. 5th
571 IFAC Symposium on Adaptive Systems in Control and Signal Processing
572 1995, Budapest, Hungary, 14-16 June, 1995.
- 573 [21] Wannes De Groote, Edward Kikken, Erik Hostens, Sofie Van Hoecke, and
574 Guillaume Crevecoeur. Neural network augmented physics models for
575 systems with partially unknown dynamics: Application to slider-crank
576 mechanism. *arXiv preprint arXiv:1910.12212*, 2019.
- 577 [22] Diederik P Kingma and Jimmy Ba. Adam A method for stochastic opti-
578 mization. *arXiv preprint arXiv:1412.6980*, 2014.
- 579 [23] Dong C Liu and Jorge Nocedal. On the limited memory bfgs method
580 for large scale optimization. *Mathematical programming*, 45(1-3):503–528,
581 1989.
- 582 [24] Ronald M Errico. What is an adjoint model? *Bulletin of the American*
583 *Meteorological Society*, 78(11):2577–2592, 1997.
- 584 [25] Grégoire Allaire. A review of adjoint methods for sensitivity analysis, un-
585 certainty quantification and optimization in numerical codes. *Ingenieurs*
586 *de l’Automobile*, 836:33–36, July 2015.

- 587 [26] Gilbert Strang. *Computational science and engineering*, volume 791.
588 Wellesley-Cambridge Press Wellesley, 2007.
- 589 [27] Alan C Hindmarsh, Peter N Brown, Keith E Grant, Steven L Lee, Radu
590 Serban, Dan E Shumaker, and Carol S Woodward. SUNDIALS: Suite of
591 nonlinear and differential/algebraic equation solvers. *ACM Transactions*
592 *on Mathematical Software (TOMS)*, 31(3):363–396, 2005.
- 593 [28] Steven G Johnson. Notes on adjoint methods for 18.335.
- 594 [29] Biswa Sengupta, Karl J Friston, and William D Penny. Efficient gradient
595 computation for dynamical models. *NeuroImage*, 98:521–527, 2014.
- 596 [30] Christopher Rackauckas, Yingbo Ma, Vaibhav Dixit, Xingjian Guo, Mike
597 Innes, Jarrett Revels, Joakim Nyberg, and Vijay Ivaturi. A comparison
598 of automatic differentiation and continuous sensitivity analysis for deriva-
599 tives of differential equation solutions. *arXiv preprint arXiv:1812.01892*,
600 2018.
- 601 [31] Michael Innes, Elliot Saba, Keno Fischer, Dhairya Gandhi, Marco Con-
602 cetto Rudilosso, Neethu Mariya Joy, Tejan Karmali, Avik Pal, and Viral
603 Shah. Fashionable modelling with flux. *CoRR*, abs/1811.01457, 2018.
- 604 [32] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury,
605 Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca
606 Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito,
607 Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner,
608 Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style,
609 high-performance deep learning library. In H. Wallach, H. Larochelle,
610 A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances*
611 *in Neural Information Processing Systems 32*, pages 8024–8035. Curran
612 Associates, Inc., 2019.
- 613 [33] Mike Innes, Alan Edelman, Keno Fischer, Chris Rackauckus, Elliot Saba,
614 Viral B Shah, and Will Tebbutt. Zygote: A differentiable programming
615 system to bridge machine learning and scientific computing. *arXiv preprint*
616 *arXiv:1907.07587*, 2019.
- 617 [34] Ralf Giering, Thomas Kaminski, and Thomas Slawig. Generating efficient
618 derivative code with taf: adjoint and tangent linear euler flow around an
619 airfoil. *Future generation computer systems*, 21(8):1345–1355, 2005.
- 620 [35] Laurent Hascoet and Valérie Pascual. The tapenade automatic differen-
621 tiation tool: Principles, model, and specification. *ACM Transactions on*
622 *Mathematical Software (TOMS)*, 39(3):20, 2013.
- 623 [36] Christopher Rackauckas and Qing Nie. Differentialequations.jl – a per-
624 formant and feature-rich ecosystem for solving differential equations in
625 julia. *The Journal of Open Research Software*, 5(1), 2017. Exported from
626 <https://app.dimensions.ai> on 2019/05/05.

