

Recent Advancements in Differential Equation Solver Software

CHRIS RACKAUCKAS

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

DifferentialEquations.jl: Research Platform for Production

- ▶ ~300 methods available, including wrappers to C/Fortran methods
 - ▶ **Platform for reproducible research and benchmarking**
- ▶ MPI+GPU Compatibility
- ▶ Implicit, IMEX, multirate, symplectic, exponential integrators, etc.
- ▶ **Adaptive high order methods for stochastic differential equations**
- ▶ Stiff state-dependent delay differential equation discontinuity tracking
- ▶ Mix in Gillespie simulation (Continuous-Time Markov Chains)
- ▶ Automatic sparsity detection and optimization
- ▶ Arbitrary code injection through callbacks

Native Julia methods routinely benchmark as one of the fastest libraries in most categories (caveat category: large (>1000) ODE/DAE systems)

3 Major Areas



Neural Differential Equations



A Hackable Model Compiler
(ModelingToolkit.jl)



Improvements to basic
numerical methods

There is a new field that is merging AI and domain-specific modeling: Scientific ML/AI

SCIENTIFIC AI: DOMAIN MODELS WITH INTEGRATED MACHINE LEARNING

[HTTPS://WWW.YOUTUBE.COM/WATCH?V=FGFX8CQHDQA](https://www.youtube.com/watch?v=FGFX8CQHDQA)

THE ESSENTIAL TOOLS OF SCIENTIFIC MACHINE LEARNING

[HTTP://WWW.STOCHASTICLIFESTYLE.COM/THE-ESSENTIAL-TOOLS-OF-SCIENTIFIC-MACHINE-LEARNING-SCIENTIFIC-ML/](http://www.stochasticlifestyle.com/the-essential-tools-of-scientific-machine-learning-scientific-ml/)

What is the
mathematical structure
of machine learning?

Neural Networks = Nonlinear Regression

- ▶ Polynomial: $e^x = a_1 + a_2x + a_3x^2 + \dots$
- ▶ Nonlinear: $e^x = 1 + \frac{a_1 \tanh(a_2x)}{a_3x - \tanh(a_4x)}$
- ▶ Neural Network: $e^x \approx W_3 \sigma(W_2 \sigma(W_1x + b_1) + b_2) + b_3$. Train the weights (W, b)

