

The impact of differentiable programming: how ∂P is enabling new science in Julia

Matt Bauman, Julia Computing









250

-250

-500





• A review of derivatives

the inverse problem

Derivatives

 $julia = 4x^3 - 3x^2 + 2x - 5$ f (generic function with 1 method)



julia> $\partial f(x) = 12x^2 - 6x + 2$ ∂f (generic function with 1 method)

julia> ∂f(2.5) 62.0

Derivatives

Derivatives

How to aim a trebuchet

220 kg counterweight60° release angle100 m to target



5m/s



How to aim a trebuchet

220 kg counterweight55° release angle100 m to target

Distance 0.68m Height 0m Time 0.02s Wind Speed 5mits Target 1.00m bisase Angle 55dep ۲



http://www.virtualtrebuchet.com/

How to simulate a trebuchet Transition 2 Transition 1 Start End Stage 2 Stage 3 Stage 1 Swinging Flying





http://www.virtualtrebuchet.com/



 $r_{1} = GL_{Acg}m_{A}\sin(A_{q}) + Gm_{P}\left(L_{Al}\sin(A_{q}) + L_{S}\sin(A_{q} + S_{q})\right) - Gm_{W}\left(L_{As}\sin(A_{q}) + L_{W}\sin(A_{q} + W_{q})\right)$ $-L_{Al}L_{S}m_{P}\sin(S_{a})\left(A_{w}^{2}-(A_{w}+S_{w})^{2}\right)-L_{As}L_{W}m_{W}\sin(W_{a})\left(A_{w}^{2}-(A_{w}+W_{w})^{2}\right)$







(*∂Trebuchet*) pkg> add Trebuchet

julia> using Trebuchet

julia> Trebuchet.shoot((5, deg2rad(55), 220))[2] 93.53880897080582

http://www.virtualtrebuchet.com/

```
Resolving package versions...
Updating `~/Projects/∂P/∂Trebuchet/Project.toml`
  [98b73d46] + Trebuchet v0.2.1
```

How to $\frac{\partial}{\partial x}$ a trebuchet

julia> shoot([5, 55, 220])

93.53880897080582

julia> shoot'([5, 55, 220])

- 3-element Array{Float64,1}:
 - 4.296867939687003
- -1.4268043075186865 0.07697967188505561

How to $\frac{\partial}{\partial x}$ a trebuchet

julia> @btime shoot(\$[5, 55, 220])
 2.081 ms (68313 allocations: 1.58 MiB)
93.53880897080582

julia> @btime shoot'(\$[5, 55, 220])
 2.551 ms (71902 allocations: 3.70 MiB)
3-element Array{Float64,1}:
 4.296867939687003

-1.4268043075186865 0.07697967188505561



Differential Geometry





Optical Constructions



ш

Compiler Support

How to aim a trebuchet





99.79941172730138

```
# Gradient descent to a single target:
function target!(f, args; N = 10, \eta = 0.1)
        grads = f'(args)
        args .−= η .* grads
```

```
julia> args = target!(x->(shoot([5, x...])-100)^2,
                       [55., 220.])
2-element Array{Float64,1}:
julia> shoot(5, args...)
```



How to quickly aim a trebuchet

julia> aim(5, 100) (49.65935546993644, 211.10317108580074)

julia> shoot(5, 49.66, 211.1) 99.62600675189528



• Once trained, this really is fast:

julia> @btime aim(\$5, \$100) 1.733 µs (18 allocations: 2.58 KiB) (49.65935546993644, 211.10317108580074)

• Even the training is fast (under a minute)

Stror (m) =5

-10



 And fun games can be had with corruption of inputs and controls to add robustness

How to understand a trebuchet

- What if I didn't know one of the terms in our system? \bullet
- But could describe most everything else?



4

$$r_{1} = GL_{Acg}m_{A}\sin(A_{q}) + q$$
$$-L_{Al}L_{S}m_{P}\sin(S_{q})(A$$
$$r_{2} = -L_{W}m_{W}\left(G\sin(A_{q} + S_{q})\right)$$
$$r_{3} = L_{S}m_{P}\left(G\sin(A_{q} + S_{q})\right)$$
$$A_{q}'' = -\frac{r_{1}M_{22}M_{33} - r_{2}M_{33}}{M_{13}M_{22}M_{31} - M_{33}}$$
$$W_{q}'' = \frac{r_{1}M_{21}M_{33} - r_{2}(M_{11}M_{33})}{M_{13}M_{22}M_{31} - M_{33}}$$
$$S_{q}'' = \frac{r_{1}M_{22}M_{31} - r_{2}M_{12}M_{33}}{M_{13}M_{22}M_{31} - M_{33}}$$