- 627 [37] Tapio Schneider, Shiwei Lan, Andrew Stuart, and Joao Teixeira. Earth
628 system modeling 2.0: A blueprint for models that learn from observations
629 and targeted high-resolution simulations. *Geophysical Research Letters*,
630 44(24):12–396, 2017.
- 631 [38] Sebastian Krämer, David Plankensteiner, Laurin Ostermann, and Helmut
632 Ritsch. Quantumoptics.jl: A julia framework for simulating open quantum
633 systems. *Computer Physics Communications*, 227:109 – 116, 2018.
- 634 [39] Amir Gholami, Kurt Keutzer, and George Biros. Anode: Uncondition-
635 ally accurate memory-efficient gradients for neural odes. *arXiv preprint*
636 *arXiv:1902.10298*, 2019.
- 637 [40] Hong Zhang, Shrirang Abhyankar, Emil Constantinescu, and Mihai An-
638 itescu. Discrete adjoint sensitivity analysis of hybrid dynamical systems
639 with switching. *IEEE Transactions on Circuits and Systems I: Regular*
640 *Papers*, 64(5):1247–1259, 2017.
- 641 [41] Thomas Lauß, Stefan Oberpeilsteiner, Wolfgang Steiner, and Karin Nach-
642 bagauer. The discrete adjoint method for parameter identification in
643 multibody system dynamics. *Multibody system dynamics*, 42(4):397–410,
644 2018.
- 645 [42] J. Revels, M. Lubin, and T. Papamarkou. Forward-mode automatic dif-
646 ferentiation in julia. *arXiv:1607.07892 [cs.MS]*, 2016.
- 647 [43] Derek Onken and Lars Ruthotto. Discretize-optimize vs. optimize-
648 discretize for time-series regression and continuous normalizing flows.
649 *arXiv preprint arXiv:2005.13420*, 2020.
- 650 [44] Feby Abraham, Marek Behr, and Matthias Heinkenschloss. The effect of
651 stabilization in finite element methods for the optimal boundary control
652 of the oseen equations. *Finite Elements in Analysis and Design*, 41(3):229
653 – 251, 2004.
- 654 [45] John T Betts and Stephen L Campbell. Discretize then optimize. *Math-*
655 *ematics for industry: challenges and frontiers*, pages 140–157, 2005.
- 656 [46] Geng Liu, Martin Geier, Zhenyu Liu, Manfred Krafczyk, and Tao Chen.
657 Discrete adjoint sensitivity analysis for fluid flow topology optimization
658 based on the generalized lattice boltzmann method. *Computers & Math-*
659 *ematics with Applications*, 68(10):1374 – 1392, 2014.
- 660 [47] Alfonso Callejo, Valentin Sonneville, and Olivier A Bauchau. Discrete
661 adjoint method for the sensitivity analysis of flexible multibody systems.
662 *Journal of Computational and Nonlinear Dynamics*, 14(2), 2019.
- 663 [48] S Scott Collis and Matthias Heinkenschloss. Analysis of the streamline
664 upwind/petrov galerkin method applied to the solution of optimal control
665 problems. 2002.

- 666 [49] Jun Liu and Zhu Wang. Non-commutative discretize-then-optimize algo-
667 rithms for elliptic pde-constrained optimal control problems. *Journal of*
668 *Computational and Applied Mathematics*, 362:596–613, 2019.
- 669 [50] E Huntley. A note on the application of the matrix riccati equation to the
670 optimal control of distributed parameter systems. *IEEE Transactions on*
671 *Automatic Control*, 24(3):487–489, 1979.
- 672 [51] Ziv Sirkes and Eli Tziperman. Finite difference of adjoint or adjoint of
673 finite difference? *Monthly weather review*, 125(12):3373–3378, 1997.
- 674 [52] Guojun Hu and Tomasz Kozlowski. Assessment of continuous and dis-
675 crete adjoint method for sensitivity analysis in two-phase flow simulations.
676 *arXiv preprint arXiv:1805.08083*, 2018.
- 677 [53] JOHANNES Kepler. Sensitivity analysis: The direct and adjoint method.
- 678 [54] F Van Keulen, RT Haftka, and NH Kim. Review of options for structural
679 design sensitivity analysis. part 1: Linear systems. *Computer methods in*
680 *applied mechanics and engineering*, 194(30-33):3213–3243, 2005.
- 681 [55] M Kouhi, G Houzeaux, F Cucchietti, M Vázquez, and F Rodriguez. Im-
682 plementation of discrete adjoint method for parameter sensitivity analysis
683 in chemically reacting flows.
- 684 [56] Siva Nadarajah and Antony Jameson. A comparison of the continuous
685 and discrete adjoint approach to automatic aerodynamic optimization. In
686 *38th Aerospace Sciences Meeting and Exhibit*, page 667.
- 687 [57] Tianyi Gou and Adrian Sandu. Continuous versus discrete advection ad-
688 joints in chemical data assimilation with cmaq. *Atmospheric environment*,
689 45(28):4868–4881, 2011.
- 690 [58] Nicolas R Gauger, Michael Giles, Max Gunzburger, and Uwe Nau-
691 mann. Adjoint methods in computational science, engineering, and fi-
692 nance (dagstuhl seminar 14371). In *Dagstuhl Reports*, volume 4. Schloss
693 Dagstuhl-Leibniz-Zentrum fuer Informatik, 2015.
- 694 [59] G. Hu and T. Kozlowski. Development and assessment of adjoint sen-
695 sitivity analysis method for transient two-phase flow simulations. pages
696 2246–2259, January 2019. 18th International Topical Meeting on Nuclear
697 Reactor Thermal Hydraulics, NURETH 2019 ; Conference date: 18-08-
698 2019 Through 23-08-2019.
- 699 [60] Dacian N. Daescu, Adrian Sandu, and Gregory R. Carmichael. Direct
700 and adjoint sensitivity analysis of chemical kinetic systems with kpp:
701 Ii—numerical validation and applications. *Atmospheric Environment*,
702 37(36):5097 – 5114, 2003.

