

Doing Scientific Machine Learning with Julia's SciML Ecosystem



CHRIS RACKAUCKAS

APPLIED MATHEMATICS INSTRUCTOR, MASSACHUSETTS INSTITUTE OF TECHNOLOGY, MATHEMATICS

SENIOR RESEARCH ANALYST, UNIVERSITY OF MARYLAND, BALTIMORE, SCHOOL OF PHARMACY

DIRECTOR OF SCIENTIFIC RESEARCH, PUMAS-AI

What to expect from this workshop

We will go through “what is scientific machine learning” all the way through some examples of how to do scientific machine learning in Julia:

- ▶ Parameter inference of differential equation model
- ▶ Add events to equations, add stochasticity, add delays
- ▶ Bayesian inference of scientific models

This tutorial is meant to be raw:

- ▶ We're going to be live coding
 - ▶ “We”, as in you and me!
- ▶ We are going to be learning the methods as we learn to write the code
- ▶ We are going to learn how to navigate the ecosystem, the documentation, the IDE, the error messages

How to Help!

- ▶ Star the repositories
 - ▶ DifferentialEquations.jl, DiffEqFlux.jl, DataDrivenDiffEq.jl, NeuralPDE.jl, ModelingToolkit.jl, Catalyst.jl
 - ▶ Flux.jl, Zygote.jl, SparseDiffTools.jl, etc.
- ▶ Report bugs
- ▶ Create tutorials
- ▶ Write blog posts
- ▶ Share data and challenge problems
- ▶ Help in chats and community channels
- ▶ Add your own packages to the common interface
- ▶ Contribute to the SciML packages



Studies in Pandemic Preparedness
Simon Frost



Pumas



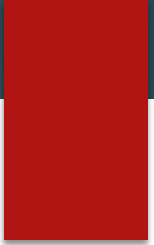
NUMFOCUS
OPEN CODE = BETTER SCIENCE

Julia
computing

We thank all of those who have previously helped SciML development

Workshop Outline

- ▶ Overview of scientific machine learning and SciML
 - ▶ What is scientific machine learning?
 - ▶ What makes the SciML ecosystem unique?
- ▶ Introduction to challenge and learning problems
 - ▶ Workshop exercises (with answers!)
 - ▶ HelicopterSciML Challenge Problem
 - ▶ **Magnetic Navigation Challenge Problem**
- ▶ Modeling with differential equations
 - ▶ Solving differential equations with DifferentialEquations.jl
 - ▶ Adding stochasticity, delays, events
- ▶ Automated model discovery via universal differential equations
 - ▶ Parameter inference on differential equations
 - ▶ Local and global optimization
 - ▶ Bayesian optimization
 - ▶ Mixing DiffEqFlux.jl and DataDrivenDiffEq.jl!
- ▶ Solving differential equations with neural networks (physics-informed neural networks)



Scientific machine learning is the mixture of scientific models with data-driven machine learning components for data-efficient model-based decision making

Universal Approximation Theorem

NEURAL NETWORKS CAN GET ϵ CLOSE TO ANY
 $R^n \rightarrow R^m$ FUNCTION

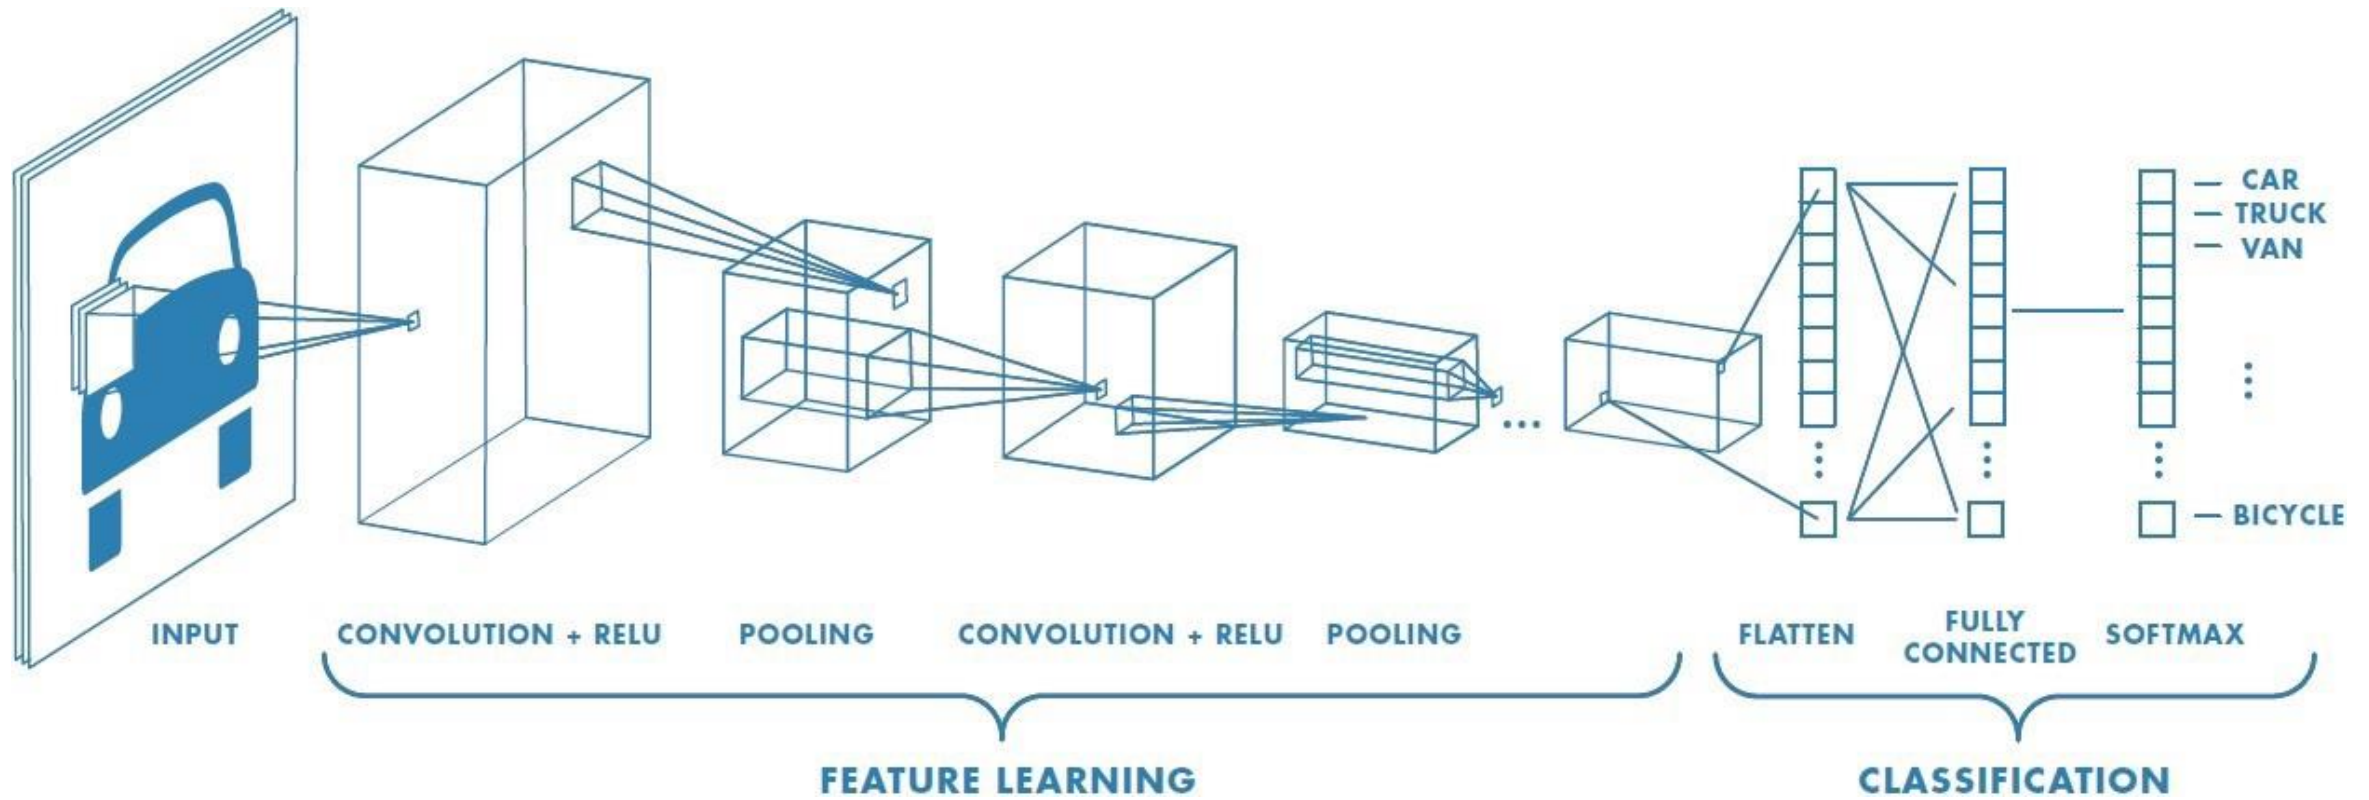
**Neural networks are just function expansions,
fancy Taylor Series like things which are good
for computing and bad for analysis.**

Neural networks work well in high dimensions

The major advances in machine learning were due to encoding more structure into the model

MORE STRUCTURE = FASTER AND BETTER FITS FROM LESS DATA

Convolutional Neural Networks Are Structure Assumptions



Universal Differential Equations:
Differential Equations defined in part
by universal approximators

Use all known scientific features, use
all numerical methods, have neural
networks cover the last mile

Demonstration of UDEs on a toy model

$$S' = -\frac{\beta_0 SF}{N} - \frac{\beta(t)SI}{N} - \mu S,$$

$$E' = \frac{\beta_0 SF}{N} + \frac{\beta(t)SI}{N} - (\sigma + \mu)E,$$

$$I' = \sigma E - (\gamma + \mu)I,$$

$$R' = \gamma I - \mu R,$$

$$N' = -\mu N,$$

$$D' = d \gamma I - \lambda D, \quad \text{and}$$

$$C' = \sigma E,$$

$$\beta(t) = \beta_0(1 - \alpha)\left(1 - \frac{D}{N}\right)^\kappa \quad \kappa = 1117.3$$

A conceptual model for the coronavirus disease 2019 (COVID-19) outbreak in Wuhan, China with individual reaction and governmental action

Lin, Qianying et al.