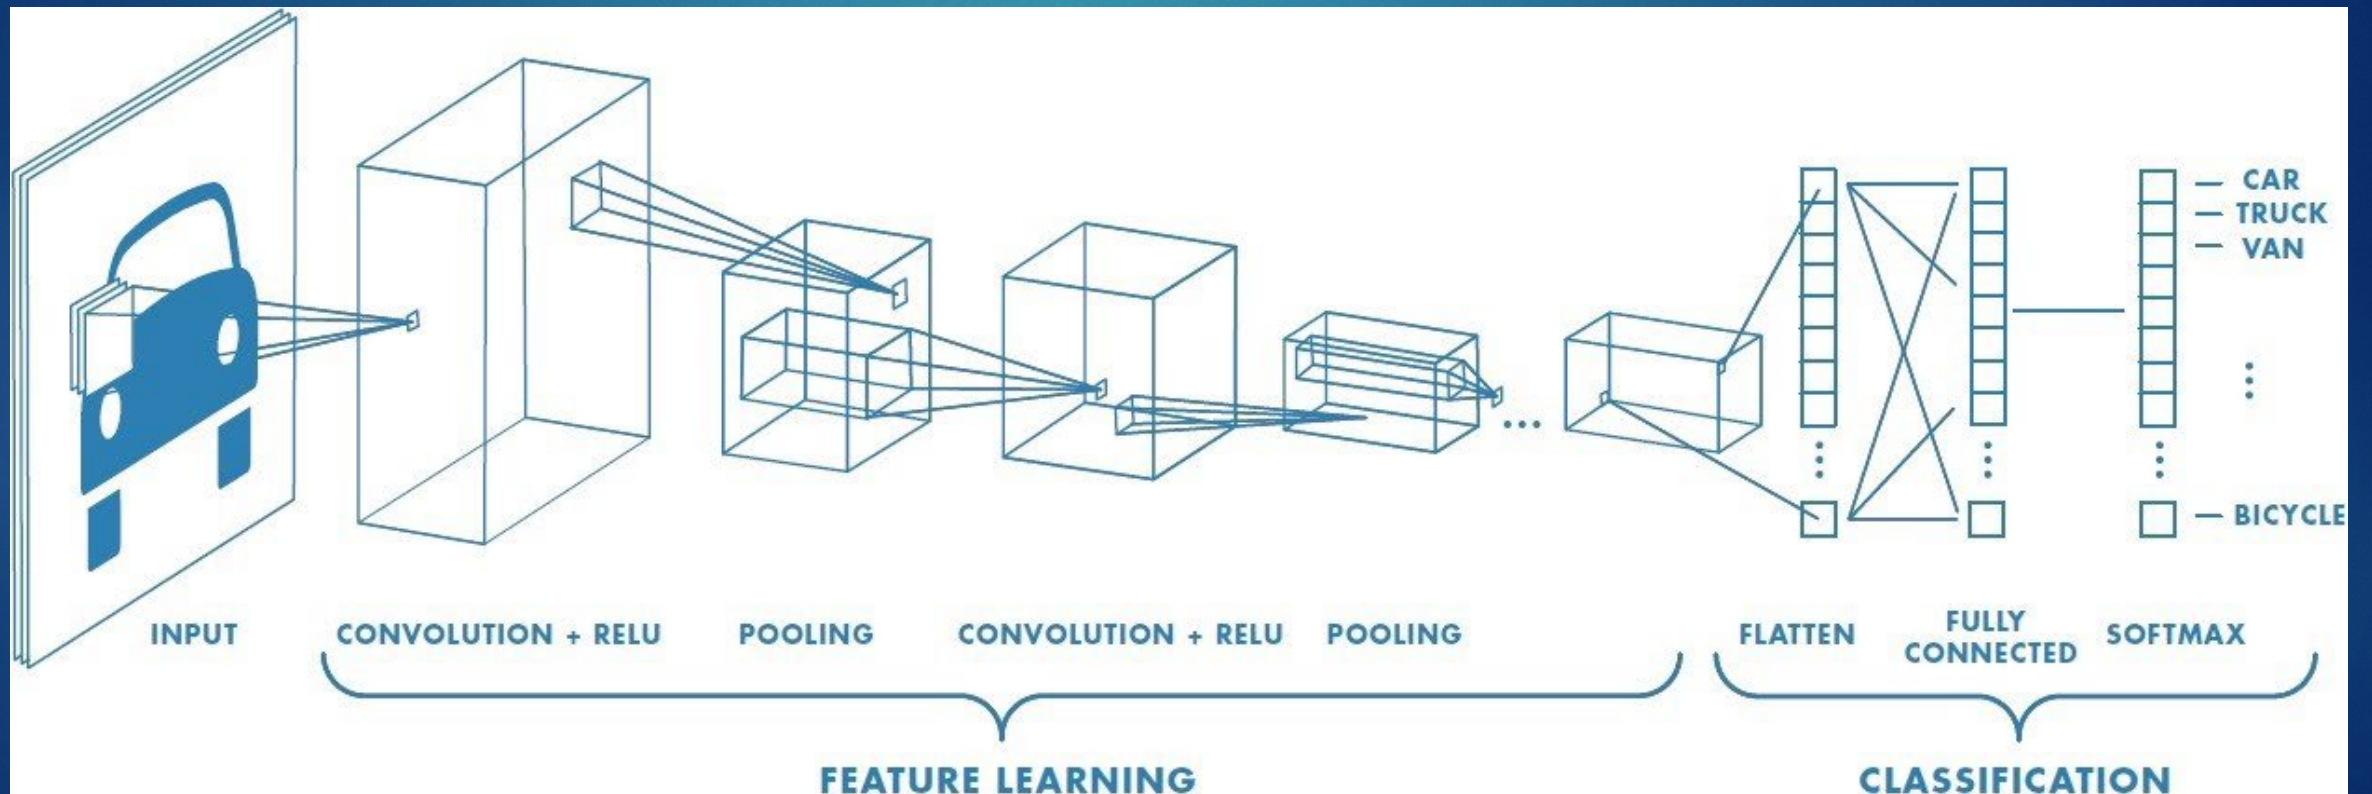
Universal Approximation Theorem

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

**Neural Networks Overcome “the curse of
dimensionality”**

Not Quite a Black Box: Convolutional Neural Networks Encode (Spatial) Structure

8



Now let's generalize
this idea to scientific
structures

Latent (Neural) Differential Equations

NEURAL ORDINARY DIFFERENTIAL EQUATION:


$$u' = f(u, p, t)$$

LET f BE A NEURAL NETWORK

Training a neural differential equation: DiffEqFlux.jl



Solve the differential equation



Compute the gradient of
the solution with respect to
the parameters defining the
neural network

Adjoint sensitivity
analysis
Differentiable
programming



Update the neural network and repeat

Automatically Learning the Model

```
File Edit View Juno Selection Find Packages Help
neural_ode.jl
1 using DiffEqFlux, OrdinaryDiffEq, Flux, Plots
2
3 # Generate data from a real ODE
4 u0 = Float32[2.; 0.]; datasize = 30
5 tspan = (0.0f0,1.5f0)
6 function trueODEfunc(du,u,p,t)
7     true_A = [-0.1 2.0; -2.0 -0.1]
8     du .= ((u.^3)'true_A)'
9 end
10 t = range(tspan[1],tspan[2],length=datasize)
11 prob = ODEProblem(trueODEfunc,u0,tspan)
12 ode_data = Array(solve(prob,Tsit5(),saveat=t))
13
14 # Define a Neural ODE
15 dudt = Chain(x -> x.^3,
16             Dense(2,75,tanh),
17             Dense(75,2))
18 n_ode(x) = neural_ode(dudt,x,tspan,AutoTsit5(Rosenbrock23(autodiff=false)),saveat=t,reltol=1e-7, abstol=1e-9)
19 function predict_n_ode()
20     n_ode(u0)
21 end
22 loss_n_ode() = sum(abs2,ode_data .- predict_n_ode())
23
24 # Train the Neural ODE to match the data
25 data = Iterators.repeated((), 200)
26 cb = function () #callback function to observe training
27     display(loss_n_ode()); cur_pred = Flux.data(predict_n_ode())
28     p1 = scatter(t,ode_data[1,:],label="data",legend=:bottomright); scatter!(p1,t,cur_pred[1,:],label="prediction")
29     p2 = scatter(t,ode_data[2,:],label="data",legend=:top); scatter!(p2,t,cur_pred[2,:],label="prediction")
30     display(plot(p1,p2,layout=(2,1)))
31 end
32 Flux.train!(loss_n_ode, Flux.params(dudt), data, Nesterov(0.0005), cb = cb)
33
REPL
julia> 
```