A derivative three ways

Derivative with respect to the control parameters



Derivative with respect to an approximation network



Derivative with respect to a subset of the model itself

$$\frac{\sin(A_q + W_q) + L_{As}\sin(W_q)A_w^2}{(A_q + S_q) - \frac{1}{33 - r_2M_{12}M_{33} - r_3M_{13}M_{22}}{r_3 - M_{13}M_{13}M_{22} - M_{12}M_{21})}$$

$$\frac{r_2(M_{11}M_{33} - M_{13}M_{31}) - r_3M_{13}M_{21}}{d_{31} - M_{33}(M_{11}M_{22} - M_{12}M_{21})}$$



error from fixed target





Differentiable Programming is going to disrupt Scientific Modeling and Simulation

Ongoing work by Julia Computing and others in pharmaceuticals, Engineering, Chemistry, Manufacturing, Batteries, Climate, ...



Scientific ML is model-based and data-efficient

Modern experimental procedures quickly captures terabytes of data



How do we simultaneously use both sources of knowledge? **Universal differential equations**



Universal Differential Equations for Scientific Machine Learning arXiv:2001.04385)





Deep Learning discovers systems models from data

Find neural networks so the model matches the data, then find the equations which implies new chemical reactions and pharmacologicallyrelevant systems

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

$$= \left(b[RA_{out}] + \delta[RA - BP] - \left(\gamma[BP] + \eta + \frac{1}{2} + \sigma dW_t + \sigma dW_t$$

Diffusion

 $d[RA - BP] = (\gamma[BP][RA_{in}] + \lambda[BP][RA - RAR] -$ -BP])dt, d[RA - RAR] = \rightarrow] - $\lambda[BP][RA - RAR])dt$, d[RAR] = ($AR] + \lambda [BP] [RA - RAR] - r [RAR]) dt,$

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

Improve QSP and PBPK

Beyond Deterministic Models in Drug Discovery and Development (Trends in Pharmacological Sciences)

DeepNLME in Practice: Automated Discovery of Dynamics

Then predict PK(/PD) profiles of new patients

Or let Pumas discover the mechanistic model:

basis = Basis([u,u^2,u^3,cos(u),sin(u),tan(u),u/(1+u)],[u]) 7 dimensional basis in ["u"] opt = STRRidge(1e-4) STRRidge{Float64}(0.0001) λs = exp10.(-10:0.1:2) 121-element Array{Float64,1}: $\Psi = SInDy(Xmat, Ymat, basis, \lambda s, maxiter = 1000, opt = opt)$ Sparse Identification Result with 1 active terms. latexify.(sum(Ψ.equations.basis)) L"\$p₁ * u\$"

and showcase diagnostics demonstrating the predicted equation is reliable!

pumas^{AI}

Deep NLME

Accelerated building energy efficiency models

Simulations of physical systems involve a combination of compiler passes, ODE / DAE solvers, and machine learning.

Differentiable Programming ties all these parts together.

 Automation of model order reduction on DAEs via neural DAE surrogate dimensional reductions

 Interaction with component-based modeling to allow for generating accelerated building models with transferred learning components

Scientific Machine Learning

https://sciml.ai

Noteworthy new capabilities

- Combine Science and Machine
 Learning
- Comprehensive Differential Equation
 Solvers
- GPU acceleration
- MTK.jl: A DSL for modeling and simulation
- Automatic differentiation

Figure 1: Training performance on the the Robertson equations. We see that our network is able to reproduce the dynamics, capturing the slow (second row) and fast transients (third row) with low error

Time (sec)

Time (sec)

Figure 4: Scaling performance of surrogate on heating system. Our surrogate's runtime grows in near-constant time while the benchmark problem shows exponential growth in runtime.

Optimization of materials for battery-powered aircraft

- GPU accelerate small (30) DAE battery models
- Utilize neural surrogates for global sensitivities
- Automatically refine equations from data
- Use these models to identify material properties
- Propose optimal experimental design
- Closed loop: direct collaboration with materials scientists, restructure model with new data

Automated Climate Parameterizations

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla)\mathbf{u} = -\nabla p + \Pr \nabla^2 \mathbf{u} + b\hat{z}$$
$$\frac{\partial b}{\partial t} + \mathbf{u} \cdot \nabla b = \nabla^2 b + Fe^z$$
Approximate only the vertical flux
$$\frac{\partial}{\partial t} + \mathbf{u} \cdot \nabla - \nabla^2 \left(c' - \frac{\partial}{\partial z} \overline{w'c'} = -w' \frac{\partial \overline{c}}{\partial z} \right)$$