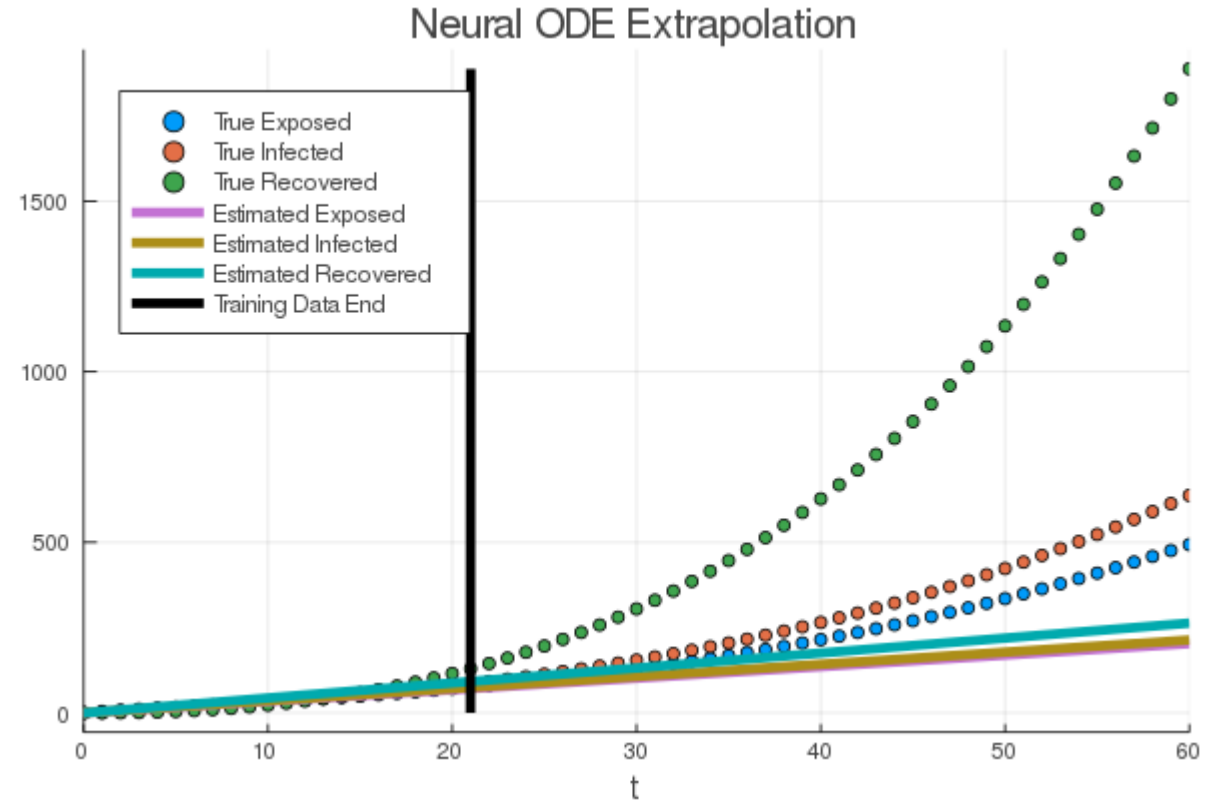
International Journal of Infectious Diseases, Volume 93, 211 - 216

Neural ODE: Learn the whole model

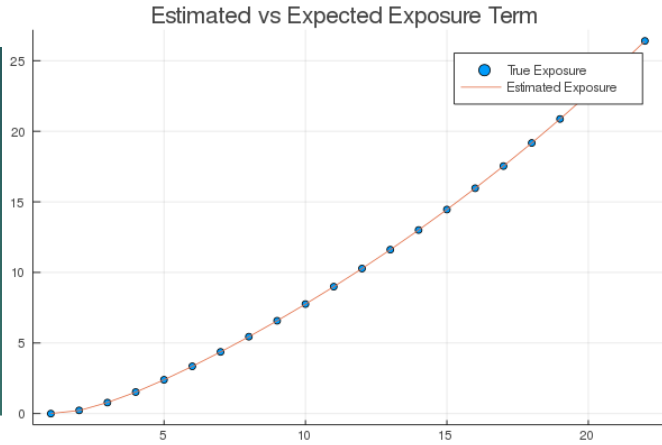
$$u' = NN(u)$$

Can fit, but not enough information to accurately extrapolate


Does not have the correct asymptotic behavior



Universal ODE




Neural ODE Extrapolation



$$\begin{aligned}
 S' &= -\frac{\beta_0 S F}{N} - \text{Replace Unknown Portion} - \mu S, & \text{Exposure: Unknown} \\
 E' &= \frac{\beta_0 S F}{N} + \text{Replace Unknown Portion} - (\sigma + \mu) E, \\
 I' &= \sigma E - (\gamma + \mu) I, \\
 R' &= \gamma I - \mu R, & \text{Infection rates: known From disease quantities} \\
 N' &= -\mu N, \\
 D' &= d \gamma I - \lambda D, & \text{and Percentage of cases known to be severe, can be estimated} \\
 C' &= \sigma E,
 \end{aligned}$$

Universal ODE Extrapolation



SInDy – Sparse Identification of Dynamical Systems

sparse vectors of coefficients $\Xi = [\xi_1 \xi_2 \dots \xi_n]$ that determine which nonlinearities are active:

$$\dot{\mathbf{X}} = \Theta(\mathbf{X})\Xi. \quad [3]$$

$$\mathbf{X} = \begin{matrix} & \text{state} \\ \begin{matrix} \mathbf{x}^T(t_1) \\ \mathbf{x}^T(t_2) \\ \vdots \\ \mathbf{x}^T(t_m) \end{matrix} & = & \begin{matrix} \begin{bmatrix} x_1(t_1) & x_2(t_1) & \dots & x_n(t_1) \\ x_1(t_2) & x_2(t_2) & \dots & x_n(t_2) \\ \vdots & \vdots & \ddots & \vdots \\ x_1(t_m) & x_2(t_m) & \dots & x_n(t_m) \end{bmatrix} \end{matrix} \downarrow \text{time} \\ \\ \dot{\mathbf{X}} = \begin{matrix} \begin{matrix} \dot{\mathbf{x}}^T(t_1) \\ \dot{\mathbf{x}}^T(t_2) \\ \vdots \\ \dot{\mathbf{x}}^T(t_m) \end{matrix} & = & \begin{matrix} \begin{bmatrix} \dot{x}_1(t_1) & \dot{x}_2(t_1) & \dots & \dot{x}_n(t_1) \\ \dot{x}_1(t_2) & \dot{x}_2(t_2) & \dots & \dot{x}_n(t_2) \\ \vdots & \vdots & \ddots & \vdots \\ \dot{x}_1(t_m) & \dot{x}_2(t_m) & \dots & \dot{x}_n(t_m) \end{bmatrix} \end{matrix} \end{matrix}$$

Next, we construct a library $\Theta(\mathbf{X})$ consisting of candidate nonlinear functions of the columns of \mathbf{X} . For example, $\Theta(\mathbf{X})$ may consist of constant, polynomial, and trigonometric terms:

$$\Theta(\mathbf{X}) = \begin{bmatrix} 1 & \mathbf{X} & \mathbf{X}^{P_2} & \mathbf{X}^{P_3} & \dots & \sin(\mathbf{X}) & \cos(\mathbf{X}) & \dots \end{bmatrix}. \quad [2]$$

Not Enough Data! Unable to achieve a sparse basis

Brunton, Steven L., Joshua L. Proctor, and J. Nathan Kutz. "Discovering governing equations from data by sparse identification of nonlinear dynamical systems." *Proceedings of the national academy of sciences* 113.15 (2016): 3932-3937.