The real power comes from incorporating known structure into the ML framework (Mixed Neural Differential Equation)

Mix Neural Networks Into DiffEqs!

```

using DiffEqFlux, Flux, OrdinaryDiffEq

u0 = param(Float32[0.8; 0.8])
tspan = (0.0f0, 25.0f0)
ann = Chain(Dense(2, 10, tanh), Dense(10, 1))

p1 = Flux.data(DiffEqFlux.destructure(ann))
p2 = Float32[-2.0, 1.1]
p3 = param([p1; p2])
ps = Flux.params(p3, u0)

function dudt_(du, u, p, t)
    x, y = u
    du[1] = DiffEqFlux.restructure(ann, p[1:41])(u)[1]
    du[2] = p[end-1]*y + p[end]*x
end

prob = ODEProblem(dudt_, u0, tspan, p3)
function predict_adjoint()
    diffeq_adjoint(p3, prob, Tsit5(), u0=u0, saveat=0.0:0.1:25.0)
end

loss_adjoint() = sum(abs2, x-1 for x in predict_adjoint())
Flux.train!(loss_adjoint, ps, Iterators.repeated((), 10), ADAM(0.1))


```

$$\frac{dx}{dt} = \text{NN}$$

$$\frac{dy}{dt} = p_1 y + p_2 x$$

Fit the “mixed neural differential equation”
using the same method!

ML-Assisted Model Discovery



The chemical reactions imply an evolution of:

$$d[RA_{out}] = (\beta - b[RA_{out}] + c[RA_{in}])dt,$$

$$d[RA_{in}] = \left(b[RA_{out}] + \delta[RA - BP] - \left(\gamma[BP] + \eta + \text{?} \quad - c \right) [RA_{in}] \right) dt + \sigma dW_t,$$

$$d[RA - BP] = (\gamma[BP][RA_{in}] + \lambda[BP][RA - RAR] - \text{?} \quad dt,$$

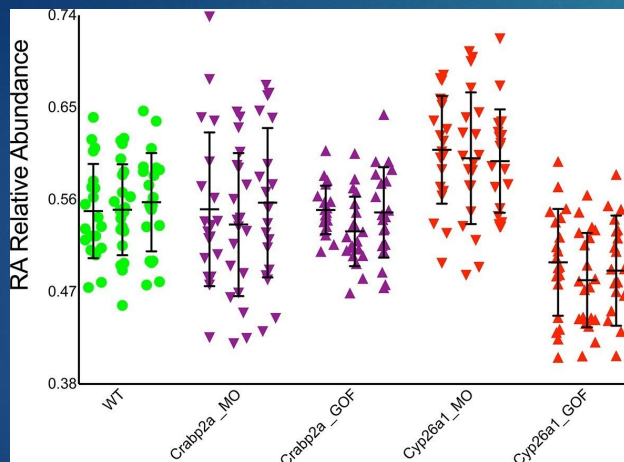
$$d[RA - RAR] = \text{?} \quad - \lambda[BP][RA - RAR])dt,$$

$$d[RAR] = \text{?} \quad dt,$$

$$d[BP] = (a - \lambda[BP][RA - RAR] - \gamma[BP][RA_{in}] + (\delta + \nu[RAR])[RA - BP] - u[BP])dt,$$

Biologically-Informed Neural Network

Data



Find neural networks so the model matches the data

$$d[RA_{out}] = (\beta - b[RA_{out}] + c[RA_{in}])dt,$$


$$d[RA_{in}] = \left(b[RA_{out}] + \delta[RA - BP] - \left(\gamma[BP] + \eta + \text{[Neural Network]} - c \right) [RA_{in}] \right) dt + \sigma dW_t,$$

$$d[RA - BP] = \left(\gamma[BP][RA_{in}] + \lambda[BP][RA - RAR] - \text{[Neural Network]} - \lambda[BP][RA - RAR] \right) dt,$$


$$d[RA - RAR] = \left(\text{[Neural Network]} - \lambda[BP][RA - RAR] \right) dt,$$

$$d[RAR] = \left(\text{[Neural Network]} \right) dt,$$

$$d[BP] = (a - \lambda[BP][RA - RAR] - \gamma[BP][RA_{in}] + (\delta + v[RAR])[RA - BP] - u[BP])dt,$$




Interpretability of Neural Differential Equations



Analyze the
Jacobian/Hessian

Data-Efficient Physics-Embedded Machine Learning



Nonlinear Optimal Control as a Mixed Neural ODE

- ▶ $x' = f(x(t), u(t), t)$
- ▶ Minimize $J = \Phi(x(t_0), t_0, x(t_f), t_f) + \int_{t_0}^{t_f} L(x(t), u(t), t) dt$
- ▶ Example: $x(t)$ is the location of an automated drone, $u(t)$ is the controller, find what the controller should be such that the vehicle goes to the right place for the least energy.
- ▶ Neural ODE Approach: Make $u(t)$ be a neural network. Find the neural network s.t. $x(t)$ correctly evolves

Neural PDEs for Acceleration of Navier-Stokes

```

37 #--- Run normal ODE solve
38
39 p = [a,b]  |> Matrix{Float32}[2]
40 prob = ODEProblem(weq, u0, (0.0, T), p)  |> ODEProblem with uType Array{Float32,1}
41 sol = solve(prob, SSPRK83(), dt=dt, progress=true, abstol=1e-6, reltol=1e-6)
42
43 function predict_adjoint()
44     diffeq_adjoint(p, prob, SSPRK83(), dt=dt, progress=true, abstol=1e-6, reltol=1e-6)
45 end  |> predict_adjoint
46
47 loss_adjoint() = 1/sum(abs2, x for x in Tracker.collect(predict_adjoint, 100))
48
49 cb = () -> display(loss_adjoint())  |λ
50 cb()  |✓
51
52 Flux.train!(loss_adjoint, params(p...), Iterators.repeated((), 3), ADAM(), cb)
53
54 #---
55
56 # field=Array(sol)
57 # vmax = maximum(abs.(field[:,1,:]))/3
58

```

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

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

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

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

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

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

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! Discretize. Result: Neural ODE.