- 703 [61] A Schwartz and E Polak. Runge-kutta discretization of optimal control
704 problems. *IFAC Proceedings Volumes*, 29(8):123–128, 1996.
- 705 [62] Kimia Ghobadi, Nedialko S Nedialkov, and Tamas Terlaky. On the dis-
706 cretize then optimize approach. *Preprint for Industrial and Systems En-*
707 *gineering*, 2009.
- 708 [63] Alain Sei and William W Symes. A note on consistency and adjointness
709 for numerical schemes. 1995.
- 710 [64] William W Hager. Runge-kutta methods in optimal control and the trans-
711 formed adjoint system. *Numerische Mathematik*, 87(2):247–282, 2000.
- 712 [65] Adrian Sandu, Dacian N Daescu, Gregory R Carmichael, and Tianfeng
713 Chai. Adjoint sensitivity analysis of regional air quality models. *Journal*
714 *of Computational Physics*, 204(1):222–252, 2005.
- 715 [66] Shashi Gowda, Yingbo Ma, Valentin Churavy, Alan Edelman, and
716 Christopher Rackauckas. Sparsity programming: Automated sparsity-
717 aware optimizations in differentiable programming. 2019.
- 718 [67] Xuechen Li, Ting-Kam Leonard Wong, Ricky T. Q. Chen, and David
719 Duvenaud. Scalable gradients for stochastic differential equations. *Inter-*
720 *national Conference on Artificial Intelligence and Statistics*, 2020.
- 721 [68] E. Hairer, S. P. Nørsett, and G. Wanner. *Solving Ordinary Differen-*
722 *tial Equations I (2nd Revised. Ed.): Nonstiff Problems*. Springer-Verlag,
723 Berlin, Heidelberg, 1993.
- 724 [69] Zichao Long, Yiping Lu, Xianzhong Ma, and Bin Dong. Pde-net: Learning
725 pdes from data. *arXiv preprint arXiv:1710.09668*, 2017.
- 726 [70] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Multistep
727 neural networks for data-driven discovery of nonlinear dynamical systems.
728 *arXiv preprint arXiv:1801.01236*, 2018.
- 729 [71] Peter J Schmid. Dynamic mode decomposition of numerical and experi-
730 mental data. *Journal of fluid mechanics*, 656:5–28, 2010.
- 731 [72] Matthew O Williams, Ioannis G Kevrekidis, and Clarence W Rowley.
732 A data-driven approximation of the koopman operator: Extending dy-
733 namic mode decomposition. *Journal of Nonlinear Science*, 25(6):1307–
734 1346, 2015.
- 735 [73] Qianxiao Li, Felix Dietrich, Erik M Bollt, and Ioannis G Kevrekidis. Ex-
736 tended dynamic mode decomposition with dictionary learning: A data-
737 driven adaptive spectral decomposition of the koopman operator. *Chaos:*
738 *An Interdisciplinary Journal of Nonlinear Science*, 27(10):103111, 2017.

