




Solving the two language problem in numerical computing

Viral B. Shah, Jeff Bezanson, Stefan Karpinski, and Alan Edelman

James H. Wilkinson Prize for Numerical Software Lecture





● C++ 47.3% ● Python 41.1% ● HTML 5.9% ● Jupyter Notebook 2.4% ● Go 1.3% ● Java 0.7% ● Other 1.3%



PYTORCH

● Python 32.1% ● C++ 29.6% ● Cuda 18.0% ● C 15.6% ● CMake 3.4% ● Fortran 0.6% ● Other 0.7%



● C++ 80.1% ● Python 9.1% ● Cuda 5.9% ● CMake 2.8% ● Matlab 0.9% ● Makefile 0.7% ● Shell 0.5%



● C++ 54.7% ● Jupyter Notebook 25.9% ● Python 11.5% ● Cuda 4.1% ● C# 1.1% ● Shell 0.9% ● Other 1.8%



Knet.jl

● Julia 84.5% ● Cuda 9.8% ● C 1.8% ● Makefile 1.7% ● Matlab 1.7% ● Python 0.4% ● Other 0.1%



● Julia 100.0%



Demmel, Dongarra et al. on Polymorphism

The Sca/LAPACK software engineering problem is to express the following metaprogram:


- (1) for all linear algebra problems (linear systems, eigenproblems, ...)
- (2) for all matrix types (general, symmetric, banded, ...)
- (3) for all data types (real, complex, single, double, higher precision)
- (4) for all machine architectures
 and communication topologies
- (5) for all programming interfaces
- (6) provide the best algorithm(s) available
 best: performance or accuracy





Celeste.jl: Julia at Peta-scale

Cori: 650,000 cores. 1.3M threads. 60 TB of data.



Most light sources are near the detection limit.

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^{||}

^{*}*Department of Electrical Engineering and Computer Sciences, University of California, Berkeley*

[†]*Parallel Computing Lab, Intel Corporation*

[‡]*Julia Computing*

[§]*Computer Science and AI Laboratories, Massachusetts Institute of Technology*

[¶]*Department of Statistics, University of California, Berkeley*


^{||}*Lawrence Berkeley National Laboratory*



Celeste: Custom sparsity patterns and storage



Matrix structure





Storage

Julia runs on Google TPUs



Performance on par with TensorFlow on TPUs

Scales to pods (512 TPU cores - 4.3 PF₁₆/s on ResNet50)



Zygote.jl: General Purpose Automatic Differentiation

```
function foo(W, Y, x)
  Z = W * Y
  a = Z * x
  b = Y * x
  c = tanh.(b)
  r = a + c
  return r
end
```



```
function ∇foo(W, Y, x)
  Z = W * Y
  a = Z * x
  b = Y * x
  c, Jtanh = ∇tanh.(b)
  a + c, function (Δr)
    Δc = Δr, Δa = Δr
    (Δtanh, Δb) = Jtanh(Δc)
    (ΔY, Δx) = (Δb * x', Y' * Δb)
    (ΔZ = Δa * x', Δx += Z' * Δa)
    (ΔW = ΔZ * Y', ΔY = W' * ΔZ')
  (nothing, ΔW, ΔY, Δx)
end
end
```



M. Innes. Don't Unroll Adjoint:
Differentiating SSA-Form Programs
([arXiv:1810.07951](https://arxiv.org/abs/1810.07951))

Composability: DifferentialEquations.jl + Measurements.jl


```
g = 9.79 ± 0.02 # Gravitational constants
L = 1.00 ± 0.01 # Length of the pendulum

# Initial speed & angle, time span
u₀ = [0 ± 0, π/60 ± 0.01]
tspan = (0.0, 6.3)

# Define the problem
function pendulum(du, u, p, t)
    θ = u[1]
    dθ = u[2]
    du[1] = dθ
    du[2] = -(g/L)*θ
end

# Pass to solvers
prob = ODEProblem(pendulum, u₀, tspan)
sol = solve(prob, Tsit5(), reltol = 1e-6)

# Analytic solution
u = u₀[2] .* cos.(sqrt(g/L) .* sol.t)
```