Solving 1000 dimensional PDEs: Hamilton-Jacobi-Bellman, Nonlinear Black-Scholes

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

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

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

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

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

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

Simplified:

- ▶ Transform it into a Backwards SDE.
- ▶ The unknown is a function!
- ▶ Learn the unknown function via neural network.
- ▶ Once learned, the PDE solution is known.

Solving high-dimensional partial differential equations using deep learning, 2018, PNAS, Han, Jentzen, E

Represent 1000 dimensional PDEs as a 1000 dimensional neural SDE

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

- ▶ Solving the PDE = training the neural network
- ▶ As a neural SDE, we can solve with higher order (less neural network evaluations), adaptivity, etc.

DiffEqFlux.jl

Unified Framework for Scientific ML

The first (mixed) Neural Differential Equation solver.
Supports:

- ▶ Neural ODEs
- ▶ Neural SDEs (SDDEs)
- ▶ Neural DAEs
- ▶ Neural DDEs
- ▶ Stiff Equations
- ▶ Hybrid Equations
- ▶ Adjoints via reverse-mode AD and adjoint sensitivity analysis
- ▶ **Data-Efficient and Physical Machine Learning**

ModelingToolkit: An Open Source Compiler + IR for Models and Transformations

Defining Principles

24

- ▶ A structured and documented IR for describing models
- ▶ “Model transforms” are compiler passes IR \rightarrow IR
- ▶ Extendable high level “context” system
- ▶ **Let others write domain-specific languages on top**
 - ▶ DSLs should not have to define simplification, differentiation, etc!

```
1 using ModelingToolkit
2
3 # Define some variables
4 @parameters t σ ρ β
5 @variables x(t) y(t) z(t)
6 @derivatives D'~t
7
8 eqs = [D(x) ~ σ*(y-x),
9        D(y) ~ x*(ρ-z)-y,
10       D(z) ~ x*y - β*z]
11
12 de = ODESystem(eqs)
13 generate_function(de, [x,y,z], [σ,ρ,β])
14
15 ## Which returns:
16 :((##363, u, p, t)->begin
17     let (x, y, z, σ, ρ, β) = (u[1], u[2], u[3], p[1], p[2], p[3])
18         ##363[1] = σ * (y - x)
19         ##363[2] = x * (ρ - z) - y
20         ##363[3] = x * y - β * z
21     end
22 end)
23
```



```

1 function lorenz(du,u,p,t)
2   du[1] = p[1]*(u[2]-u[1])
3   du[2] = u[1]*(p[2]-u[3]) - u[2]
4   du[3] = u[1]*u[2] - p[3]*u[3]
5 end   > lorenz
6
7 using ModelingToolkit   ✓
8 @parameters t σ ρ β   (t(), σ(), ρ(), β())
9 @variables x(t) y(t) z(t)   (x(t()), y(t()), z(t()))
10 u = [x,y,z]   > Operation[3]
11 du = similar(u)   > Operation[3]
12 p = [σ,ρ,β]   > Operation[3]
13 lorenz(du,u,p,t)   x(t()) * y(t()) - β() * z(t())
14 du   ▼ Operation[3]
15     σ() * (y(t()) - x(t()))
16     x(t()) * (ρ() - z(t())) - y(t())
17     x(t()) * y(t()) - β() * z(t())
18 @derivatives Dx'~x Dy'~y Dz'~z   (> Differential, > Differential, > Differential)
19 J = [Dx(du[1]) Dy(du[1]) Dz(du[1])
20      Dx(du[2]) Dy(du[2]) Dz(du[2])
21      Dx(du[3]) Dy(du[3]) Dz(du[3])]   > 3x3 Array{Operation,2}:
22 J = expand_derivatives.(J)   ▼ 3x3 Array{Expression,2}:
23     σ() * -1           σ()   Constant(0)
24     ρ() - z(t())   Constant(-1)   x(t()) * -1
25     y(t())           x(t())       -1 * β()
26

```

Automatically
convert
numerical
functions to
symbolic

Current Research

26

```
1  using ModelingToolkit, DiffEqOperators, DiffEqBase, LinearAlgebra
2
3  # Define some variables
4  @parameters t x
5  @variables u(..)
6  @derivatives Dt'~t
7  @derivatives Dxx''~x
8  eq = Dt(u(t,x)) ~ Dxx(u(t,x))
9  bcs = [u(0,x) ~ - x * (x-1) * sin(x),
10         u(t,0) ~ 0, u(t,1) ~ 0]
11
12  domains = [t ∈ IntervalDomain(0.0,1.0),
13             x ∈ IntervalDomain(0.0,1.0)]
14
15  pdesys = PDESystem(eq,bcs,domains,[t,x],[u])
16  discretization = MOLFiniteDifference(0.1)
17  prob = discretize(pdesys,discretization)
18
19  using OrdinaryDiffEq
20  sol = solve(prob,Tsit5(),saveat=0.1)
```

- ▶ Components and Pantelides for large DAE systems
- ▶ An extendable framework for automated PDE discretizations
- ▶ Methods for (nonlinear) model order reduction
- ▶ Transformations for SDEs
- ▶ Tearing and other structural optimizations

Now assume the form
of the differential
equation is “good”

CAN WE IMPROVE THE SOLVERS?