- ▶ Operation $\cos(u_1) * -0.0013108600297508188 + \cos(u_2) * 0.001048733466930909 + \sin(u_3) * 0.002524237642240494 + 4.582000697122147 + u_3 * 48.22745315102507 + u_3 \wedge 2 * -0.5293305992835255 + u_2 * 39.085961651678964 + u_2 * u_3 * -0.6742175940650399 + u_2 * u_3 \wedge 2 * 0.0018086945606415868 + u_2 \wedge 2 * -0.7760315827702667 + u_2 \wedge 2 * u_3 * -0.00827007707292397 + u_2 \wedge 2 * u_3 \wedge 2 * -4.8420203054602525e-5 + u_1 * 0.6927075862062384 + u_1 * u_3 * 2.5477896384187675 + u_1 * u_3 \wedge 2 * -0.007633697801342265 + u_1 * u_2 * -0.8050223920175605 + u_1 * u_2 * u_3 * -0.005893734488035572 + u_1 * u_2 * u_3 \wedge 2 * -4.205818407350913e-5 + u_1 * u_2 \wedge 2 * 0.05154776022562611 + u_1 * u_2 \wedge 2 * u_3 * 0.00011401535262358879 + u_1 * u_2 \wedge 2 * u_3 \wedge 2 * -1.8409670007515867e-7 + u_1 \wedge 2 * -1.480917344589218 + u_1 \wedge 2 * u_3 * 0.022834435321810845 + u_1 \wedge 2 * u_3 \wedge 2 * -7.10505011605666e-5 + u_1 \wedge 2 * u_2 * -0.0811262292209696 + u_1 \wedge 2 * u_2 * u_3 * 1.2503710381374686e-5 + u_1 \wedge 2 * u_2 * u_3 \wedge 2 * -1.5835869421530206e-7 + u_1 \wedge 2 * u_2 \wedge 2 * 0.0003756078420420898 + u_1 \wedge 2 * u_2 \wedge 2 * u_3 * 2.0403671083190194e-6 + u_1 \wedge 2 * u_2 \wedge 2 * u_3 \wedge 2 * -4.0790059067580516e-10, \cos(u_1) * 0.0018236630124880049 + \sin(u_3) * -0.002857556410244201 + 0.738713743952307 + u_3 * -45.316633125282735 + u_3 \wedge 2 * 0.4976552341495027 + u_2 * -36.669905096040644 + u_2 * u_3 * 0.63405194300575 + u_2 * u_3 \wedge 2 * -0.001699189499009162 + u_2 \wedge 2 * 0.7292234161358288 + u_2 \wedge 2 * u_3 * 0.007782847250932861 + u_2 \wedge 2 * u_3 \wedge 2 * 4.5537832343115385e-5 + u_1 * -0.662837140886116 + u_1 * u_3 * -2.3955577736237044 + u_1 * u_3 \wedge 2 * 0.007174813124917316 + u_1 * u_2 * 0.7564652530371222 + u_1 * u_2 * u_3 * 0.005539740817006857 + u_1 * u_2 * u_3 \wedge 2 * 3.952859749575076e-5 + u_1 * u_2 \wedge 2 * -0.04846972496409705 + u_1 * u_2 \wedge 2 * u_3 * -0.00010714683124587004 + u_1 * u_2 \wedge 2 * u_3 \wedge 2 * 1.7315253185547634e-7 + u_1 \wedge 2 * 1.3922758705496125 + u_1 \wedge 2 * u_3 * -0.021478161074782457 + u_1 \wedge 2 * u_3 \wedge 2 * 6.675620535553527e-5 + u_1 \wedge 2 * u_2 * 0.07628907557295377 + u_1 \wedge 2 * u_2 * u_3 * -1.174623626431566e-5 + u_1 \wedge 2 * u_2 * u_3 \wedge 2 * 1.4858536352836396e-7 + u_1 \wedge 2 * u_2 \wedge 2 * -0.0003531614272747699 + u_1 \wedge 2 * u_2 \wedge 2 * u_3 * -1.9178976768869506e-6 + u_1 \wedge 2 * u_2 \wedge 2 * u_3 \wedge 2 * 3.8405659245262027e-10, -0.04932474700217403 + u_2 * 0.17406814677977456 + u_1 \wedge 2 * u_2 * -1.4594144102122378e-6]$


Universal ODE -> Internal Sparse Regression

Sparse Identification on only the missing term

Operation [$u_2 * 0.10234428543435758 + u_1 * u_2 * 0.11371750552005416 + u_1^2 * u_2 * 0.12635459799855597$] of $u=(S/N, I, D/N)$

$$\begin{aligned}
 S' &= -\frac{\beta_0 S F}{N} - \mu S, \\
 E' &= \frac{\beta_0 S F}{N} + (\sigma + \mu) E, \\
 I' &= \sigma E - (\gamma + \mu) I, \\
 R' &= \gamma I - \mu R, \\
 N' &= -\mu N, \\
 D' &= d \gamma I - \lambda D, \quad \text{and} \\
 C' &= \sigma E,
 \end{aligned}$$

Replace Unknown Portion



ML-Augmented Scientific Modeling


1. IDENTIFY KNOWN PARTS OF A MODEL, BUILD A UODE
2. TRAIN A NEURAL NETWORK (OR OTHER APPROXIMATOR) TO CAPTURE THE MISSING MECHANISMS
3. SPARSE IDENTIFY THE MISSING TERMS TO MECHANISTIC TERMS
4. VERIFY THE MECHANISMS ARE SCIENTIFICALLY PLAUSIBLE
5. EXTRAPOLATE, DO ASYMPTOTIC ANALYSIS, PREDICT BIFURCATIONS
6. GET MORE DATA TO VERIFY THE NEW TERMS

UTILIZE ALL ADVANCED NUMERICAL METHODS WITH ML!

U-ODE's for eVTOL Battery Modeling: 19% Increase in Degradation Modeling Accuracy



Coupled Electrochemical-Thermal Performance Model



U-ODE Degradation Model

Tesla Model S/X Mileage vs Remaining Battery Capacity



Data-Driven Quantification of Quarantine Strength

$$\frac{dS(t)}{dt} = -\frac{\beta S(t) I(t)}{N}$$


$$\frac{dI(t)}{dt} = \frac{\beta S(t) I(t)}{N} - (\gamma + Q(t)) I(t) = \frac{\beta S(t) I(t)}{N} - (\gamma + NN(W, U)) I(t)$$

$$\frac{dR(t)}{dt} = \gamma I(t) + \delta T(t)$$


$$\frac{dT(t)}{dt} = Q(t) I(t) - NN(W, U) I(t) - \delta T(t).$$

A machine learning aided global diagnostic and comparative tool to assess effect of quarantine control in Covid-19 spread


Raj Dandekar,¹ Chris Rackauckas,² and George Barbastathis^{3,4†}



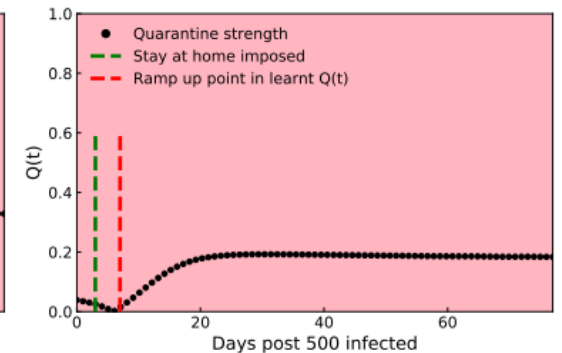
(a) New York



(b) New Jersey

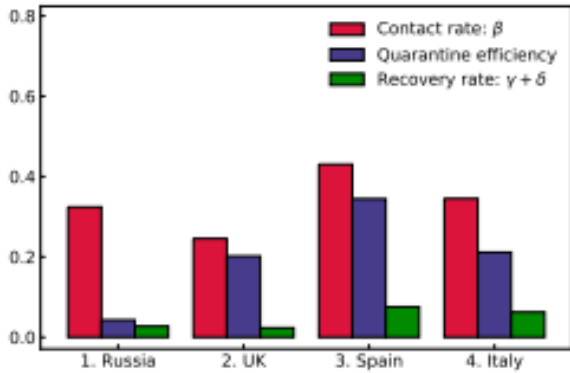


(c) Illinois

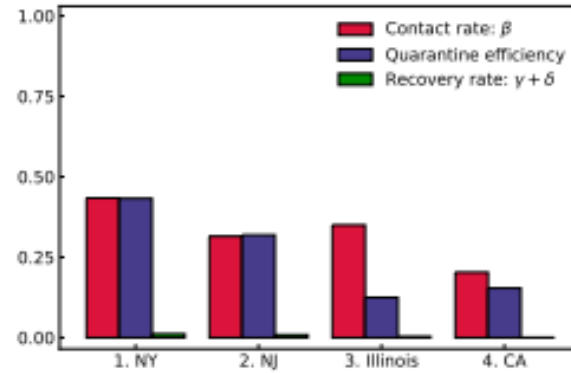


(d) California

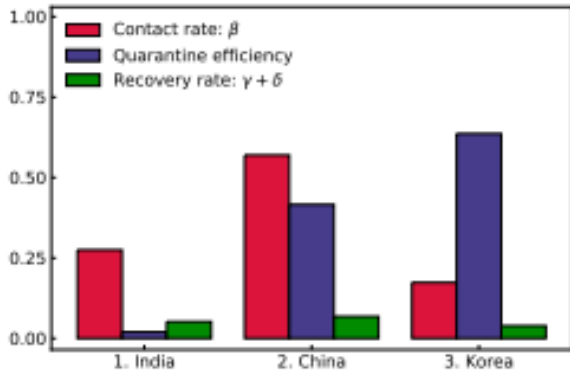
Diagnostics En Masse Reveal Interesting Trends



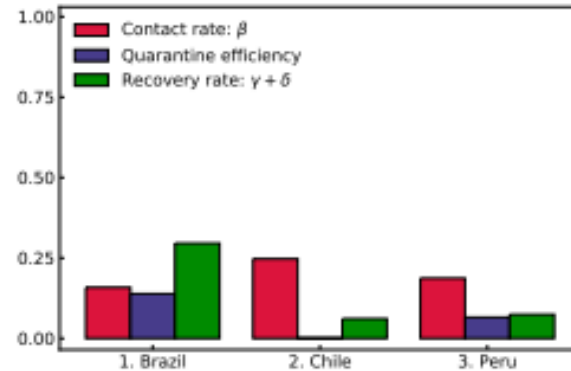
(a) Europe



(b) USA



(c) Asia



(d) South America

Figure 15: Global comparison of infection, recovery rates and quarantine efficiency.

But ODEs are simple, lets
move to more difficult
equations

Warning: these next few slides may be information overload if you're not familiar with scientific computing. That's okay! Take in what you can.