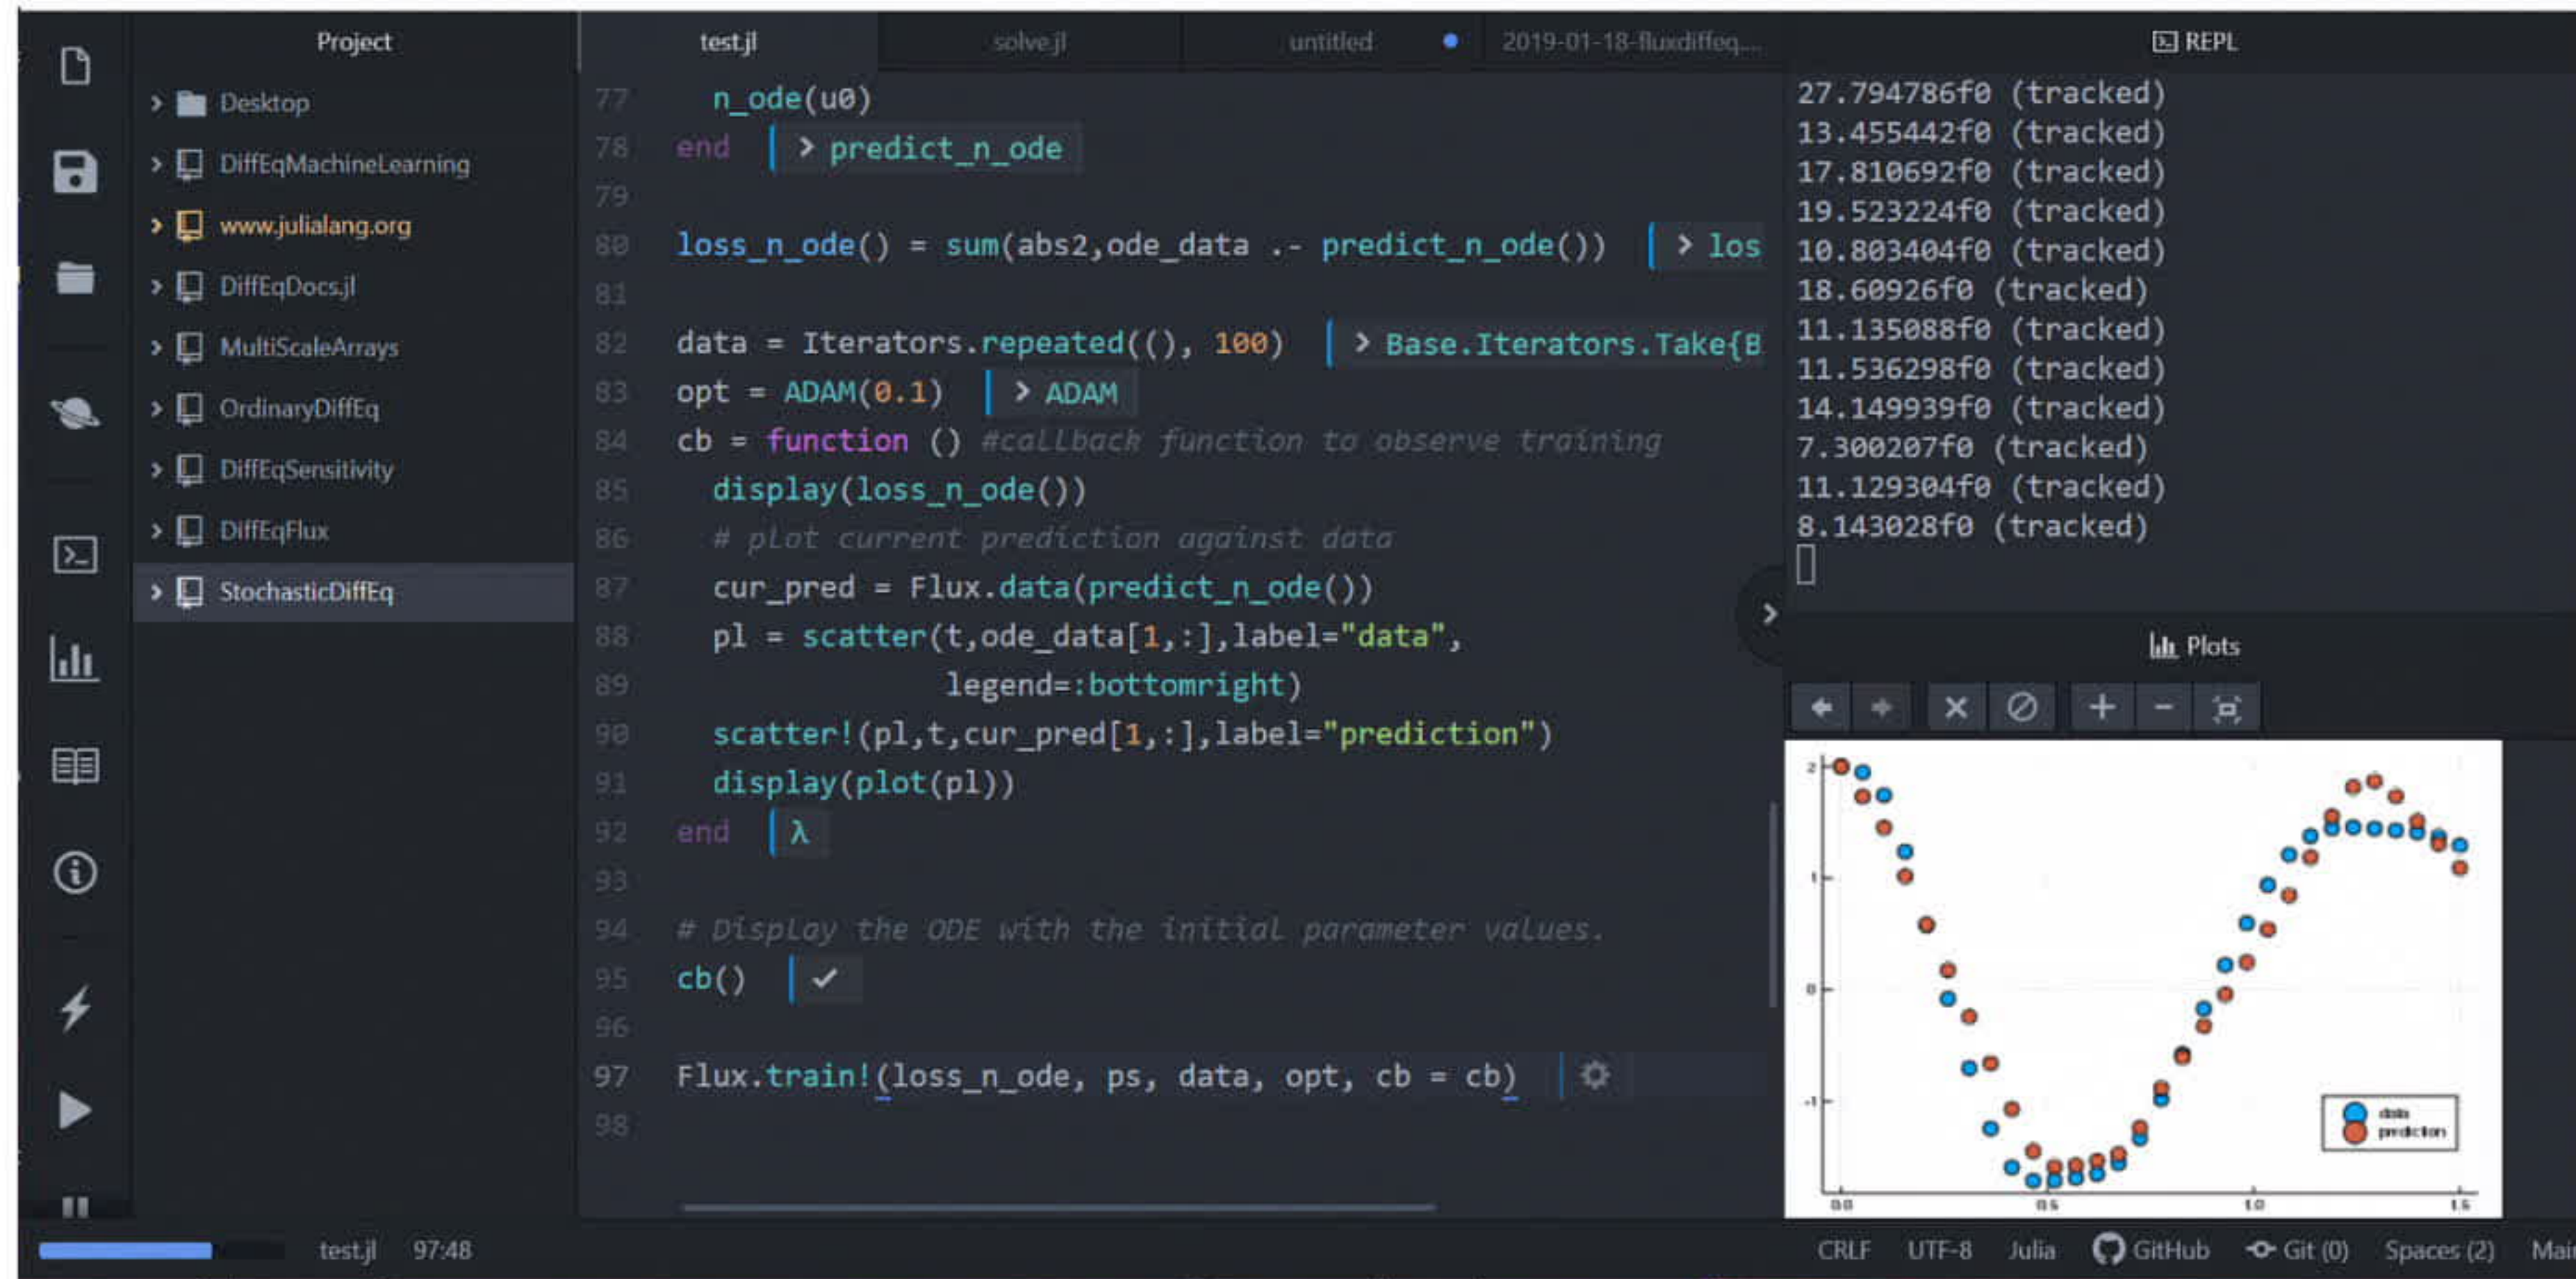


Rackauckas et al. *DifferentialEquations.jl – A Performant and Feature-Rich Ecosystem for Solving Differential Equations in Julia*. 2017. ([Journal of Open Research Software](#))



Giordano. *Uncertainty propagation with functionally correlated quantities* ([arXiv:1610.08716](#))

Composability: DifferentialEquations.jl + Flux.jl (Neural ODEs)



The screenshot displays a Julia IDE interface with a project explorer on the left, a code editor in the center, and a REPL window on the right. The code in the editor defines a neural ODE model and its training process. The REPL window shows the output of the training process, including a list of tracked loss values. A plot window in the bottom right corner shows the training data (blue dots) and the model's predictions (red dots), which form a smooth curve fitting the data points.

```
Project
├── Desktop
├── DiffEqMachineLearning
├── www.julia.org
├── DiffEqDocs.jl
├── MultiScaleArrays
├── OrdinaryDiffEq
├── DiffEqSensitivity
├── DiffEqFlux
└── StochasticDiffEq

test.jl
77 n_ode(u0)
78 end | > predict_n_ode
79
80 loss_n_ode() = sum(abs2,ode_data .- predict_n_ode()) | > los
81
82 data = Iterators.repeated((), 100) | > Base.Iterators.Take{B
83 opt = ADAM(0.1) | > ADAM
84 cb = function () #callback function to observe training
85     display(loss_n_ode())
86     # plot current prediction against data
87     cur_pred = Flux.data(predict_n_ode())
88     pl = scatter(t,ode_data[1,:],label="data",
89                 legend=:bottomright)
90     scatter!(pl,t,cur_pred[1,:],label="prediction")
91     display(plot(pl))
92 end | λ
93
94 # Display the ODE with the initial parameter values.
95 cb() | ✓
96
97 Flux.train!(loss_n_ode, ps, data, opt, cb = cb) | ⚙
98
```

27.794786f0 (tracked)