- 739 [74] Naoya Takeishi, Yoshinobu Kawahara, and Takehisa Yairi. Learning koop-
740 man invariant subspaces for dynamic mode decomposition. In *Advances*
741 *in Neural Information Processing Systems*, pages 1130–1140, 2017.
- 742 [75] Hayden Schaeffer. Learning partial differential equations via data discov-
743 ery and sparse optimization. *Proceedings of the Royal Society A: Mathe-*
744 *matical, Physical and Engineering Sciences*, 473(2197):20160446, 2017.
- 745 [76] Markus Quade, Markus Abel, Kamran Shafi, Robert K Niven, and
746 Bernd R Noack. Prediction of dynamical systems by symbolic regression.
747 *Physical Review E*, 94(1):012214, 2016.
- 748 [77] Hongqing Cao, Lishan Kang, Yuping Chen, and Jingxian Yu. Evolution-
749 ary modeling of systems of ordinary differential equations with genetic
750 programming. *Genetic Programming and Evolvable Machines*, 1(4):309–
751 337, 2000.
- 752 [78] Khalid Raza and Rafat Parveen. Evolutionary algorithms in genetic reg-
753 ulatory networks model. *CoRR*, abs/1205.1986, 2012.
- 754 [79] Hayden Schaeffer, Giang Tran, and Rachel Ward. Extracting sparse high-
755 dimensional dynamics from limited data. *SIAM Journal on Applied Math-*
756 *ematics*, 78(6):3279–3295, 2018.
- 757 [80] Ramakrishna Tipireddy, Paris Perdikaris, Panos Stinis, and Alexandre M.
758 Tartakovsky. A comparative study of physics-informed neural network
759 models for learning unknown dynamics and constitutive relations. *CoRR*,
760 abs/1904.04058, 2019.
- 761 [81] Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Discovering gov-
762 erning equations from data by sparse identification of nonlinear dynamical
763 systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–
764 3937, 2016.
- 765 [82] Niall M Mangan, Steven L Brunton, Joshua L Proctor, and J Nathan
766 Kutz. Inferring biological networks by sparse identification of nonlinear
767 dynamics. *IEEE Transactions on Molecular, Biological and Multi-Scale*
768 *Communications*, 2(1):52–63, 2016.
- 769 [83] Niall M Mangan, J Nathan Kutz, Steven L Brunton, and Joshua L Proc-
770 tor. Model selection for dynamical systems via sparse regression and in-
771 formation criteria. *Proceedings of the Royal Society A: Mathematical,*
772 *Physical and Engineering Sciences*, 473(2204):20170009, 2017.
- 773 [84] Peng Zheng, Travis Askham, Steven L. Brunton, J. Nathan Kutz, and
774 Aleksandr Y. Aravkin. A unified framework for sparse relaxed regularized
775 regression: SR3. 7:1404–1423. Conference Name: IEEE Access.

- 776 [85] Kathleen Champion, Peng Zheng, Aleksandr Y. Aravkin, Steven L. Brunton,
777 and J. Nathan Kutz. A unified sparse optimization framework to
778 learn parsimonious physics-informed models from data.
- 779 [86] Alexander Bills, Shashank Sripad, William Leif Fredericks, Matthew
780 Guttenberg, Devin Charles, Evan Frank, and Venkatasubramanian
781 Viswanathan. Universal Battery Performance and Degradation Model
782 for Electric Aircraft. 7 2020.
- 783 [87] Raj Dandekar and Lydia Bourouiba. Splash upon impact on a deep pool.
784 In preparation.
- 785 [88] Yang Cao, Shengtai Li, Linda Petzold, and Radu Serban. Adjoint sensi-
786 tivity analysis for differential-algebraic equations: The adjoint dae sys-
787 tem and its numerical solution. *SIAM journal on scientific computing*,
788 24(3):1076–1089, 2003.
- 789 [89] R. A. Fisher. The wave of advance of advantageous genes. *Annals of*
790 *Eugenics*, 7(4):355–369, 1937.
- 791 [90] P. Grindrod. *The Theory and Applications of Reaction-diffusion Equa-*
792 *tions: Patterns and Waves*. Oxford applied mathematics and computing
793 science series. Clarendon Press, 1996.
- 794 [91] Justin Sirignano and Konstantinos Spiliopoulos. Dgm: A deep learning
795 algorithm for solving partial differential equations. *Journal of Computa-*
796 *tional Physics*, 375:1339–1364, Dec 2018.
- 797 [92] Isaac E Lagaris, Aristidis Likas, and Dimitrios I Fotiadis. Artificial neural
798 networks for solving ordinary and partial differential equations. *IEEE*
799 *transactions on neural networks*, 9(5):987–1000, 1998.
- 800 [93] E Weinan, Jiequn Han, and Arnulf Jentzen. Deep learning-based numeri-
801 cal methods for high-dimensional parabolic partial differential equations
802 and backward stochastic differential equations. *Communications in Math-*
803 *ematics and Statistics*, 5(4):349–380, 2017.
- 804 [94] Jiequn Han, Arnulf Jentzen, and Weinan E. Solving high-dimensional par-
805 tial differential equations using deep learning. *Proceedings of the National*
806 *Academy of Sciences*, 115(34):8505–8510, 2018.
- 807 [95] Yaohua Zang, Gang Bao, Xiaojing Ye, and Haomin Zhou. Weak adver-
808 sarial networks for high-dimensional partial differential equations. *arXiv*
809 *preprint arXiv:1907.08272*, 2019.
- 810 [96] Côme Huré, Huyên Pham, and Xavier Warin. Some machine learn-
811 ing schemes for high-dimensional nonlinear pdes. *arXiv preprint*
812 *arXiv:1902.01599*, 2019.