Universal Differential-Algebraic Equations: Encoding Physical Constraints

Utilize known chemical kinetics

$$y_1' = -0.04y_1 + NN_1(y_1, y_2, y_3)$$


$$y_2' = 0.04y_1 + NN_2(y_1, y_2, y_3)$$

$$1 = y_1 + y_2 + y_3$$

With known conservation laws

$$Mu' = f(u) + NN(u)$$

Convert to a mass-matrix DAE
(singular mass matrix) and fit



Learn highly stiff equations: **Hessian condition number 10^{13}**

Discretized PDE Operators are Convolutions

1 _{x1}	1 _{x0}	1 _{x1}	0	0
0 _{x0}	1 _{x1}	1 _{x0}	1	0
0 _{x1}	0 _{x0}	1 _{x1}	1	1
0	0	1	1	0
0	1	1	0	0

Image

4		

Convolved Feature

$$\frac{u(x + \Delta x, y) - 2u(x, y) + u(x - \Delta x, y)}{\Delta x^2} + \frac{u(x, y + \Delta y) - 2u(x, y) + u(x, y - \Delta y)}{\Delta y^2}$$

Is equivalent to the stencil

0	1	0
1	-4	1
0	1	0

$$\frac{u(x + \Delta x) - 2u(x) + u(x - \Delta x)}{\Delta x^2} = u''(x) + \mathcal{O}(\Delta x^2)$$

$$\Delta u = u_{xx} + u_{yy}$$

Automatically Learning PDEs from Data: Universal PDEs for Fisher-KPP

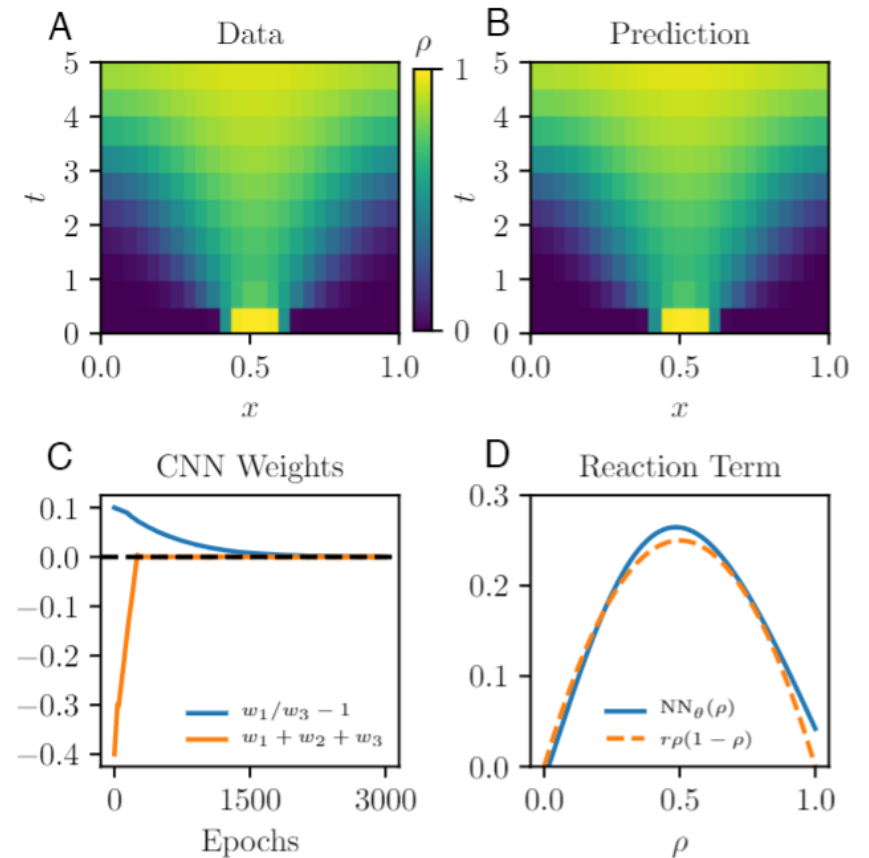
Truth: Fisher-KPP Equations

$$\rho_t = r\rho(1 - \rho) + D\rho_{xx},$$

Truth: Universal Differential Equation

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

Automatically recover that the dynamical system has a diffusion operator and a quadratic reaction term!



Embedding Neural Networks
into Scientific Simulation Can
Also Be Used To Accelerate!

Universal ODEs Accelerate Non-Newtonian Fluid Simulations

$$\sigma(t) = \int_{-\infty}^t G(t-s)F(\dot{\gamma}(s)) ds, \quad [15]$$

which depends on the history of deformation, with some memory function G (55). This is equivalent to the following instantaneous form:

$$\sigma(t) = \phi_1(t), \quad [16]$$

$$\frac{d\phi_1}{dt} = G(0)F(\dot{\gamma}) + \phi_2, \quad [17]$$

$$\frac{d\phi_2}{dt} = \frac{dG(0)}{dt}F(\dot{\gamma}) + \phi_3, \quad [18]$$

⋮

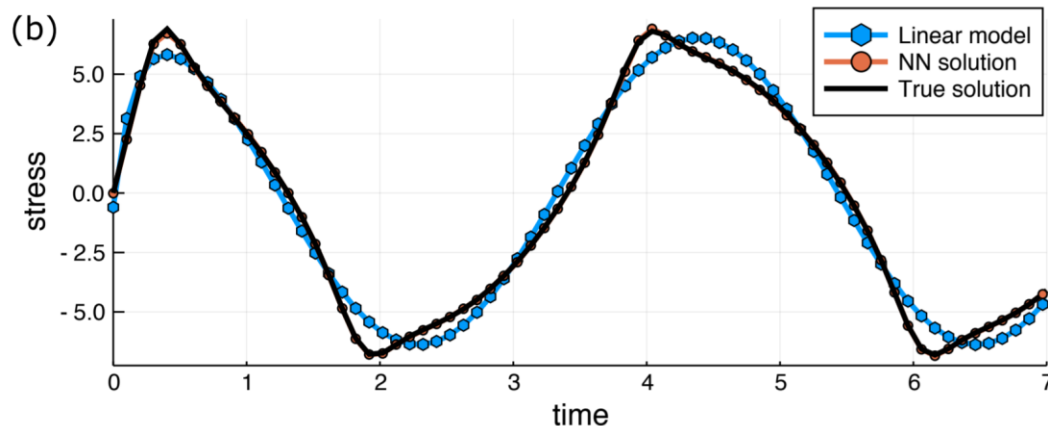
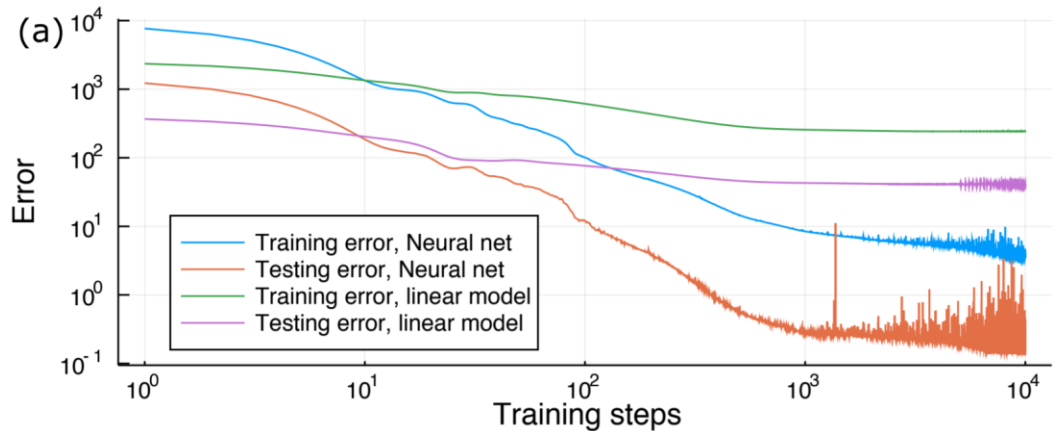
$$\sigma(t) = U_0(\dot{\gamma}, \phi_1, \dots, \phi_N),$$

$$\frac{d\phi_1}{dt} = U_1(\dot{\gamma}, \phi_1, \dots, \phi_N),$$

⋮

$$\frac{d\phi_N}{dt} = U_N(\dot{\gamma}, \phi_1, \dots, \phi_N),$$

Transform a system
Of DAEs into
Parameterized
system of ODEs,
2x acceleration



Universal PDEs for Acceleration: Automated Climate Parameterizations




Fig. S1. Reduction of the Boussinesq equations. On the left is the comparison between the training data (blue) and the trained UPDE (orange) over space at the 10th fitting time point, and on the right is the same comparison shown over time at spatial midpoint.

- ▶ Boussinesq Equations (Navier-Stokes) are used in climate models

$$\nabla \cdot \mathbf{u} = 0$$

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \text{Pr} \nabla^2 \mathbf{u} + b \hat{z}$$


$$\frac{\partial b}{\partial t} + \mathbf{u} \cdot \nabla b = \nabla^2 b + Fe^z$$

- ▶ People attempt to solve this by “parameterizing”, i.e. getting a 1-dimensional approximation through averaging:

$$\left(\frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla - \nabla^2 \right) \overline{c'} - \frac{\partial}{\partial z} \overline{w'c'} = -w' \frac{\partial \bar{c}}{\partial z}$$

where $\overline{w'c'}$ is unknown.

- ▶ Instead of picking a form for $\overline{w'c'}$ (the current method), replace it with a neural network and learn it from small scale simulations!




Universal Differential
Equations extend
previous physics-informed
neural network and deep
BSDE algorithms

UDE Methods Cover Accelerated Physics-Informed Neural Network Methods

$$l_n = \sum_{m=0}^{M-1} [\alpha_m y_{n-m} + \Delta t \beta_m f^{NN}(y_{n-m})], \quad n = M, \dots, N.$$

$$\text{loss}(f^{NN}(y^{NN}(t))) = \frac{1}{N_y} \sum_{n=1}^{N_y} (y^{NN}(t_n) - y^*(t_n))^2$$

$$+ \frac{1}{N_f} \sum_{n=1}^{N_f} \left(\frac{dy^{NN}}{dt}(t_n) - f^{NN}(y^{NN}(t_n)) \right)^2,$$



This methodology can be seen as a universal differential equation with a multistep integrator where adaptive=false

The UDE methodology thus gives an generalization to:

- ▶ Implicit methods, SSP methods
- ▶ Runge-Kutta-Chebyshev methods
- ▶ SDEs, DAEs, DDEs, etc.