13.455442f0 (tracked)
17.810692f0 (tracked)
19.523224f0 (tracked)
10.803404f0 (tracked)
18.60926f0 (tracked)
11.135088f0 (tracked)
11.536298f0 (tracked)
14.149939f0 (tracked)
7.300207f0 (tracked)
11.129304f0 (tracked)
8.143028f0 (tracked)
[]

Plots

Legend: data (blue dots), predictors (red dots)

test.jl 97:48 CRLF UTF-8 Julia GitHub Git (0) Spaces (2) Main



Rackauckas et al. *DiffEqFlux.jl - A Julia Library for Neural Differential Equations*
([arXiv:1902.02376](https://arxiv.org/abs/1902.02376))

Other special matrix types

- Diagonal
- UniformScaling
- Symmetric, Hermitian
- LowerTriangular, UpperTriangular
- Bidiagonal, Tridiagonal, SymTridiagonal
- Adjoint, Transpose

Recent work: type system formalization



Julia Subtyping: a Rational Reconstruction

F. Zappa Nardelli, J. Belyakova, A. Pelenitsyn, B. Chung, J. Bezanson, J. Vitek

[OOPSLA 2018](#)

Paper: <https://www.di.ens.fr/~zappa/projects/lambdajulia/>

$$\begin{array}{c}
\text{[TOP]} \\
\hline
E \vdash t <: \text{Any} \vdash E
\end{array}
\qquad
\begin{array}{c}
\text{[REFL]} \\
\hline
E \vdash a <: a \vdash E
\end{array}
\qquad
\begin{array}{c}
\text{[TUPLE]} \\
E \vdash a_1 <: a'_1 \vdash E_1 \quad \dots \quad E_{n-1} \vdash a_n <: a'_n \vdash E_n \\
\text{consistent}(E_n) \\
\hline
E \vdash \text{Tuple}\{a_1, \dots, a_n\} <: \text{Tuple}\{a'_1, \dots, a'_n\} \vdash E_n
\end{array}$$

$$\begin{array}{c}
\text{[TUPLE_LIFT_UNION]} \\
t' = \text{lift_union}(\text{Tuple}\{a_1, \dots, a_n\}) \\
E \vdash t' <: t \vdash E' \\
\hline
E \vdash \text{Tuple}\{a_1, \dots, a_n\} <: t \vdash E'
\end{array}
\qquad
\begin{array}{c}
\text{[TUPLE_UNLIFT_UNION]} \\
t' = \text{unlift_union}(\text{Union}\{t_1, \dots, t_n\}) \\
E \vdash t <: t' \vdash E' \\
\hline
E \vdash t <: \text{Union}\{t_1, \dots, t_n\} \vdash E'
\end{array}$$