- 813 [97] H Lamba. An adaptive timestepping algorithm for stochastic differential
814 equations. *Journal of computational and applied mathematics*, 161(2):417–
815 430, 2003.
- 816 [98] Benoit Cushman-Roisin and Jean-Marie Beckers. Chapter 4 - equations
817 governing geophysical flows. In Benoit Cushman-Roisin and Jean-Marie
818 Beckers, editors, *Introduction to Geophysical Fluid Dynamics*, volume 101
819 of *International Geophysics*, pages 99 – 129. Academic Press, 2011.
- 820 [99] Zhihua Zhang and John C. Moore. Chapter 11 - atmospheric dynamics.
821 In Zhihua Zhang and John C. Moore, editors, *Mathematical and Physical*
822 *Fundamentals of Climate Change*, pages 347 – 405. Elsevier, Boston, 2015.
- 823 [100] Section 1.3 - governing equations. In Lakshmi H. Kantha and Carol Anne
824 Clayson, editors, *Numerical Models of Oceans and Oceanic Processes*, vol-
825 ume 66 of *International Geophysics*, pages 28–46. Academic Press, 2000.
- 826 [101] Stephen M Griffies and Alistair J Adcroft. Formulating the equations of
827 ocean models. 2008.
- 828 [102] Stephen M. Griffies, Michael Levy, Alistair J. Adcroft, Gokhan Danaba-
829 soglu, Robert W. Hallberg, Doug Jacobsen, William Large, , and Todd
830 Ringler. Theory and Numerics of the Community Ocean Vertical Mixing
831 (CVMix) Project. Technical report, 2015. Draft from March 9, 2015. 98
832 + v pages.
- 833 [103] F.A. Morrison and A.P.C.E.F.A. Morrison. *Understanding Rheology*. Ray-
834 mond F. Boyer Library Collection. Oxford University Press, 2001.
- 835 [104] P.J. Oliveira. Alternative derivation of differential constitutive equa-
836 tions of the oldroyd-b type. *Journal of Non-Newtonian Fluid Mechanics*,
837 160(1):40 – 46, 2009. Complex flows of complex fluids.