A comparative study of physics-informed neural network models for learning unknown dynamics and constitutive relations Ramakrishna Tipireddy, Paris Perdikaris, Panos Stinis and Alexandre Tartakovsky

Multistep Neural Networks for Data-driven Discovery of Nonlinear Dynamical Systems Maziar Raissi, Paris Perdikaris, and George Em Karniadakis

Our results indicate that the accuracy of the trained neural network models is much higher for the cases where we only have to learn a constitutive relation instead of the whole dynamics.

Solving 1000 dimensional Hamilton-Jacobi-Bellman via Universal SDEs

- Semilinear Parabolic Form (Diffusion-Advection Equations, Hamilton-Jacobi-Bellman, Black-Scholes)

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

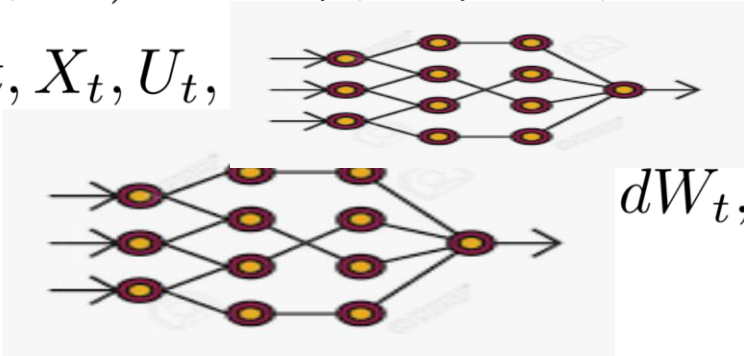
Then the solution of Eq. 1 satisfies the following BSDE (cf., e.g., refs. 8 and 9):

$$\begin{aligned} u(t, X_t) - u(0, X_0) &= - \int_0^t f(s, X_s, u(s, X_s), \sigma^T(s, X_s) \nabla u(s, X_s)) ds \\ &\quad + \int_0^t [\nabla u(s, X_s)]^T \sigma(s, X_s) dW_s. \end{aligned} \quad [3]$$

- Make $(\sigma^T \nabla u)(t, X)$ a neural network.
- Solve the resulting SDEs and learn $\sigma^T \nabla u$ via:

$$l(\theta) = \mathbb{E} \left[|g(X_{t_N}) - \hat{u}(\{X_{t_n}\}_{0 \leq n \leq N}, \{W_{t_n}\}_{0 \leq n \leq N})|^2 \right].$$

Forward-Backward Stochastic Neural Networks: Deep Learning of High-dimensional Partial Differential Equations
Maziar Raissi

$$\begin{aligned} dX_t &= \mu(t, X_t) dt + \sigma(t, X_t) dW_t, \\ dU_t &= f(t, X_t, U_t, \text{Neural Network}(\sigma^T \nabla u)) dt \\ &\quad + \text{Neural Network}(\sigma^T \nabla u) dW_t, \end{aligned}$$


Use high order, implicit, adaptive SDE solvers
Train a solution in minutes

Using non-adaptive explicit 0.5th order Euler-Maruyama matches the state-of-the-art deep BSDE methods from the literature

Solving high-dimensional partial differential equations using deep learning

Jiequn Han, Arnulf Jentzen, and Weinan E

UDEs are a BLAS/LAPACK of SciML

Scientific Machine Learning requires efficient
and accurate training of UDEs

Efficient and robust software for UDEs in the
Julia language result in efficient and robust
implementations for many algorithms



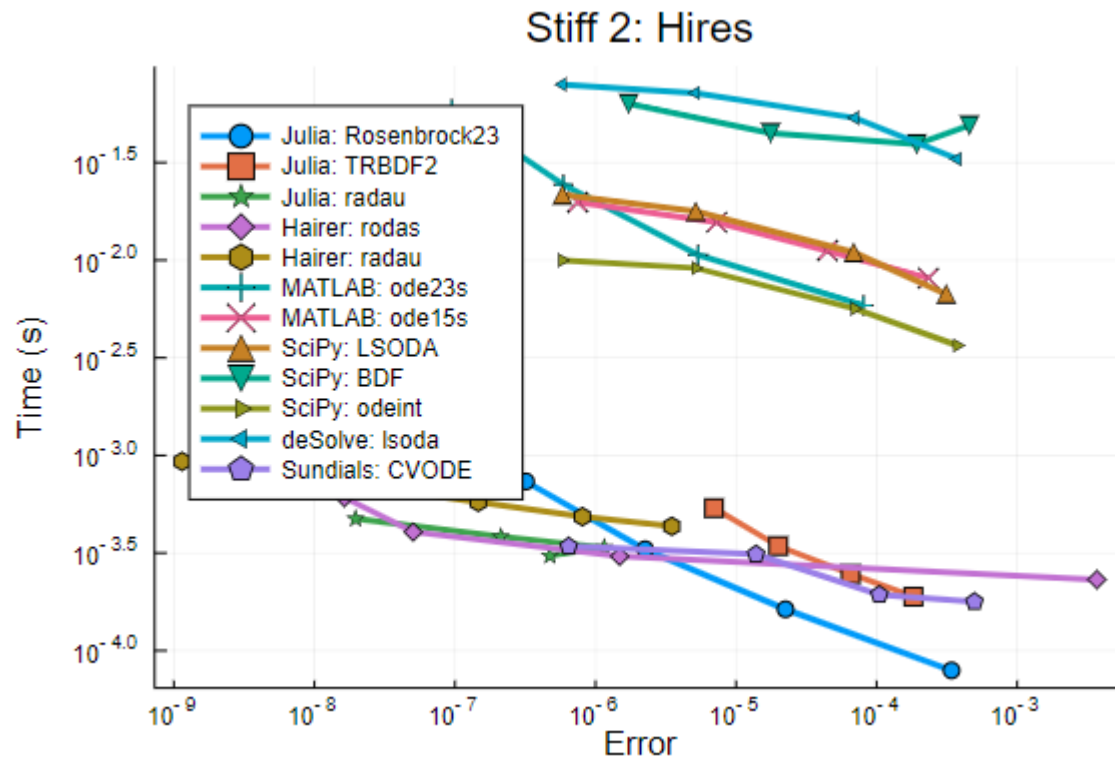
SciML Open Source Software Organization

sciml.ai


- ▶ DifferentialEquations.jl: high-performance differential equation solvers
- ▶ DiffEqFlux.jl: universal differential equation training optimizers, sensitivity analysis, and layer functions
- ▶ ModelingToolkit.jl: symbolic-numeric optimizations and automated parallelism
- ▶ NeuralPDE.jl: neural network solvers for PDEs, including automated physics-informed neural networks and deep BSDE methods for high dimensional PDEs
- ▶ Catalyst.jl: high-performance differentiable modeling of chemical reaction networks
- ▶ NBodySimulator.jl: high-performance differentiable molecular dynamics
- ▶ DataDrivenDiffEq.jl: Koopman Dynamic mode decomposition (DMD) methods and sparse identification (SInDy)

And 50 more libraries that cannot be fit!

SciML tools outperform ecosystems in high and low level languages



<https://github.com/SciML/SciMLBenchmarks.jl>



DifferentialEquations.jl's stiff ODE solvers can outperform SUNDIALS CVODE (C++) and Fortran methods like Radau

Speed drives researchers to Julia's SciML

Test problem: Lorenz equation

- ▶ DifferentialEquations.jl: 1.675 ms
- ▶ Jax: 3.66ms (*from author of Jax)
- ▶ Torchscript torchdifreq: 48 seconds

Simple neural ODE training (2-dimensional neural ODE from Neural Ordinary Differential Equations Chen et al.):

- ▶ DifferentialEquations.jl: ~3 seconds (will show live!)
- ▶ Torchdifreq: ~300 seconds

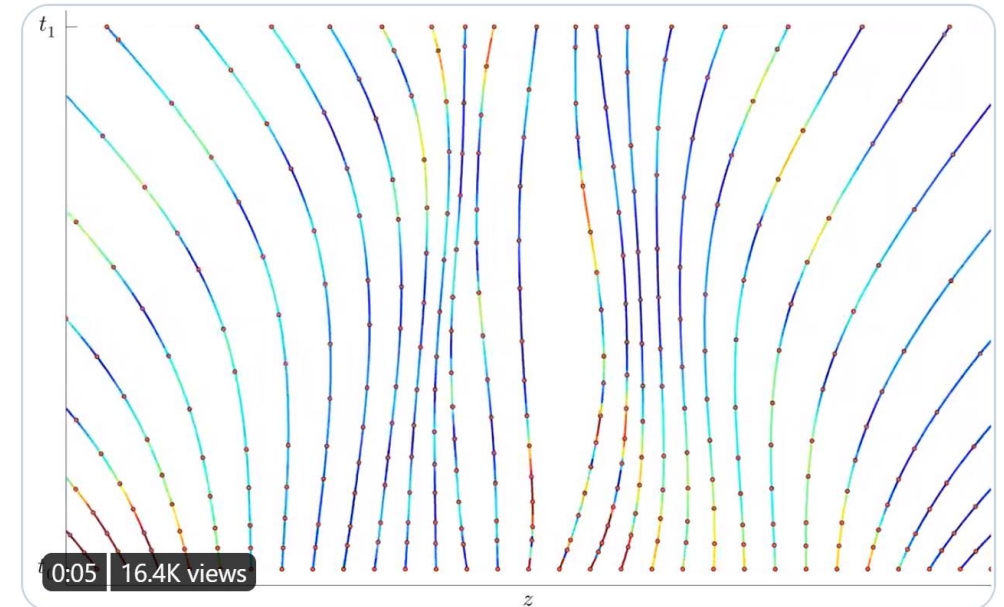


David Duvenaud @DavidDuvenaud · Jul 17

Neural ODEs are slow. We speed them up by regularizing their higher derivatives, learning ODEs that are easy to solve:

arxiv.org/pdf/2007.04504...

with @jacobjinkelly @jessebett @SingularMatrix



6

124

622



Jesse Bettencourt
@jessebett

Replying to @jessebett @TheyCallMeMr_ and 3 others

Here the training/evaluation of the dynamics neural network and its higher order derivatives are in JAX (python).