$$\begin{array}{c}
\text{[UNION_LEFT]} \\
E \vdash t_1 <: t \vdash E_1 \quad \dots \quad \text{reset_occ}_E(E_{n-1}) \vdash t_n <: t \vdash E_n \\
\hline
E \vdash \text{Union}\{t_1, \dots, t_n\} <: t \vdash \text{max_occ}_{E_1 \dots E_n}(E_n)
\end{array}
\qquad
\begin{array}{c}
\text{[UNION_RIGHT]} \\
\exists j. E \vdash t <: t_j \vdash E' \\
\hline
E \vdash t <: \text{Union}\{t_1, \dots, t_n\} \vdash E'
\end{array}$$

$$\begin{array}{c}
\text{[APP_INV]} \\
n \leq m \quad E_0 = \text{add}(\text{Barrier}, E) \\
\forall 0 < i \leq n. E_{i-1} \vdash a_i <: a'_i \vdash E'_i \quad \wedge \quad E'_i \vdash a'_i <: a_i \vdash E_i \\
\hline
E \vdash \text{name}\{a_1, \dots, a_m\} <: \text{name}\{a'_1, \dots, a'_n\} \vdash \text{del}(\text{Barrier}, E_n)
\end{array}
\qquad
\begin{array}{c}
\text{[APP_SUPER]} \\
\text{name}\{\top_1, \dots, \top_m, \dots\} <: t'' \in \text{tds} \\
E \vdash t''[a_1/\top_1 \dots a_m/\top_m] <: t' \vdash E' \\
\hline
E \vdash \text{name}\{a_1, \dots, a_m\} <: t' \vdash E'
\end{array}$$

$$\begin{array}{c}
\text{[L_INTRO]} \\
\hline
\text{add}({}^L\top_{t_1}^{t_2}, E) \vdash t <: t' \vdash E' \\
\hline
E \vdash t \text{ where } t_1 <: \top <: t_2 <: t' \vdash \text{del}(\top, E')
\end{array}
\qquad
\begin{array}{c}
\text{[R_INTRO]} \\
\hline
\text{add}({}^R\top_{t_1}^{t_2}, E) \vdash t <: t' \vdash E' \\
\text{consistent}(\top, E') \\
\hline
E \vdash t <: t' \text{ where } t_1 <: \top <: t_2 \vdash \text{del}(\top, E')
\end{array}$$

[R_L]

[L_LEFT]

$$\frac{\begin{array}{l} search(\top, E) = {}^L\top_l^u \\ E \vdash u <: t \vdash E' \end{array}}{E \vdash \top <: t \vdash E'}$$

[L_RIGHT]

$$\frac{\begin{array}{l} search(\top, E) = {}^L\top_l^u \\ E \vdash t <: l \vdash E' \end{array}}{E \vdash t <: \top \vdash E'}$$

$$\frac{\begin{array}{l} search(\top_1, E) = {}^R\top_1_{l_1}^{u_1} \quad search(\top_2, E) = {}^L\top_2_{l_2}^{u_2} \\ outside(\top_1, \top_2, E) \Rightarrow E \vdash u_2 <: l_2 \vdash E' \\ E \vdash u_1 <: l_2 \vdash E'' \end{array}}{E \vdash \top_1 <: \top_2 \vdash upd({}^R\top_1_{l_1}^{u_1} \text{Union}\{\top_1, l_1\}, E')}$$

[R_RIGHT]

[R_LEFT]

$$\frac{\begin{array}{l} search(\top, E) = {}^R\top_l^u \\ E \vdash l <: t \vdash E' \end{array}}{E \vdash \top <: t \vdash upd({}^R\top_l^u, E')}$$

$$\frac{\begin{array}{l} search(\top, E) = {}^R\top_l^u \\ (is_var(t) \wedge search(t, E) = {}^L S_{l_1}^{u_1}) \Rightarrow \neg outside(\top, S, E) \\ E \vdash t <: u \vdash E' \end{array}}{E \vdash t <: \top \vdash upd({}^R\top_l^u \text{Union}\{l, t\}, E')}$$

[TYPE_LEFT]

$$\frac{\begin{array}{l} \neg is_var(a_1) \\ E \vdash typeof(a_1) <: t_2 \vdash E' \end{array}}{E \vdash \text{Type}\{a_1\} <: t_2 \vdash E'}$$

[TYPE_RIGHT]

$$\frac{\begin{array}{l} is_kind(t_1) \quad is_var(t_2) \\ E \vdash \text{Type}\{\top\} \text{ where } \top <: \text{Type}\{t_2\} \vdash E' \end{array}}{E \vdash t_1 <: \text{Type}\{t_2\} \vdash E'}$$

[TYPE_TYPE]

$$\frac{add(\text{Barrier}, E) \vdash a_1 <: a_2 \vdash E' \quad E' \vdash a_2 <: a_1 \vdash E''}{E \vdash \text{Type}\{a_1\} <: \text{Type}\{a_2\} \vdash del(\text{Barrier}, E')}$$

Reproducibility by default: Manifest.toml files

```
# This file is machine-generated - editing it directly is not advised
```

```
[[UnicodePlots]]
```

```
deps = ["Dates", "Random", "SparseArrays", "StatsBase", "Test"]
```

```
git-tree-sha1 = "b8e58d4390ccebfa4f3bf502a45e08066eec3bf9"
```

```
uuid = "b8865327-cd53-5732-bb35-84acbb429228"
```

```
version = "1.1.0"
```

```
[[OrderedCollections]]
```

```
deps = ["Random", "Serialization", "Test"]
```

```
git-tree-sha1 = "85619a3f3e17bb4761fe1b1fd47f0e979f964d5b"
```

```
uuid = "bac558e1-5e72-5ebc-8fee-abe8a469f55d"
```

```
version = "1.0.2"
```

```
[[SortingAlgorithms]]
```

```
deps = ["DataStructures", "Random", "Test"]
```

```
git-tree-sha1 = "03f5898c9959f8115e30bc7226ada7d0df554ddd"
```

```
uuid = "a2af1166-a08f-5f64-846c-94a0d3cef48c"
```

```
version = "0.3.1"
```

Fix deprecations #111



Merged

ararslan merged 2 commits into `master` from `fbot/deps` on Sep 25, 2018

Conversation 4

Commits 2

Checks 1

Files changed 1




femtocleaner bot commented on Aug 8, 2018

Contributor

+ 😊 ...