Non-Stiff Methods
are still being
improved!

The Structure of a Runge-Kutta Method

$$y_{n+1} = y_n + h \sum_{i=1}^s b_i k_i,$$

where

$$k_1 = f(t_n, y_n),$$

$$k_2 = f(t_n + c_2 h, y_n + h(a_{21} k_1)),$$

$$k_3 = f(t_n + c_3 h, y_n + h(a_{31} k_1 + a_{32} k_2)),$$

$$\vdots$$

$$k_s = f(t_n + c_s h, y_n + h(a_{s1} k_1 + a_{s2} k_2 + \cdots + a_{s,s-1} k_{s-1})).$$

0					
c_2	a_{21}				
c_3	a_{31}	a_{32}			
\vdots	\vdots		\ddots		
c_s	a_{s1}	a_{s2}	\cdots	$a_{s,s-1}$	
	b_1	b_2	\cdots	b_{s-1}	b_s


Ways to Judge an RK Method

30

Optimization of next order coefficients

$$\begin{aligned} b_2 a_{21} + b_3 [a_{31} + a_{32}] + b_4 [a_{41} + a_{42} + a_{43}] &= 1/2 \\ b_2 a_{21}^2 + b_3 [a_{31} + a_{32}]^2 + b_4 [a_{41} + a_{42} + a_{43}]^2 &= 1/3 \\ b_2 a_{22} + b_3 [a_{21} a_{32} + a_{33}] + b_4 [a_{21} a_{42} + a_{43} (a_{31} + a_{32}) + a_{44}] &= 1/6 \\ b_2 a_{21}^3 + b_3 [a_{31} + a_{32}]^3 + b_4 [a_{41} + a_{42} + a_{43}]^3 &= 1/4 \\ b_2 a_{21} a_{22} + b_3 \left[\frac{1}{2} a_{21}^2 a_{32} + (a_{31} + a_{32}) (a_{21} a_{32} + a_{33}) \right] + \frac{1}{2} b_4 [a_{21}^2 a_{42} \\ + a_{43} (a_{31} + a_{32})^2 + 2 (a_{41} + a_{42} + a_{43}) (a_{21} a_{42} + (a_{31} + a_{32}) a_{43} + a_{44})] &= 1/6 \\ b_3 a_{22} a_{32} + b_4 [a_{21} a_{32} a_{43} + a_{22} a_{42} + a_{33} a_{43}] &= 1/24 \\ b_2 a_{21}^4 + b_3 [a_{31} + a_{32}]^4 + b_4 [a_{41} + a_{42} + a_{43}]^4 &= 1/5 \\ 3 b_2 a_{21}^2 a_{22} + b_3 [a_{21}^3 a_{32} + 3 (a_{31} + a_{32})^2 (a_{21} a_{32} + a_{33})] + b_4 [a_{21}^3 a_{42} \\ + (a_{31} + a_{32})^3 a_{43} + 3 (a_{41} + a_{42} + a_{43})^2 (a_{21} a_{42} + (a_{31} + a_{32}) a_{43} + a_{44})] &= 7/20 \\ b_3 a_{21}^2 a_{32} (a_{31} + a_{32}) + b_4 [(a_{41} + a_{42} + a_{43}) (a_{21}^2 a_{42} + (a_{31} + a_{32})^2 a_{43})] &= 1/15 \\ \frac{1}{2} b_2 a_{22}^2 + b_3 [a_{21} a_{32} (\frac{1}{2} a_{21} a_{32} + a_{22} + a_{33}) + a_{22} a_{32} (a_{31} + a_{32}) + \frac{1}{2} a_{33}^2] \\ + b_4 [\frac{1}{2} a_{21}^2 (a_{32} a_{43} + a_{42}^2) + (a_{31} + a_{32}) (a_{21} (a_{32} a_{43} + a_{42} a_{43}) + a_{43} (a_{33} + a_{44}) \\ + \frac{1}{2} (a_{31} + a_{32}) a_{43}^2) + a_{21} a_{42} (a_{22} + a_{44}) + (a_{21} a_{32} a_{43} + a_{22} a_{42} \\ + a_{33} a_{43}) (a_{41} + a_{42} + a_{43}) + \frac{1}{2} a_{44}^2] &= 11/120 \\ b_4 a_{22} a_{32} a_{43} &= 1/120 \end{aligned}$$

Stability



Dormand-Prince 5th Order (1980)