Figures

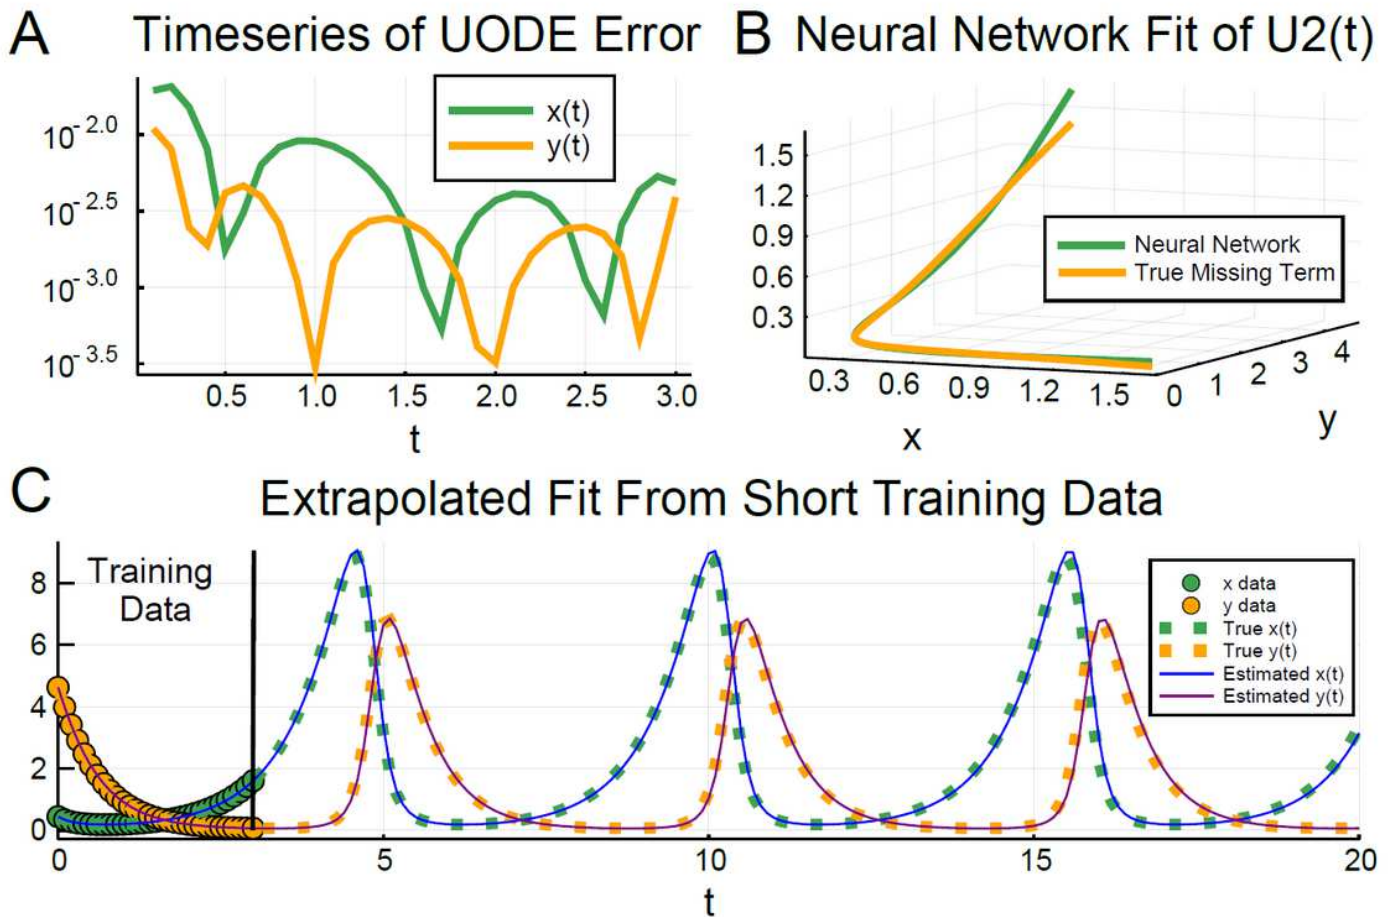


Figure 1

Automated Lotka-Volterra equation discovery with UODE-enhanced SInDy. (A) The error in the trained UODE against $x(t)$ and $y(t)$ in green and yellow respectively. (B) The measured values of the missing term $U_2(x; y)$ throughout the time series, with the neural network approximate in green and the true value xy in yellow. (C) The extrapolation of the knowledge-enhanced SInDy fit series. The green and yellow dots show the data that was used to fit the UODE, and the dots show the true solution of the Lotka-Volterra Equations 8 beyond the training data. The blue and purple lines show the extrapolated solution how the UODE-enhanced SInDy recovered equations.