Not a problem in Julia, I used PyCall.jl to import the conda-environment + JAX code, calling my JAX neural nets, vmaps, and jets from within the Julia ODE solver!

4:12 PM · Jul 20, 2020 · Twitter Web App

Feature	SciML	Sundials (C++)	PETSc TS (C++)	torchdiffeq	Jax
Stiff ODEs and DAEs	Hundreds of methods tested and tuned on hundreds of problems	Yes (CVODE_BDF and IDA)	Yes (Rosenbrock-W methods, BDFs, etc.)	None	None (one in progress, ~200 times slower than SciPy according to the author!)
Adjoint Methods	8 choices tuned for different scenarios, including stabilized checkpointing, differentiate the solver, reversing adjoint	Stabilized checkpointing	Discrete sensitivity analysis (equivalent to differentiate through the solver)	Requires reversing the ODE or differentiate the solver	Requires reversing the ODE
Parallelism	GPU, MPI, multithreading	GPU, MPI, multithreading	GPU, MPI, and multithreading	GPU	GPU
Event handing	Yes	Yes	Yes	None	None
SDEs	Lots of methods, including stabilized, methods for stiff equations, high strong order, high weak order	None	None	torchsde, only diagonal noise (or order 0.5), requires reversing the SDE	None
Delays	All ODE methods	None	None	None	None

What do these features mean?

SciML is not just for speed

SCIML IS FOR FLEXIBILITY, ACCURACY, AND CORRECTNESS

WARNING: SOME "EXPERT" TALK HERE

SciML is meticulously tested

✓ **master** Merge pull request #1215 from utkarsh530/extplmethods

Multi-threading for Implicit Extrapolation Methods

Commit a09a464

Compare 96c3949...a09a464

Branch master

Christopher Rackauckas

🔄 #5442 passed

🕒 Ran for 1 hr 26 min 58 sec

🕒 Total time 6 hrs 16 min 8 sec

📅 18 hours ago

🔄 Restart build

Match behaviors exactly in pure Julia, and **fix bugs from the widely used Fortran code (deSolve, SciPy)**

Build jobs View config

✓ # 5442.1	AMD64	Xenial	</> Julia: 1	GROUP=Interfacel
✓ # 5442.2	AMD64	Xenial	</> Julia: 1	GROUP=Interfacell
✓ # 5442.3	AMD64	Xenial	</> Julia: 1	GROUP=Integrators_I
✓ # 5442.4	AMD64	Xenial	</> Julia: 1	GROUP=Integrators_II
✓ # 5442.5	AMD64	Xenial	</> Julia: 1	GROUP=Regression_I
✓ # 5442.6	AMD64	Xenial	</> Julia: 1	GROUP=Regression_II
✓ # 5442.7	AMD64	Xenial	</> Julia: 1	GROUP=AlgConvergence_I
✓ # 5442.8	AMD64	Xenial	</> Julia: 1	GROUP=AlgConvergence_II
✓ # 5442.9	AMD64	Xenial	</> Julia: 1	GROUP=AlgConvergence_III
✓ # 5442.10	AMD64	Xenial	</> Julia: 1	GROUP=Downstream
✓ # 5442.11	AMD64	Xenial	</> Julia: 1	GROUP=ODEInterfaceRegression

Full test suite is over a day of events, gradients, GPUs, convergence, stochastic distributions, etc.



🕒 44 min 59 sec
🕒 31 min 34 sec
🕒 34 min 53 sec
🕒 41 min 23 sec
🕒 31 min 45 sec
🕒 26 min 47 sec
🕒 37 min 57 sec
🕒 44 min 16 sec
🕒 28 min 43 sec
🕒 41 min 42 sec
🕒 12 min 9 sec

```
sol1 =solve(prob,DP5(),dt=1/8)
sol2 =solve(prob,dopri5(),dt=1/8)
@test sol1.t ≈ sol2.t
```




SciML's tools do not rely on properties which can fail to hold

```
using DifferentialEquations, ODEInterfaceDiffEq, Plots
function rober(du,u,p,t)
    y1,y2,y3 = u
    k1,k2,k3 = p
    du[1] = -k1*y1+k3*y2*y3
    du[2] = k1*y1-k2*y2^2-k3*y2*y3
    du[3] = k2*y2^2
    nothing
end
prob = ODEProblem(rober,[1.0,0.0,0.0],(0.0,1e5),(0.04,3e7,1e4))
sol = solve(prob,DP5(),reltol=1e-10, abstol=1e-10) # Fails!

sol = solve(prob,Rosenbrock23())
p1 = plot(sol,tspan=(1e-2,1e5),xscale=:log10)
prob_reverse = ODEProblem(rober,sol[end],(1e5,0.0),(0.04,3e7,1e4))
sol2 = solve(prob_reverse,Rosenbrock23())
p2 = plot(sol2,tspan=(1e-2,1e5),xscale=:log10)
plot(p1,p2)
```



Ill-Conditioned Gradients Cause Difficulties in Scientific Machine Learning



Understanding and mitigating gradient pathologies
in physics-informed neural networks
Sifan Wang, Yujun Teng, Paris Perdikaris

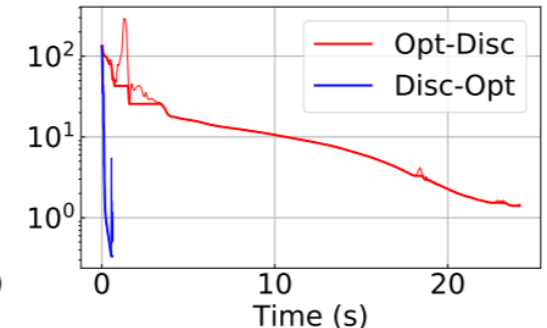

**Off-the-shelf ML tools will not work on stiff
scientific machine learning problems!**

DiffEqFlux has the features to handle stiff ill-conditioned scientific problems

- ▶ 300 highly optimized differential equation solvers for highly stiff equations:
 - ▶ ROCK methods
 - ▶ Implicit methods (ODEs, SDEs, DAEs, DDEs)
 - ▶ Multistep methods
 - ▶ SSP methods (hyperbolic PDEs)
 - ▶ Adaptive SDE solvers (implicit, high order)
 - ▶ Event handling
- ▶ And these implementations are well-optimized:
 - ▶ DiffEqFlux trains the neural ODE from the original neural ODE paper in ~3 seconds
 - ▶ torchscript torchdiffeq : ~300 seconds

Hessian condition number 10^{13} effectively trained in tutorials

- ▶ Mixed AD Hessian-free Newton-Krylov for robust second order optimization
- ▶ Discrete and continuous sensitivity analysis (checkpointed, stabilized, etc.)



Discretize-Optimize vs. Optimize-Discretize for Time-Series Regression and Continuous Normalizing Flows
Derek Onken, Lars Ruthotto

DiffEqFlux.jl has the bells and whistles to solve “real” problems

Neural ODE with batching on the GPU (without internal data transfers) with high order adaptive implicit ODE solvers for stiff equations using matrix-free Newton-Krylov via preconditioned GMRES and trained using checkpointed adjoint equations.

```
using OrdinaryDiffEq, Flux, DiffEqFlux, DiffEqOperators, CuArrays
x = Float32[2.; 0.]|>gpu
tspan = Float32.((0.0f0,25.0f0))
dudt = Chain(Dense(2,50,tanh),Dense(50,2))|>gpu
p = DiffEqFlux.destructure(dudt)
dudt_(du,u::TrackedArray,p,t) = u .= DiffEqFlux.restructure(dudt,p)(u)
dudt_(du,u::AbstractArray,p,t) = u .= Flux.data(DiffEqFlux.restructure(dudt,p)(u))
ff = ODEFunction(dudt_,jac_prototype = JacVecOperator(dudt_,x))
prob = ODEProblem(ff,x,tspan,p)
diffeq_adjoint(p,prob,KenCarp4(linsolve=LinSolveGMRES());u0=x,
              saveat=0.0:0.1:25.0,backsolve=false)
```


Workshop Outline

- ▶ ~~Overview of scientific machine learning and SciML~~
 - ▶ ~~What is scientific machine learning?~~
 - ▶ ~~What makes the SciML ecosystem unique?~~
- ▶ Modeling with differential equations
 - ▶ Solving differential equations with DifferentialEquations.jl
 - ▶ Adding stochasticity, delays, events
- ▶ Introduction to challenge and learning problems
 - ▶ Workshop exercises (with answers!)
 - ▶ HelicopterSciML Challenge Problem
 - ▶ **Magnetic Navigation Challenge Problem**
- ▶ Automated model discovery via universal differential equations
 - ▶ Parameter inference on differential equations
 - ▶ Local and global optimization
 - ▶ Bayesian optimization
 - ▶ Mixing DiffEqFlux.jl and DataDrivenDiffEq.jl!
- ▶ Solving differential equations with neural networks (physics-informed neural networks)

Now let's get to coding

SOLVING DIFFERENTIAL EQUATIONS WITH STOCHASTICITY, DELAYS, AND EVENTS

AND THEN ADD SOME PARALLELISM

Let's start coding some models: Lotka-Volterra Equations

43


$$\frac{d \text{🐰}}{dt} = \alpha \text{🐰} - \beta \text{🐰} \text{🐱}$$

Exponential growth Gets eaten by wolves

$$\frac{d \text{🐱}}{dt} = \delta \text{🐰} \text{🐱} - \gamma \text{🐱}$$

Increases with more food Decreases with competition

The Lotka-Volterra Equations: Model of Rabbits and Wolves



ODE Solver Packages on the Common Interface

You do not need to change your code to use solves on the common interface! Julia might have the largest number of active developers in the field! Other great solvers you should check out:

- ▶ OrdinaryDiffEq.jl: The workhorse
- ▶ Sundials.jl: CVODE_BDF is a great stiff ODE solver
- ▶ ODEInterfaceDiffEq.jl: radau is great for stiff ODEs at low tolerances ($<1e-8$)
- ▶ LSODA.jl: lsoda is all-around good for smaller ODEs (<100)
- ▶ IRKGaussLegendre.jl: IRKGL16 is 16th order and symplectic, great for physical problems at really low tolerances ($<1e-12$)
 - ▶ **JuliaCon: Implicit RK solver for high precision numerical integration**
- ▶ TaylorIntegration.jl: Great at low tolerances, can give error bounds
- ▶ NeuralPDE.jl: parallelized-in-time physics-informed neural network methods
 - ▶ **JuliaCon: Julia Track Google Code In and Beyond**
 - ▶ **JuliaCon: Minisymposium on Partial Differential Equations**
- ▶ GeometricIntegratorsDiffEq.jl: Great fixed time-step methods for small ODEs (symplectic)
- ▶ QuDiffEq.jl: Great ODE solvers if you have a quantum computer and need to output QASM
- ▶ TimeMachine.jl: A priori time stepping from Clima, great for multi-node MPI problems

(Differentiable) Modeling Frameworks

- ▶ ModelingToolkit.jl: symbolic-numeric for accelerated modeling
 - ▶ **JuliaCon: Auto-Optimization and Parallelism in DifferentialEquations.jl**
- ▶ Catalyst.jl: chemical reaction networks
- ▶ Petri.jl and AlgebraicPetri.jl: Petri networks and applied category theory
- ▶ NetworkDynamics.jl: dynamics on networks
 - ▶ **JuliaCon: NetworkDynamics.jl - Modeling dynamical systems on networks**
- ▶ PowerSimulationsDynamics.jl: dynamics of power grids
 - ▶ **JuliaCon: Crash Course in Energy Systems Modeling and Analysis with Julia**
- ▶ JuSDL.jl: causal modeling that can mix the various differential equations
 - ▶ **JuliaCon: Jusdl.jl - Julia Based System Description Language**
- ▶ BioEnergeticFoodWebs.jl: simulations of biomass flows
- ▶ QuantumOptics.jl: simulations of quantum systems
- ▶ DynamicalSystems.jl: dynamical systems and chaos analysis
- ▶ RigidBodySim.jl: simulations of rigid-body dynamics and robotics

And so many more!

github.com/epirecipes/sir-julia
Various implementations of the classical SIR model in Julia

Workshop Outline

- ▶ ~~Overview of scientific machine learning and SciML~~
 - ▶ ~~What is scientific machine learning?~~
 - ▶ ~~What makes the SciML ecosystem unique?~~
- ▶ ~~Modeling with differential equations~~
 - ▶ ~~Solving differential equations with DifferentialEquations.jl~~
 - ▶ ~~Adding stochasticity, delays, events~~
- ▶ Introduction to challenge and learning problems
 - ▶ Workshop exercises (with answers!)
 - ▶ HelicopterSciML Challenge Problem
 - ▶ **Magnetic Navigation Challenge Problem**
- ▶ Automated model discovery via universal differential equations
 - ▶ Parameter inference on differential equations
 - ▶ Local and global optimization
 - ▶ Bayesian optimization
 - ▶ Mixing DiffEqFlux.jl and DataDrivenDiffEq.jl!
- ▶ Solving differential equations with neural networks (physics-informed neural networks)



SciML challenge and learning problems

Workshop Exercise Sheet

- ▶ https://tutorials.sciml.ai/html/exercises/01-workshop_exercises.html
 - ▶ Lots of exercises, from beginner to advanced
 - ▶ Problems on performance optimization, parameter inference, neural ODEs

Solutions: https://tutorials.sciml.ai/html/exercises/02-workshop_solutions.html

HelicopterSciML Challenge Problem: Learn missing physics!



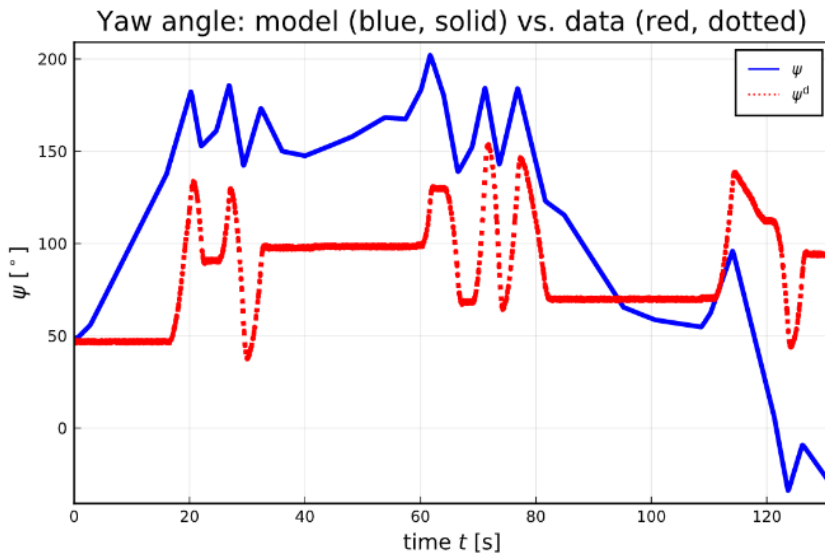
Goal: Discover the unexplained physics of this system

Figure 1: Laboratory helicopter, Sharma (2020).

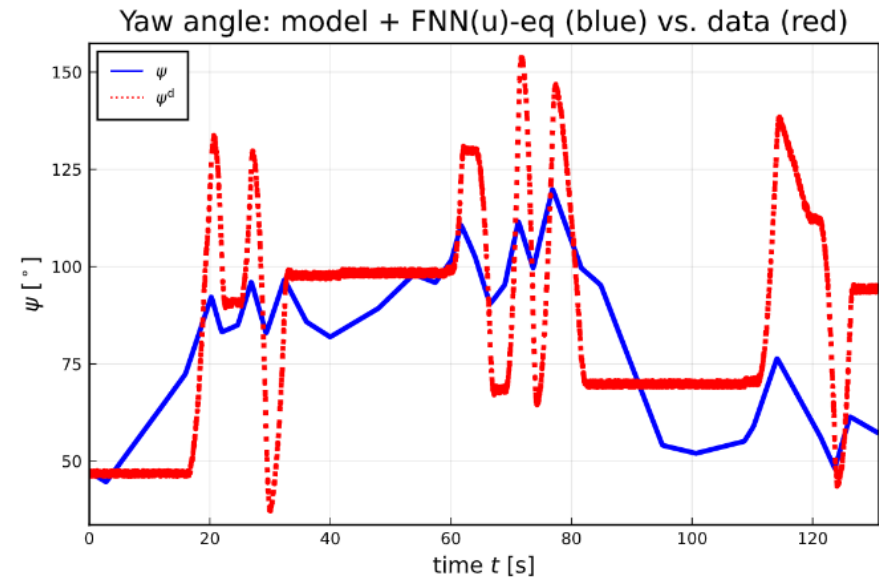
<https://github.com/SciML/HelicopterSciML.jl>

HelicopterSciML Challenge Problem Example Solution

Before Augmentation



After Discovery



Discovered missing higher
order friction terms

$$\begin{aligned} \text{FNN}(u; p)_1 &\approx -4.37 \cdot 10^{-4} \cos(u_\psi) + 4.02 \cdot 10^{-4} \sin(u_\psi) \\ \text{FNN}(u; p)_2 &\approx -1.35 \cdot 10^{-2} \cos(u_\psi) + 7.74 \cdot 10^{-3} u_\psi^2. \end{aligned}$$

Magnetic Navigation Challenge Problem

[HTTPS://GITHUB.COM/MIT-AI-ACCELERATOR/MAGNAV.JL](https://github.com/mit-ai-accelerator/magnav.jl)

Workshop Outline

- ▶ ~~Overview of scientific machine learning and SciML~~
 - ▶ ~~What is scientific machine learning?~~
 - ▶ ~~What makes the SciML ecosystem unique?~~
- ▶ ~~Modeling with differential equations~~
 - ▶ ~~Solving differential equations with DifferentialEquations.jl~~
 - ▶ ~~Adding stochasticity, delays, events~~
- ▶ ~~Introduction to challenge and learning problems~~
 - ▶ ~~Workshop exercises (with answers!)~~
 - ▶ ~~HelicopterSciML Challenge Problem~~
 - ▶ ~~**Magnetic Navigation Challenge Problem**~~
- ▶ Automated model discovery via universal differential equations
 - ▶ Parameter inference on differential equations
 - ▶ Local and global optimization
 - ▶ Bayesian optimization
 - ▶ Mixing DiffEqFlux.jl and DataDrivenDiffEq.jl!
- ▶ Solving differential equations with neural networks (physics-informed neural networks)