I fixed a number of deprecations for you

Reviewers

 ararslan

Assignees

No one—assign

Labels

None yet

Projects

None yet



femtocleaner bot commented on Sep 14, 2018

Author

Contributor

+ 😊 ...

My code has been updated. I now view the world differently.

Am I still the same bot I was before?

In any case, I've updated this PR to reflect my new knowledge. I hope you like it.

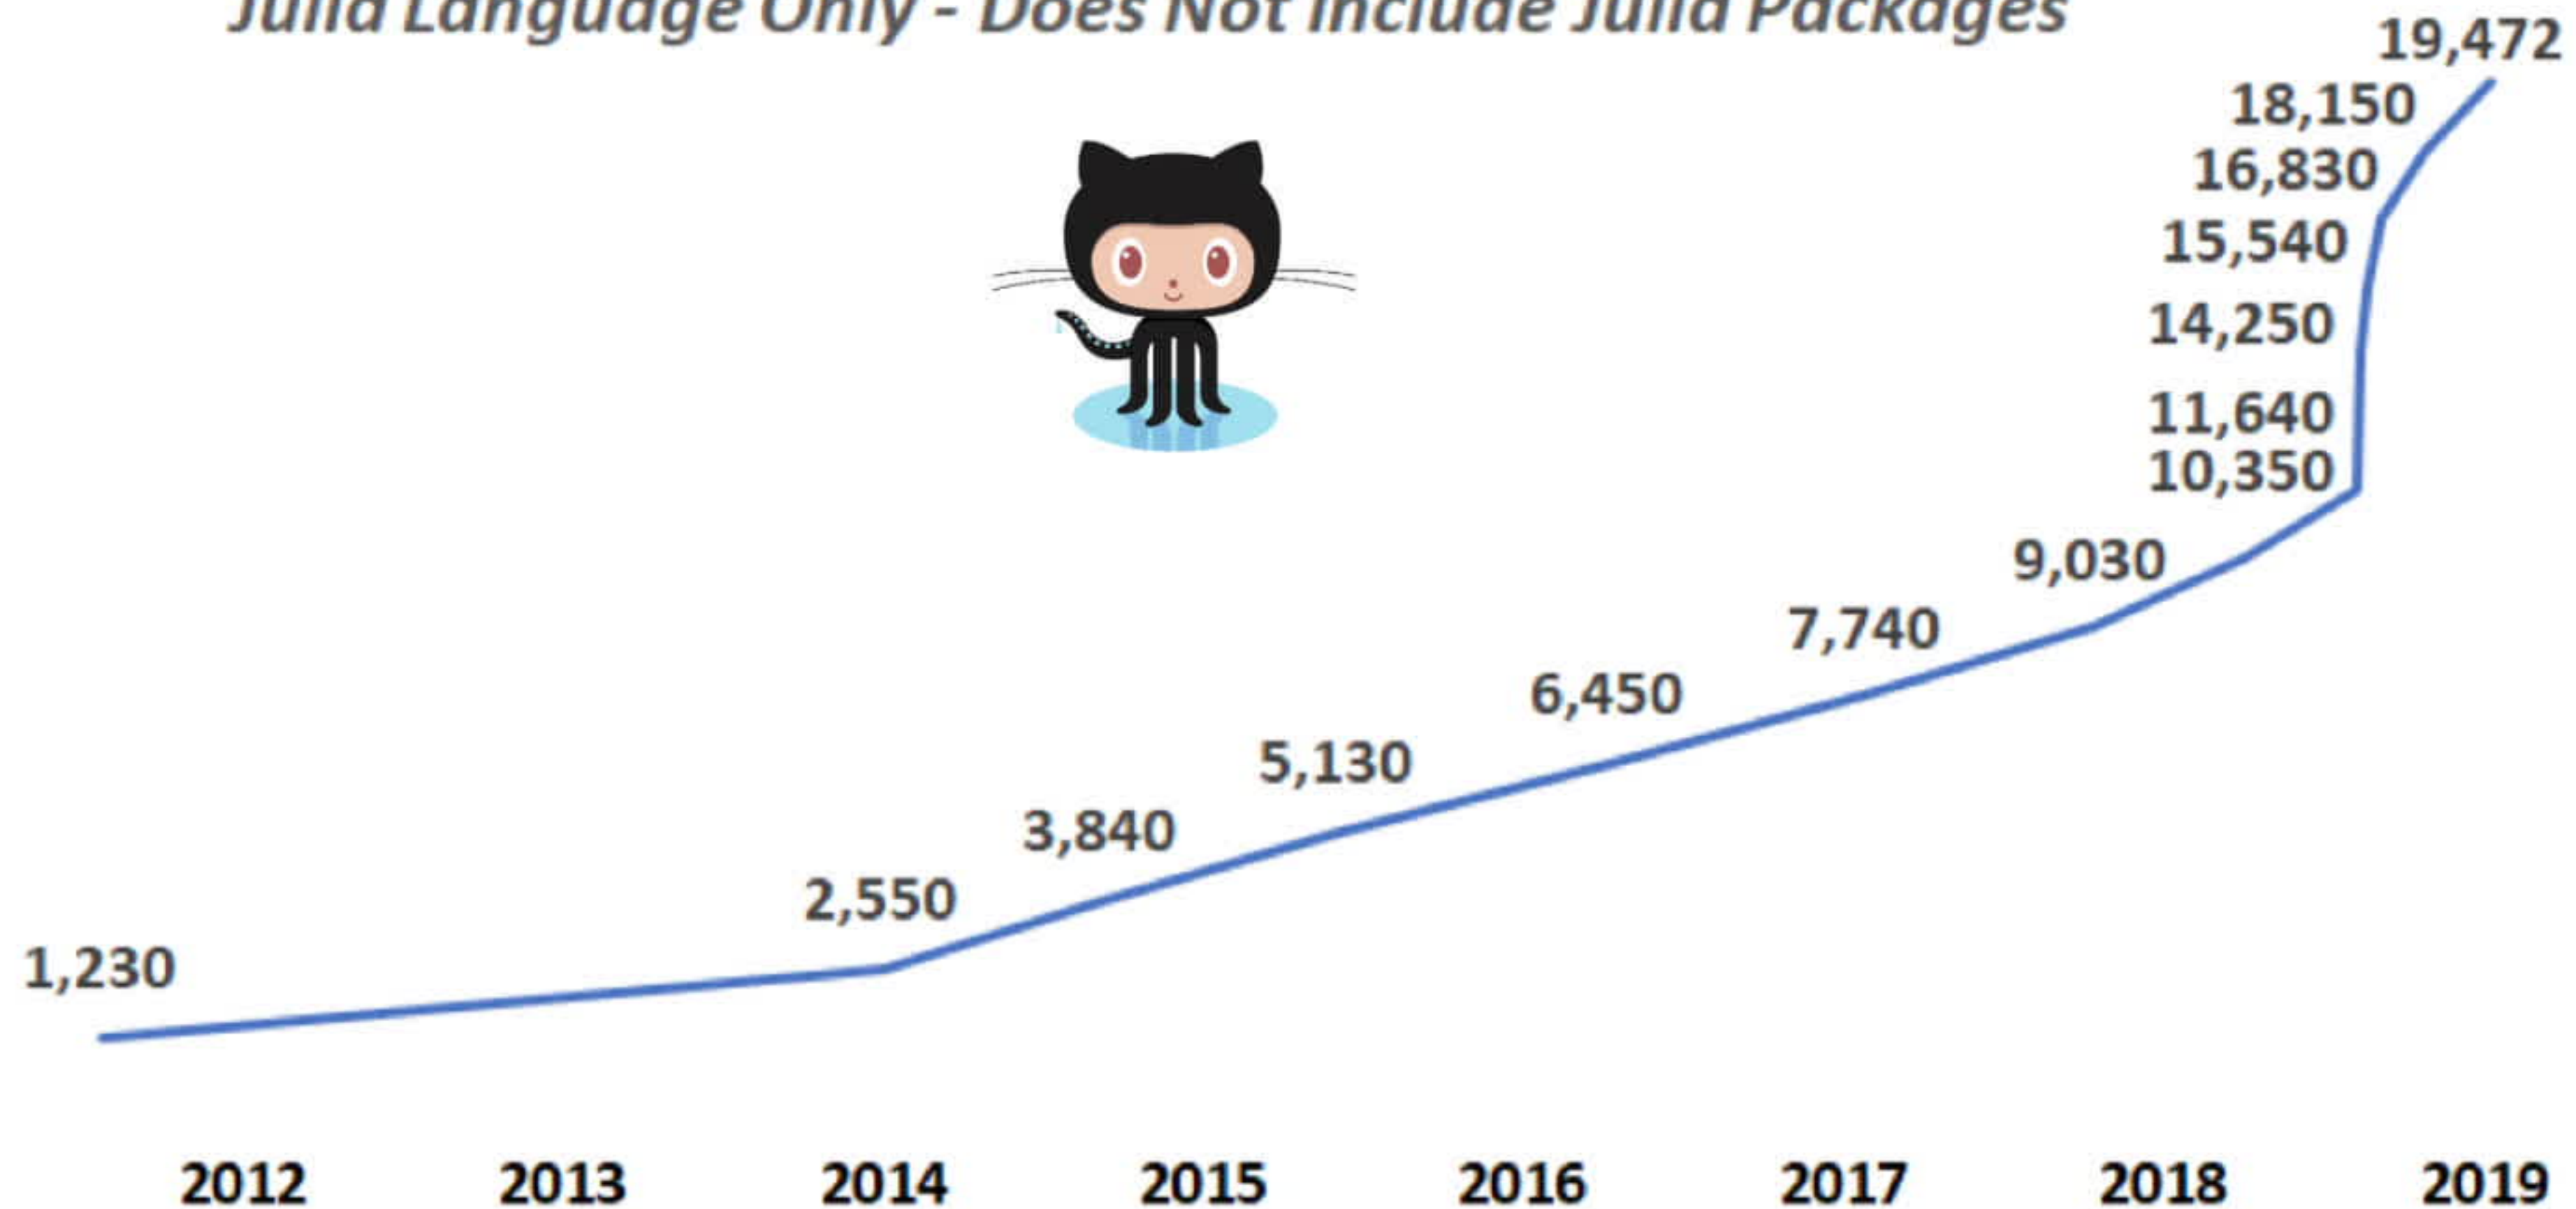


femtocleaner bot force-pushed the `fbot/deps` branch from `879a084` to `44c9b89` on Sep 14, 2018