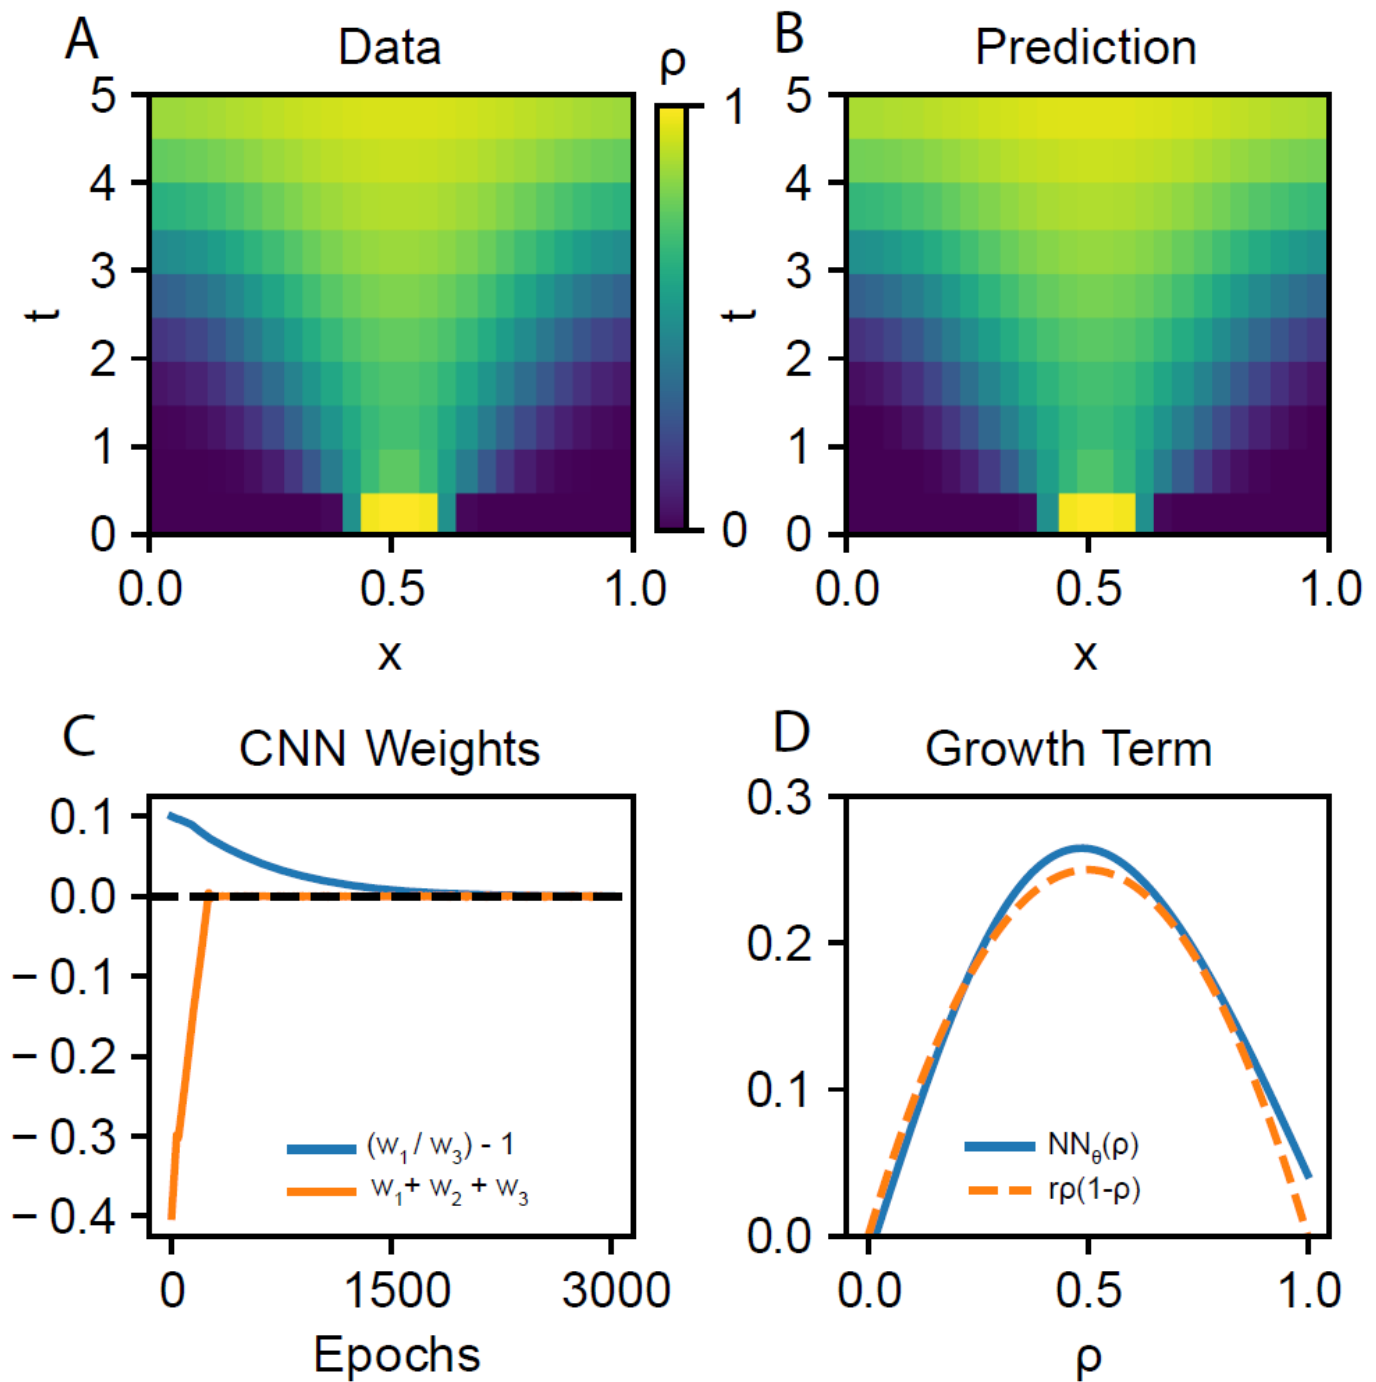


Figure 2

Recovery of the UPDE for the Fisher-KPP equation. (A) Training data and (B) prediction of the UPDE for $\rho(x, t)$. (C) Curves for the weights of the CNN filter $[w_1, w_2, w_3]$ indicate the recovery of the $[1, -2, 1]$ stencil for the 1-dimensional Laplacian. (D) Comparison of the learned (blue) and the true growth term (orange) showcases the learned parabolic form of the missing nonlinear equation.