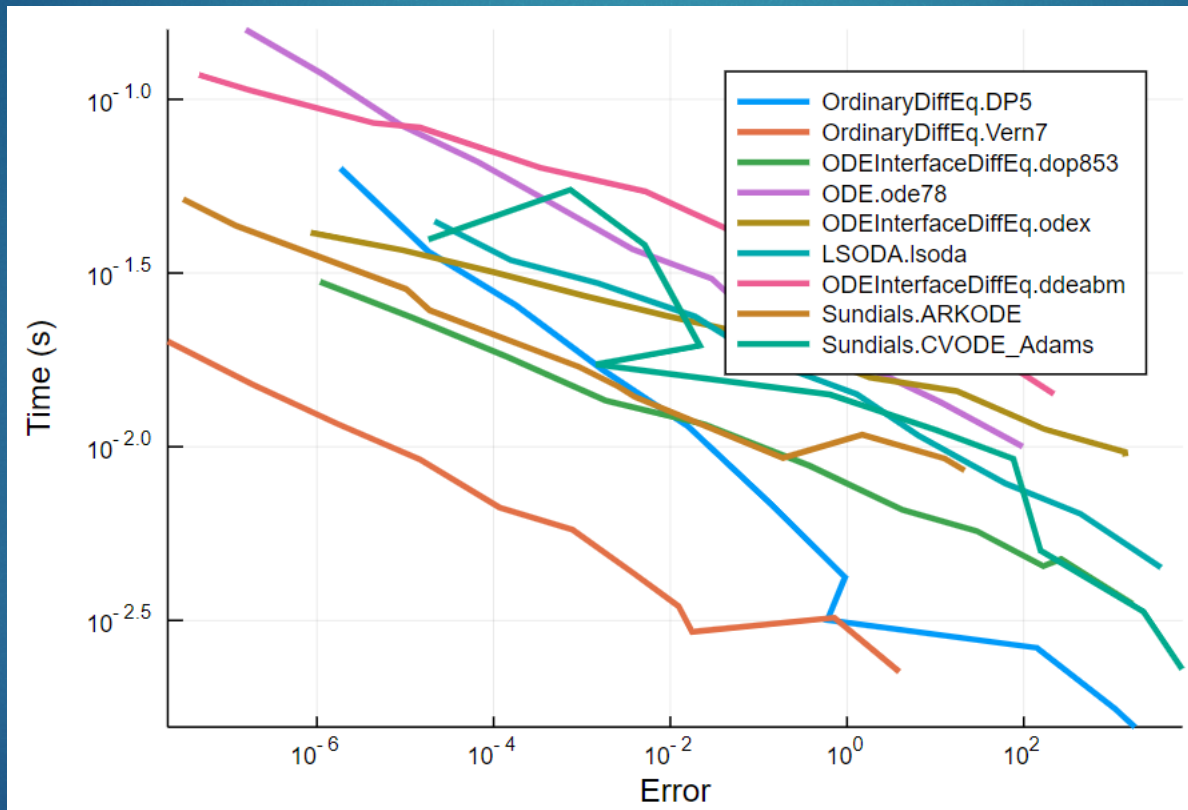
0							
1/5	1/5						
3/10	3/40	9/40					
4/5	44/45	-56/15	32/9				
8/9	19372/6561	-25360/2187	64448/6561	-212/729			
1	9017/3168	-355/33	46732/5247	49/176	-5103/18656		
1	35/384	0	500/1113	125/192	-2187/6784	11/84	
	35/384	0	500/1113	125/192	-2187/6784	11/84	0
	5179/57600	0	7571/16695	393/640	-92097/339200	187/2100	1/40

Advancements since 2010

Recent methods, Tsit5 and Vern#, reduce the number of assumptions made in coefficient optimization, leading to more optimal solutions (>2010)


Methods specialized for wave equations, low-dispersion results, extended monotonicity equation for PDEs (SSPRK), etc. are hot topics in new high order Runge-Kutta methods (2017)

100x100 Linear ODEs



3-Body Problem (CVODE_Adams fails)

34



And parallelism is
not well exploited!


Pervasive Allowance of Within-Method parallelism through Julia

Zero GPU/Distributed message passing done by the solver!

```
@muladd function perform_step!(integrator, cache::Tsit5Cache, repeat_step=false)
    @unpack t,dt,uprev,u,f,p = integrator
    @unpack c1,c2,c3,c4,c5,c6,a21,a31,a32,a41,a42,a43,a51,a52,a53,a54,a61,a62,a63,a64,a65,a
    @unpack k1,k2,k3,k4,k5,k6,k7,utilde,tmp,atmp = cache
    a = dt*a21
    @. tmp = uprev+a*k1
    f(k2, tmp, p, t+c1*dt)
    @. tmp = uprev+dt*(a31*k1+a32*k2)
    f(k3, tmp, p, t+c2*dt)
    @. tmp = uprev+dt*(a41*k1+a42*k2+a43*k3)
    f(k4, tmp, p, t+c3*dt)
    @. tmp = uprev+dt*(a51*k1+a52*k2+a53*k3+a54*k4)
    f(k5, tmp, p, t+c4*dt)
    @. tmp = uprev+dt*(a61*k1+a62*k2+a63*k3+a64*k4+a65*k5)
```

Multithreading Extrapolation

Simultaneous
Euler stepping
of different
step sizes



Parallel Runge-Kutta methods

38

$$\mathbf{A} = \begin{bmatrix} 0 & & & & \\ \times & 0 & & & \\ \times & 0 & 0 & & \\ \times & \times & \times & 0 & \\ \times & \times & \times & 0 & 0 \end{bmatrix}$$

5 stages

But only 3 steps in parallel

DiffEqGPU.jl: Automatic GPU-based parameter parallelism of high level code

39

```
using DiffEqGPU, CuArrays, OrdinaryDiffEq, Test

function lorenz(du,u,p,t)
  @inbounds begin
    du[1] = p[1]*(u[2]-u[1])
    du[2] = u[1]*(p[2]-u[3]) - u[2]
    du[3] = u[1]*u[2] - p[3]*u[3]
  end
  nothing
end

CuArrays.allowscalar(false)
u0 = Float32[1.0;0.0;0.0]
tspan = (0.0f0,100.0f0)
p = (10.0f0,28.0f0,8/3f0)
prob = ODEProblem(lorenz,u0,tspan,p)
prob_func = (prob,i) -> remake(prob,p=rand(Float32,3).*p)
monteprob = EnsembleProblem(prob, prob_func = prob_func)

#Performance check with nvvp
# CUDAnative.CUDAdrv.@profile
@time solve(monteprob,Tsit5(),EnsembleGPUArray(),trajectories=100_000,saveat=1.0f0)
```