A Growing Community

Julia GitHub Stars

Julia Language Only - Does Not Include Julia Packages



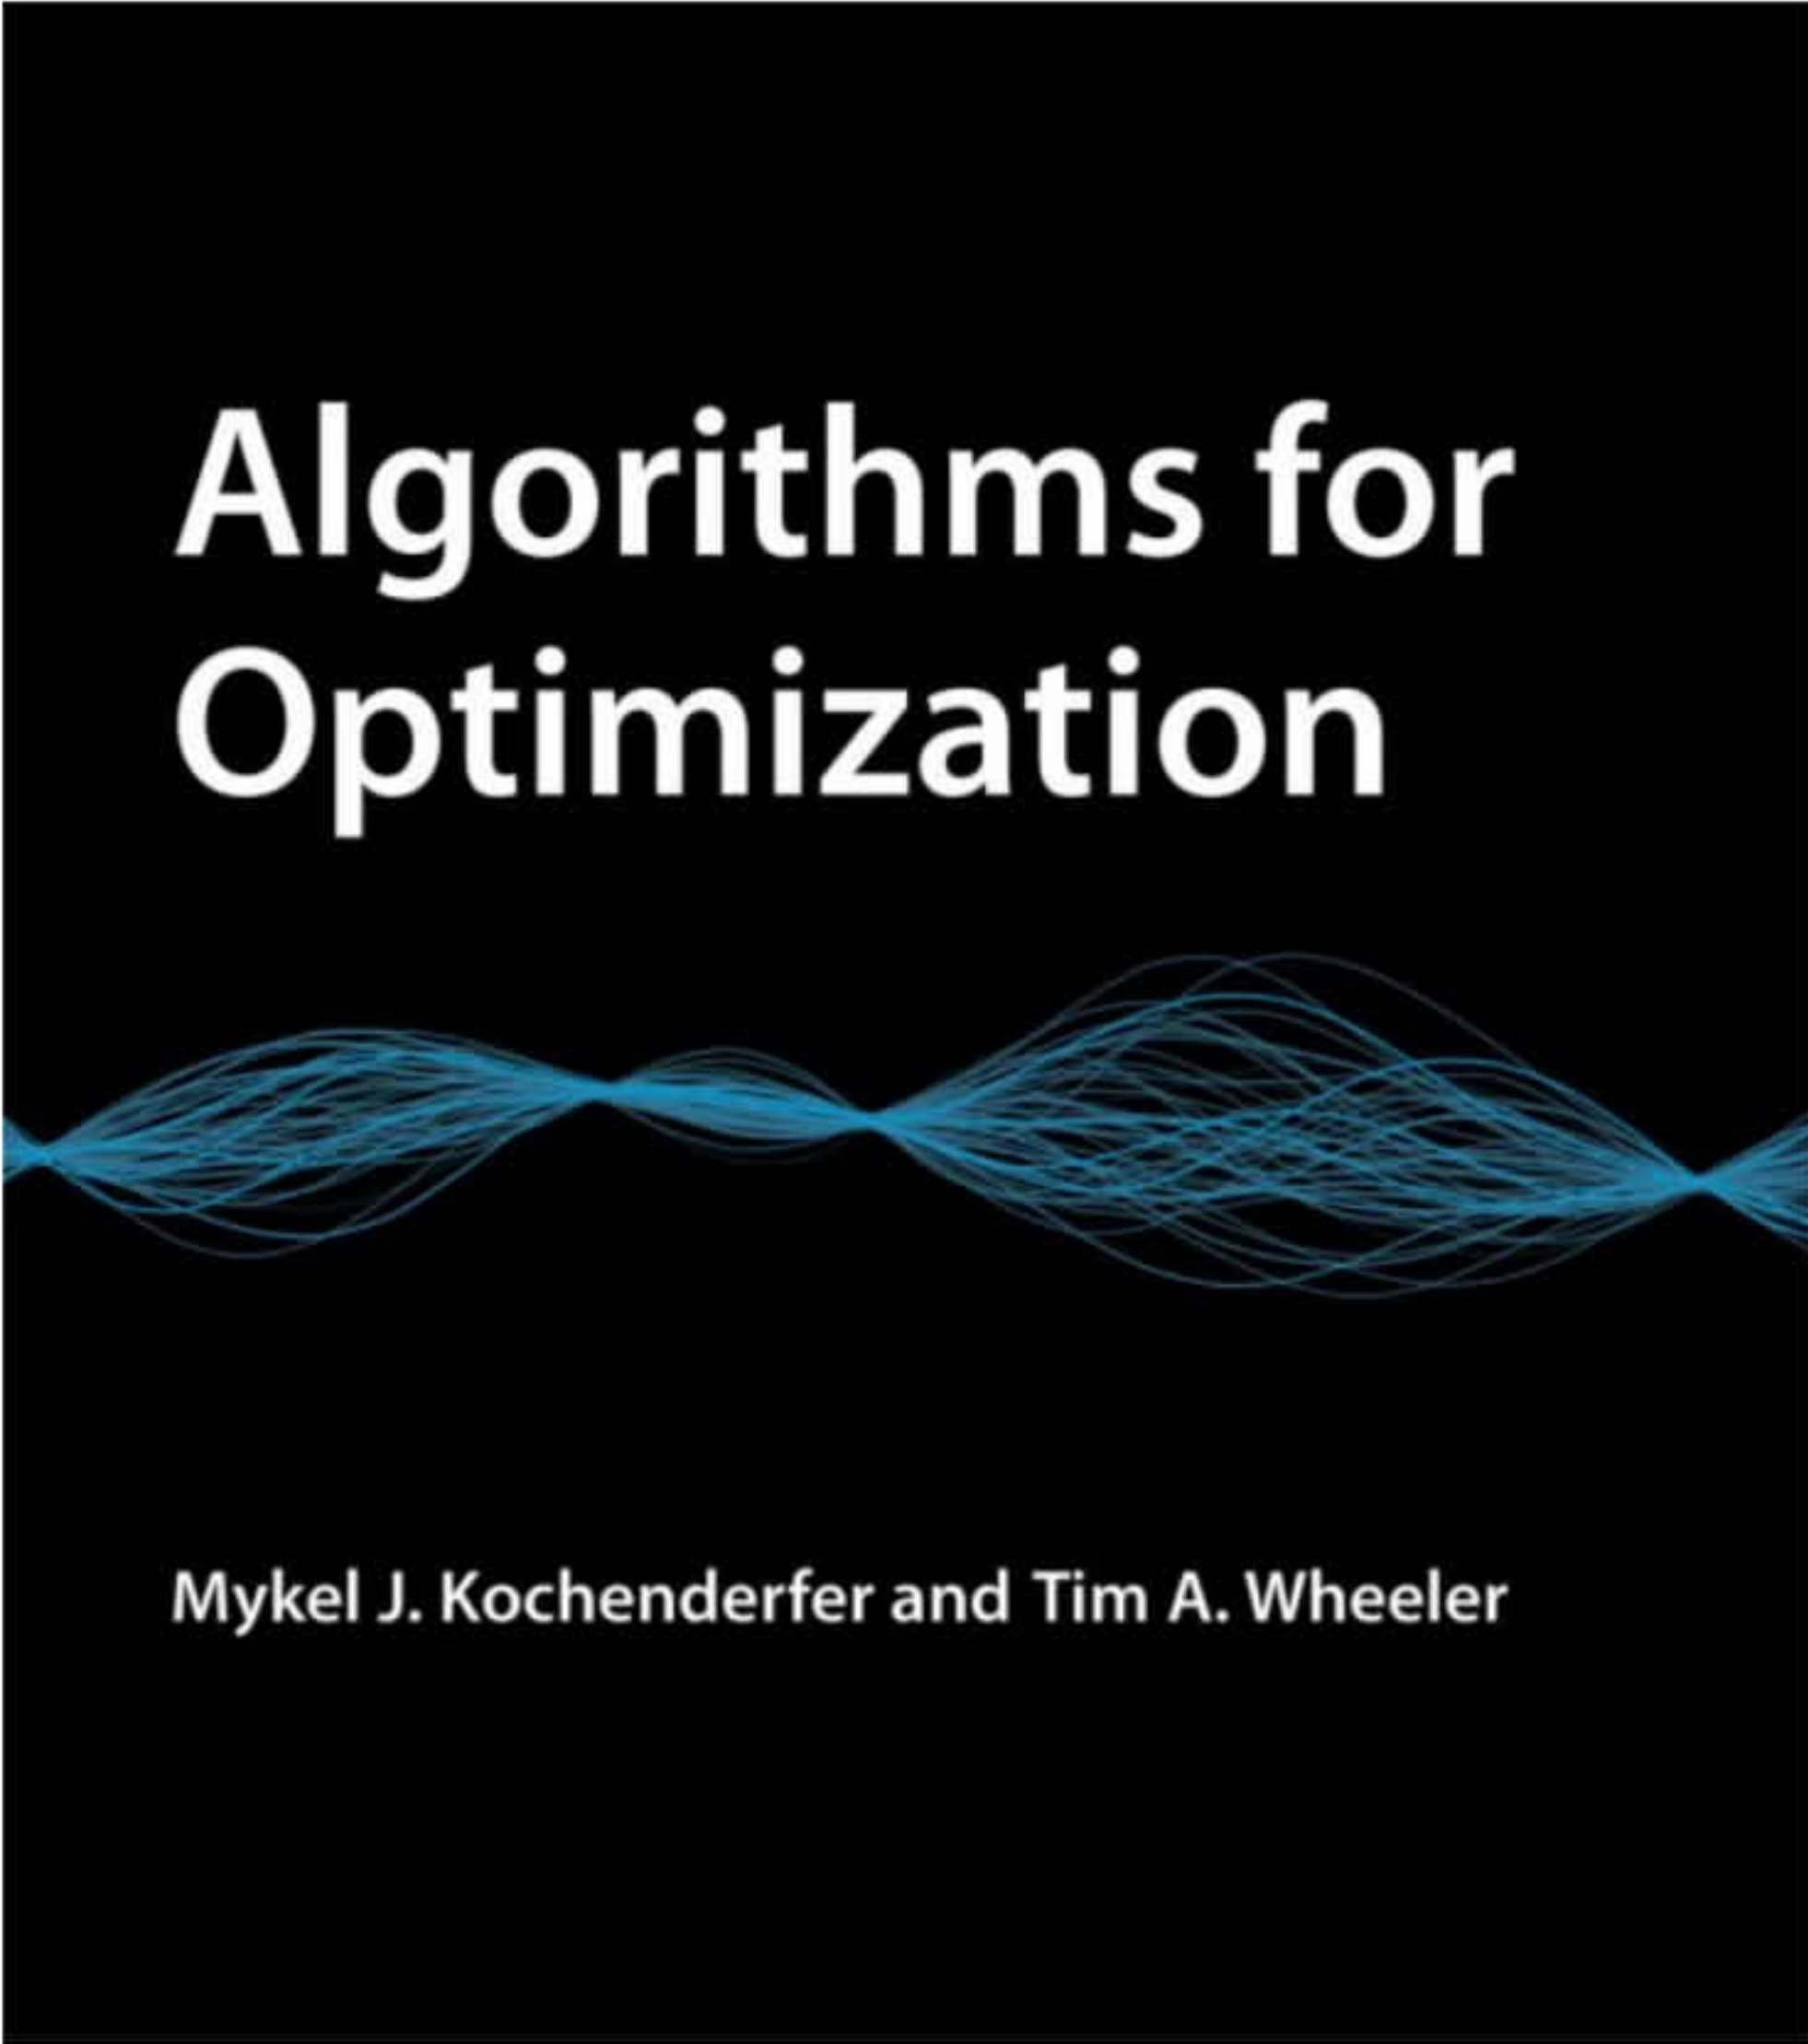
Books



Books

Algorithms for Optimization

Mykel J. Kochenderfer and Tim A. Wheeler






Figure 4.1. The sufficient decrease condition, the first Wolfe condition, can always be satisfied by a sufficiently small step size along a descent direction.

If \mathbf{d} is a valid descent direction, then there must exist a sufficiently small step size that satisfies the sufficient decrease condition. We can thus start with a large step size and decrease it by a constant reduction factor until the sufficient decrease condition is satisfied. This algorithm is known as *backtracking line search*⁶ because of how it backtracks along the descent direction. Backtracking line search is shown in figure 4.2 and implemented in algorithm 4.2. We walk through the procedure in example 4.2.

⁶ Also known as *Armijo line search*, L. Armijo, "Minimization of Functions Having Lipschitz Continuous First Partial Derivatives," *Pacific Journal of Mathematics*, vol. 16, no. 1, pp. 1–3, 1966.

```

function backtracking_line_search(f, ∇f, x, d, α; p=0.5, β=1e-4)
    y, g = f(x), ∇f(x)
    while f(x + α*d) > y + β*α*(g·d)
        α *= p
    end
    α
end

```

Algorithm 4.2. The backtracking line search algorithm, which takes objective function f , its gradient ∇f , the current design point \mathbf{x} , a descent direction \mathbf{d} , and the maximum step size α . We can optionally specify the reduction factor p and the first Wolfe condition parameter β .