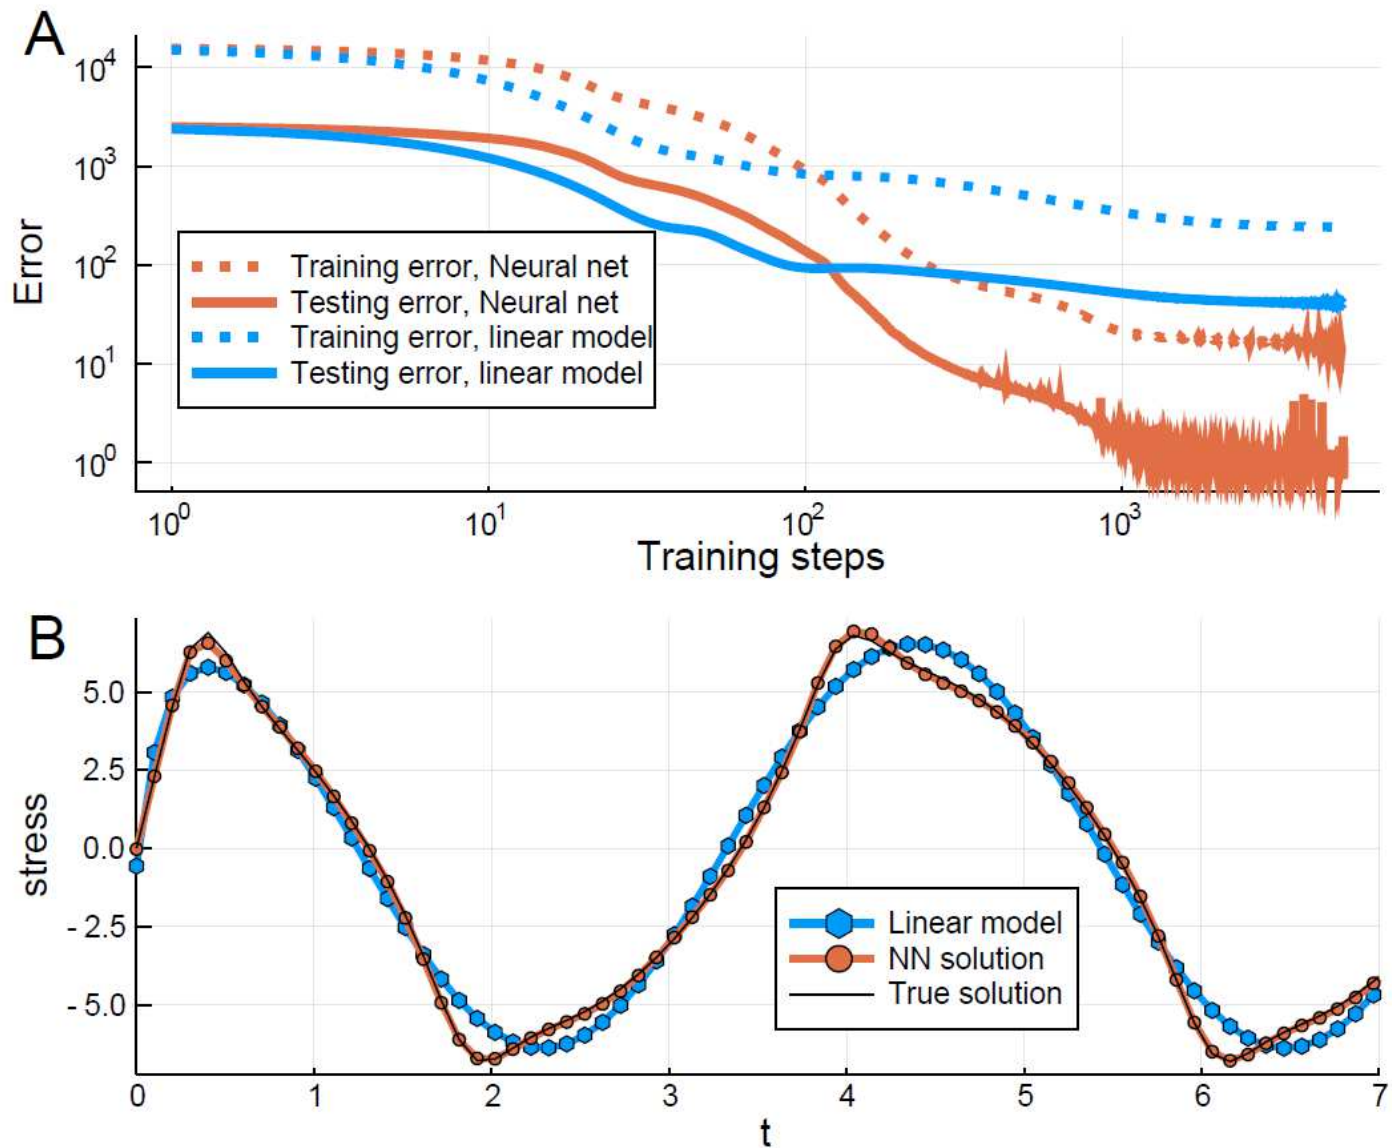


Figure 3

Convergence of neural closure relations for a non-Newtonian Fluid. (A) Error between the approximated σ using the linear approximation Equation 7 and the neural network closure relation Equation 26 against the full FENE-P solution. The error is measured for the strain rates $\dot{\gamma} = 12 \cos \omega t$ for $\omega = 1, 1.2, \dots, 2$ and tested with the strain rate $\dot{\gamma} = 12 \cos 1.5t$. (B) Predictions of stress for testing strain rate for the linear approximation and UODE solution against the exact FENE-P stress.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [UniversalDifferentialEquationsSupplement.pdf](#)