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

High fidelity 15,000x acceleration over direct

3D simulation. Traditionally done by hand!

hand! Capturing missing physics in climate model parameterizations using neural differential equations

Ali Ramadhan, John Marshall, Andre Souza, Gregory LeClaire Wagner, Manvitha Ponnapati, Christopher Rackauckas

Reinforcement Learning with AlphaZero.jl

fish /home/jonathan/AlphaZero.jl 💿 🖲 🥹	julia /home/jonathan/AlphaZero.jl
Starting iteration 1 Starting self-play	Red plays: 1 2 3 4 5 6 7
Time spent on inference: 65% Generating 38 samples per second on average Average exploration depth: 5.3 MCTS memory footprint: 444.63MB Experience buffer size: 103,806 (78,664 distinct boards)	<pre>- · · · · · · · · · · · · · · · · · · ·</pre>
Memory Analysis	Nncts Nnen Vnet
Loss Lv Lp Lreg Linv Hpnet Hp Htot Nb NS 1.8375 0.8683 0.7120 0.1123 0.1448 1.7167 1.0191 84,848 78,664 103,806 all samples 1.8368 0.8681 0.7117 0.1123 0.1450 1.7166 1.0191 84,848 78,664 103,806 latest batch 2.1392 0.9533 0.8035 0.1123 0.2700 1.5119 0.7403 25,614 25,114 25,952 1 to 7 turns left 1.9495 0.9524 0.7274 0.1123 0.1573 1.7300 1.0213 24,237 23,246 25,952 7 to 14 turns left 1.7513 0.8953 0.6778 0.1123 0.0642 1.8387 1.1647 21,446 19,751 25,952 14 to 22 turns left 1.2854 0.6665 0.4950 0.1123 0.0118 1.8884 1.3847 17,139 14,596 25,950 22 to 42 turns left Starting learning Optimizing the loss Loss Lv Lp Lreg Linv Hp Hpnet 2.4152 1.0240 1.1380 0.1123 0.1411 1.0081 1.3999 1.4595 0.6575 0.6532 0.1122 0.0367 1.0081 1.3999	1,132 10 +0.00 +0.25 Pncts Pnet UCT Pnen Qncts Qnet 4 97.4% 75.7% 69.9% +0.24 41.7% +0.19 +0.91 2 1.8% 10.2% 17.4% +0.15 35.0% +0.05 -0.53 5 0.4% 4.6% 5.0% +0.12 20.4% +0.06 +0.25 6 0.3% 4.4% 2.8% +0.19 1.4% +0.15 +0.12 3 0.1% 2.8% 1.4% +0.21 1.2% +0.18 -0.12 1 0.0% 1.6% 2.3% +0.08 0.1% -0.00 -0.03 7 0.0% 0.7% 1.3% -0.08 0.4% -0.18 -0.29 > do 4
Percentage of Won Games	Percentage of Won Games
80-60-	80- 60-
40 - 20 - AlphaZero / MCTS (1000 rollouts)	40 - 20 - Network Only / N
AlphaZero / MirMax (depth 5)	0 0 5 10 15 Iteration number
Jonathar	n Laurent

https://github.com/jonathan-laurent/

<u>AlphaZero.jl</u>

- **Simple**: Core algorithm is 2000 lines of Julia
- **Extensible**: Generic interfaces allow new games and learning

J	7
CTS (1000 rollo InMax (depth 5)	JRS)
20	8

- Fast: 10-100x faster than other high-level alternatives
- Scalable: Combines distributed, multithreaded and multi-GPU parallelism

Celeste was our first peta-scale Bayesian Inference application Turing.jl makes Probabilistic Programming available more broadly

	Π

Cataloging the Visible Universe through Bayesian Inference at Petascale

Jeffrey Regier*, Kiran Pamnany[†], Keno Fischer[‡], Andreas Noack[§], Maximilian Lam^{*}, Jarrett Revels[§], Steve Howard[¶], Ryan Giordano[¶], David Schlegel^{||}, Jon McAuliffe[¶], Rollin Thomas^{||}, Prabhat^{||}

Thank You

Most light sources are near the detection limit

ullet

lacksquare

Electrification, batteries, climate change and the pandemic are reshaping entire industries

Differentiable Programming is at the heart of improving the pace of innovation

In Julia, ∂P is being driven by real-world applications in modeling and simulation

These advances are accelerating product design for engineers and scientists

In short, Julia is building a language ecosystem and compiler toolchain for modern science