```

function direct(f, a, b,  $\epsilon$ , k_max)
  g = reparameterize_to_unit_hypercube(f, a, b)
  intervals = Intervals()
  n = length(a)
  c = fill(0.5, n)
  interval = Interval(c, g(c), fill(0, n))
  add_interval!(intervals, interval)
  c_best, y_best = copy(interval.c), interval.y

  for k in 1 : k_max
    S = get_opt_intervals(intervals,  $\epsilon$ , y_best)
    to_add = Interval[]
    for interval in S
      append!(to_add, divide(g, interval))
      dequeue!(intervals[min_depth(interval)])
    end
    for interval in to_add
      add_interval!(intervals, interval)
      if interval.y < y_best
        c_best, y_best = copy(interval.c), interval.y
      end
    end
  end

  return rev_unit_hypercube_parameterization(c_best, a, b)
end

```

Algorithm 7.8. DIRECT, which takes the multidimensional objective function f , vector of lower bounds a , vector of upper bounds b , tolerance parameter ϵ , and number of iterations k_{\max} . It returns the best coordinate.





Figure 7.20. The DIRECT method after 16 iterations on the Branin function, appendix B.3. Each cell is bordered by white lines. The cells are much denser around the minima of the Branin function, as the DIRECT method procedurally increases its resolution in those regions.

Some of the Universities Teaching Julia





“ *A main goal in designing a language should be to plan for growth. The language must start small, and the language must grow as the set of users grows.* ”

Guy Steele, “Growing a language”, 1998