- ▶ Demonstrates a 12x-90x speedup over multithreaded ODE solves by using just 1 GPU.
- ▶ Handles stiff and non-stiff hybrid ODEs/SDEs/DDEs/DAEs with adaptive timestepping.
- ▶ Support coming soon:
 - ▶ Multiple GPUs
 - ▶ Mixed with multithreading/distributed
- ▶ Result: take existing hybrid ODE simulations written in Julia and automatically GPU+distributed+multithreaded parallelize it for the user!
- ▶ Requires relatively small systems (<50 DEs?)


Stiff ODEs: Fall of the BDF

WHAT'S COMING TO GET GEAR'S METHOD.


Evolution of Gear's Method

41


- ▶ GEAR: Original code. Adaptive order adaptive time via interpolation
 - ▶ Lowers the stability!
- ▶ LSODE series: update of GEAR
 - ▶ Adds rootfinding, Krylov, etc
- ▶ VODE: Variable-coefficient form
 - ▶ No interpolation necessary.
- ▶ CVODE: VODE rewritten in C++
 - ▶ Adds sensitivity analysis




BDF1




BDF2




BDF3



BDF4



BDF5



BDF6

Problems with BDF

BDF is a multistep method

- Needs “Startup Steps”

- Inefficient with events


It is only L-stable up to 2nd order

Has high truncation error coefficients


Implicit

- Requires good step predictors


BDF Integrator Stepping with Events



Rodas5 Integrator Stepping with Events



Orego Benchmarks



Rosenbrock Methods

Aren't new! (ode23s)

Can fix a lot of problems:

- Exploit sparse factorization

- No step predictions required

- Can optimize coefficients to high order

Con: Needs accurate Jacobians

- Answer: AD (or symbolic)

$$Wk_1 = F(y_n)$$

$$Wk_2 = F\left(y_n + \frac{2}{3}hk_1\right) - \frac{4}{3}hdJk_1$$

$$y_{n+1} = y_n + \frac{h}{4}(k_1 + 3k_2)$$

ODE Problems can fall into different classes

Physical Modeling

SecondOrderODEProblem($f, u_0, tspan, p$)

- ▶ $u'' = f(u, p, t)$

PartitionedODEProblem($f_1, f_2, v_0, u_0, tspan, p$)

- ▶ $v' = f_1(t, u)$

- ▶ $u' = f_2(v)$

HamiltonianODEProblem($H, p_0, q_0, tspan, p$)

- ▶ $H(p, q)$

PDE Discretizations

SplitODEProblem($f_1, f_2, u_0, tspan, p$) (IMEX)

- ▶ $u' = f_1(u, p, t) + f_2(u, p, t)$

SemilinearODEProblem($A, f_2, u_0, tspan, p$)

- ▶ $u' = Au + f(u, p, t)$

LocalSemilinearODEProblem($A, f_2, u_0, tspan, p$)

$$u' = Au + f(u, p, t)$$

SDIRK Methods can treat half of the problem as explicit,
decreasing the nonlinear solver cost

Exponential Runge-Kutta

Explicit methods for stiff equations

Small enough: Build matrix exponential

Large enough: Krylov $\exp(t^*A)^*v$

$$U_{ni} = e^{c_i h_n L_n} u_n + h_n \sum_{j=1}^{i-1} a_{ij}(h_n L_n) N_n(U_{nj}),$$
$$u_{n+1} = e^{h_n L_n} u_n + h_n \sum_{i=1}^s b_i(h_n L_n) N_n(U_{ni})$$

Crossover Question: Can we automatically divide equations for IMEX/Multirate methods?