Now let's do some model inference

LEARN THE PARAMETERS OF A DIFFERENTIAL EQUATION
THEN LEARN THE MISSING PIECES OF A DIFFERENTIAL EQUATION

Universal ODEs learn and extrapolate other dynamical behaviors

Truth

$$\dot{x} = \alpha x - \beta xy,$$

$$\dot{y} = \gamma xy - \delta y.$$

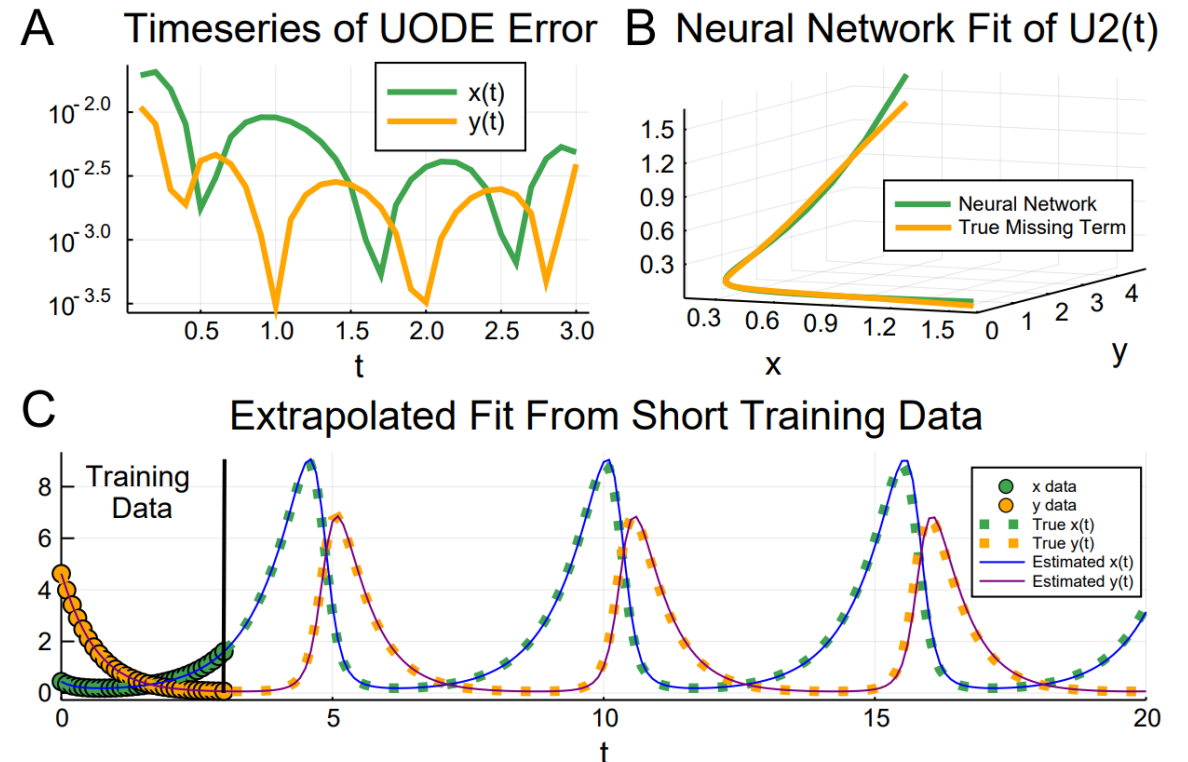
Partially-known neural embedded equations

$$\dot{x} = \alpha x - U_1(x, y),$$

$$\dot{y} = -\delta y + U_2(x, y),$$

Automatically recover the long-term behavior from less than half of a period in a cyclic time series!

Turn neural networks back into equations with SInDy. Let's do this example!



Packages for model inference

- ▶ DiffEqFlux.jl: helpers for performing inference on models. Interface over:
 - ▶ Optim.jl: workhorse optimizers like BFGS
 - ▶ Flux.jl: specialized neural network optimizers like ADAM
 - ▶ BlackBoxOptim.jl: very robust global optimizers
 - ▶ Evolutionary.jl: genetic algorithms and CMA
 - ▶ And many more!
- ▶ DataDrivenDiffEq.jl: methods for Koopman DMD and SInDy (turning data into equations!)
- ▶ Turing.jl: Bayesian estimation
- ▶ Gen
- ▶ GalacticOptim.jl: differentiable local+global optimizer interface. Coming soon!

Workshop Outline

- ▶ ~~Overview of scientific machine learning and SciML
 - ▶ ~~What is scientific machine learning?~~
 - ▶ ~~What makes the SciML ecosystem unique?~~~~
- ▶ ~~Modeling with differential equations
 - ▶ ~~Solving differential equations with DifferentialEquations.jl~~
 - ▶ ~~Adding stochasticity, delays, events~~~~
- ▶ ~~Introduction to challenge and learning problems
 - ▶ ~~Workshop exercises (with answers!)~~
 - ▶ ~~HelicopterSciML Challenge Problem~~
 - ▶ **Magnetic Navigation Challenge Problem**~~
- ▶ ~~Automated model discovery via universal differential equations
 - ▶ ~~Parameter inference on differential equations
 - ▶ ~~Local and global optimization~~
 - ▶ ~~Bayesian optimization~~~~
 - ▶ ~~Mixing DiffEqFlux.jl and DataDrivenDiffEq.jl!~~~~
- ▶ Solving differential equations with neural networks (physics-informed neural networks)

NeuralPDE.jl: Automated PDE Solving via Neural Networks

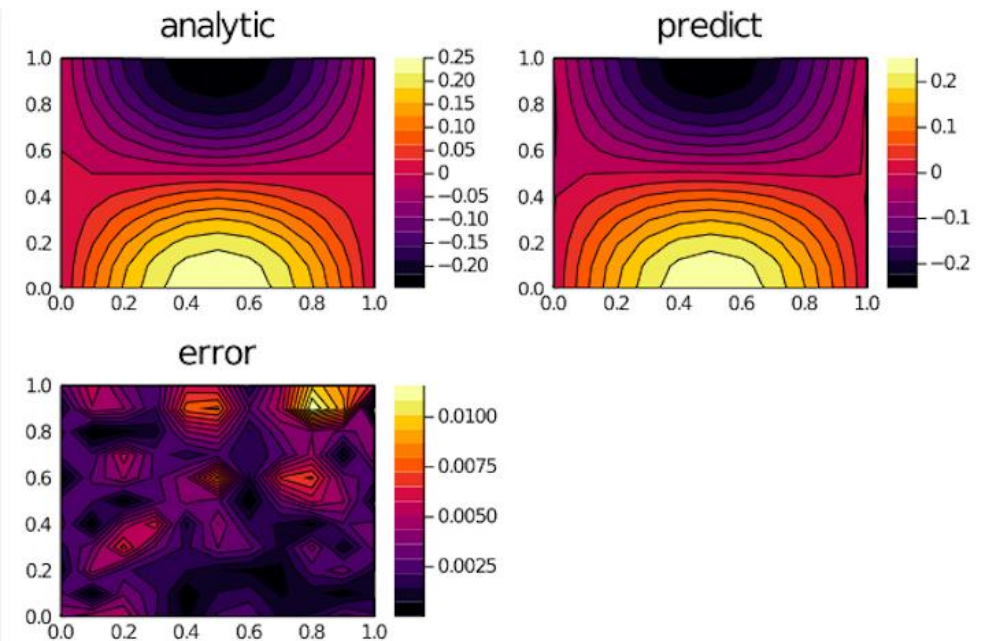
```
@parameters x y θ
@variables u(..)
@derivatives Dxx''~x Dyy''~y

eq = Dxx(u(x,y,θ)) + Dyy(u(x,y,θ)) ~ -sin(pi*x)*sin(pi*y)
bcs = [u(0,y) ~ 0.fθ, u(1,y) ~ -sin(pi*1)*sin(pi*y),
       u(x,0) ~ 0.fθ, u(x,1) ~ -sin(pi*x)*sin(pi*1)]
domains = [x ∈ IntervalDomain(0.0,1.0), y ∈ IntervalDomain(0.0,1.0)]
discretization = PhysicsInformedNN(0.1)

opt = Flux.ADAM(0.02)
chain = FastChain(FastDense(2,16,Flux.σ),FastDense(16,16,Flux.σ),FastDense(16,1))

pde_system = PDESystem(eq,bcs,domains,[x,y],[u])
prob = discretize(pde_system,discretization)
alg = NNDE(chain,opt,autodiff=false)

phi,res = solve(prob,alg,verbose=true, maxiters=5000)
```



What is this library doing?

- ▶ The deep BSDE method
 - ▶ Mentioned earlier: can be transformed into a universal stochastic differential equation and solved via `DiffEqFlux.jl`
- ▶ Physics-informed neural networks

For understanding, let's build the simplest physics-informed neural network from scratch!

Solve an ODE with a neural network

- ▶ Let $u' = f(u, t)$ with $u(0) = u_0$. We want to build a neural network $NN(t)$ that is the solution to this differential equation.
- ▶ By definition then, we must have that $NN'(t) = f>NN(t), t$ and $NN(0) = u_0$
- ▶ Define $\mathcal{C}(\theta) = \sum_t \|NN'(t) - f>NN(t), t\|$ for θ the parameters of the ODE
 - ▶ Then this cost is zero when $NN(t)$ is the solution to the ODE
 - ▶ Therefore minimize this loss to get the solution!
- ▶ Extra trick: $g(t) = tNN(t) - u_0$ is an approximator that always satisfies the boundary condition

Why Physics-Informed Neural Networks?

- ▶ $\mathcal{C}(\theta) = \mathcal{C}_{pde}(\theta) + \mathcal{C}_{boundary}(\theta) + \mathcal{C}_{data}(\theta)$ can nudge a model towards data
 - ▶ Equivalent to regularizing the neural network by a scientific equation
- ▶ Can train fast continuous surrogates by making the neural network parameter dependent

Time to build a physics-
informed neural network in
Flux!

Thank you! For more information, check out JuliaCon starting next week!

Probabilistic Optimization with the Koopman Operator, July 29th

SciML: Automatic Discovery of droplet fragmentation Physics, July 29th

Exploring Disease Vector Dynamics Under Environmental Change, July 29th

NetworkDynamics.jl - Modeling dynamical systems on networks, July 30th

Automated optimization and parallelism with DifferentialEquations.jl, July 31st

all feature SciML tools, along with many, many more!

Mix neural networks
with FDM, FVM, FEM,
pseudospectral
methods, implicit ODE
solvers, high order
adaptive SDE solver, ...

Julia's SciML software
ecosystem is built to
handle the sparse, stiff,
and ill-conditioned
problems of real
science