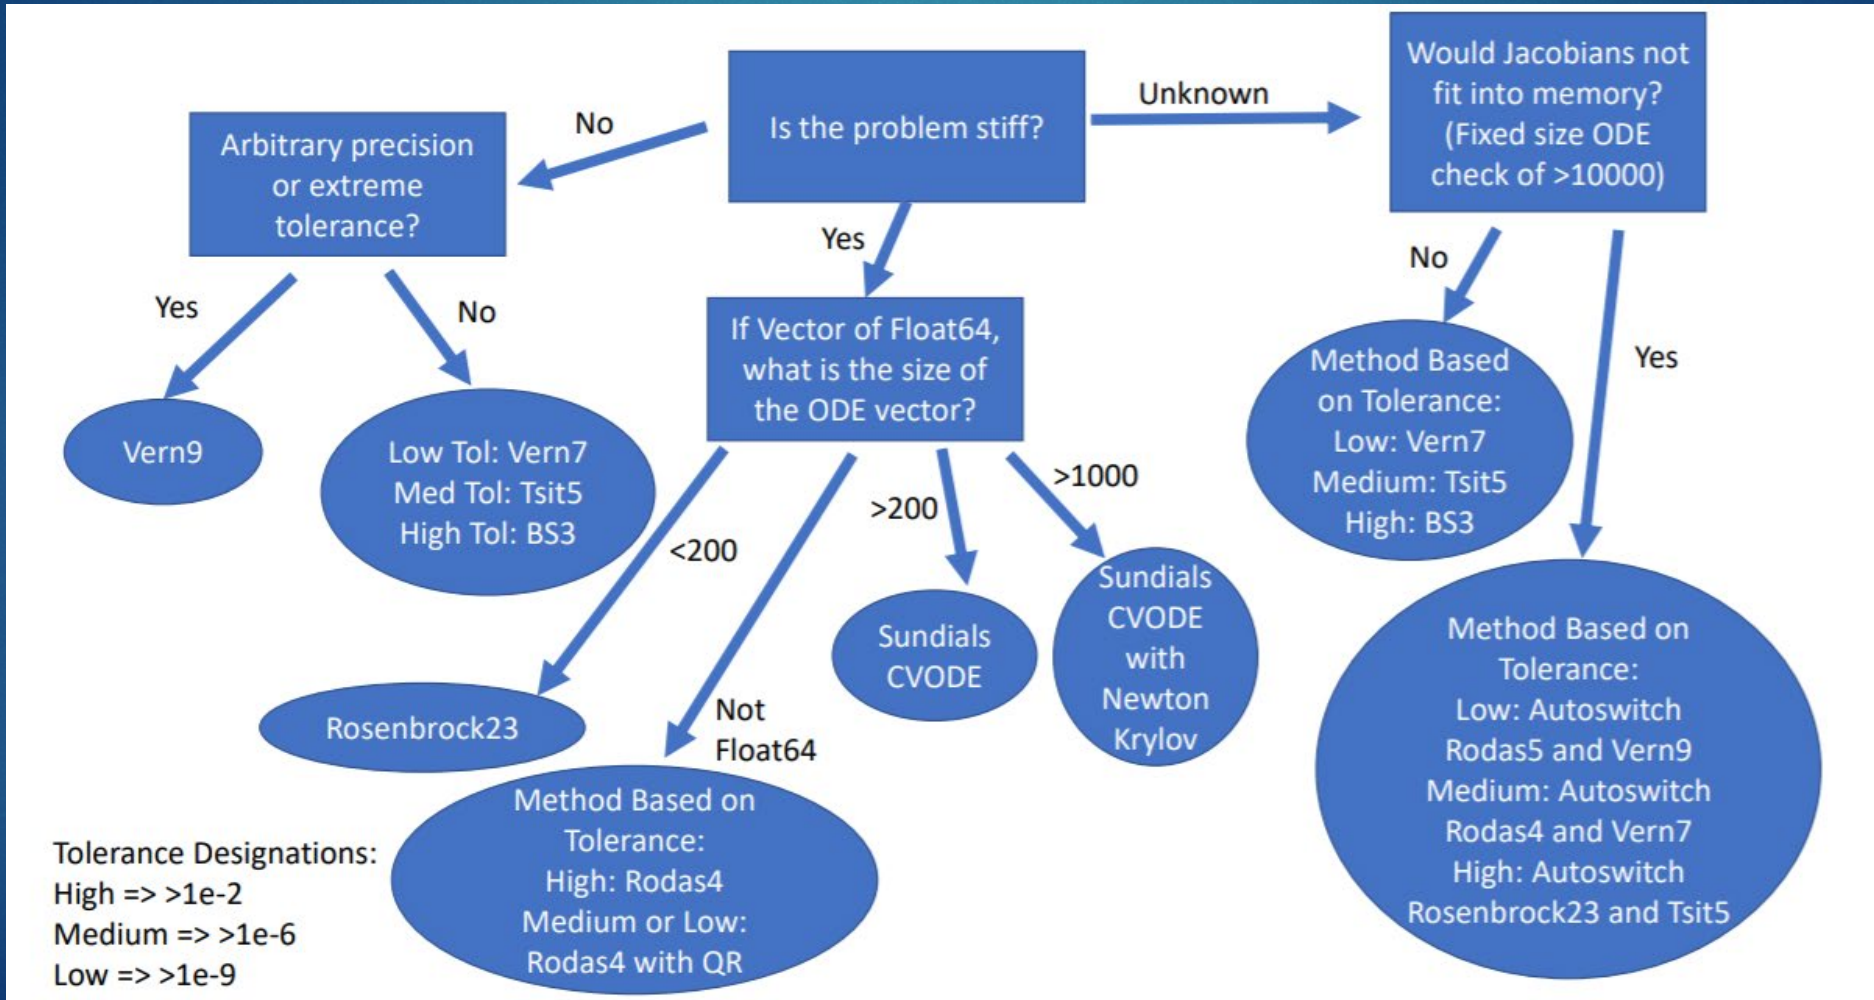
Current Results (DiffEqBenchmarks.jl)

48

- ▶ <20-30 stiff ODEs = High order Rosenbrock
- ▶ <2000 ODEs = Optimized SDIRK Methods
- ▶ >2000 (general) ODEs = SUNDIALS BDF ☹
 - ▶ Stabilized Explicit methods in PDE contexts
 - ▶ IMEX methods when a split is known
 - ▶ ...

Ongoing research to crack the code for more types of systems

Putting it together for users: polyalgorithms



Conclusion

- ▶ Today you can solve differential equations
- ▶ Tomorrow you will likely be able to solve them much faster
 - ▶ Neural-embedded methods for simplifying models
 - ▶ A compiler for researching transformations methods in context
 - ▶ Ongoing improvements to numerical methods

Students, want a paid summer position?

- ▶ Contact me for Google Summer of Code development.
- ▶ No Julia experience is required.
- ▶ <https://julialang.org/soc/ideas-page>

Industry interested in this research?

- ▶ Contact me to help fund development in JuliaDiffEq!
- ▶ We need industry sponsors/interest for CSSI grants, please let us know