# Matnix Calculus lecture notes: How can we use so many derivatives? <br> ... a couple of applications ... and the "adjoint method" 

Matrix Calculus, IAP 2023<br>Profs. Steven G. Johnson \& Alan Edelman, MIT

## Newton's method: Nonlinear equations via Linearization



## Multidimensional Newton's method: Real world is nonlinear!

vector out
18.06: solving $f(x)=0$ where $x \in \mathbb{R}^{n}$ (input=vector) and $f$ and $0 \in$ 䲩 $^{n}$ (output=vector)

1. Linearize: Jacobian
2. Linearize:

$$
f(x+\delta x) \approx f(x)+f^{\prime}(x) \delta x
$$

2. Solve linear equation

$$
\begin{aligned}
f(x) & +f^{\prime}(x) \delta x=0 \\
\Rightarrow \delta x & =\underset{\text { irikersse } f(x)}{\text { Jacobian }}
\end{aligned}
$$

3. Update $x$

$$
x \leftarrow x-f^{\prime}(x)^{-1} f(x)
$$

That's it! Once we have the Jacobian, just solve a linear system on each step.

Converges amazingly fast:
doubles \#digits (squares error) on each step ("quadratic convergence")!

Caveat: needs a starting guess close enough to root (google "Newton fractal"...)

## Nonlinear optimization: $\min f(x), x \in$ 政n $^{n}$

(or maximize)
$-\nabla f$ points downhill (steepest descent)
Even if we have $n=10^{6}$ parameters $x$, we can evolve them all simultaneously in the downhill direction.

Reverse-mode / adjoint / left-to-right / backpropagation: computing $\nabla \mathrm{f}$ costs about same as evaluating $f(x)$ once.

Makes large-scale optimization practical: training neural nets, optimizing shape of airplane wing, portfolio optimization...


## Nonlinear optimization: Lots of complications

- How far do we "step" in $-\nabla \mathrm{f}$ direction?
- Line search: $\min _{\alpha} f(x-\alpha \nabla f)$ - backtrack if not improved
- and/or Limit step size to trust region, grow/shrink as needed
- Details are tricky to get right
- Constraints: $\min \mathrm{f}(\mathrm{x})$ subject to $\mathrm{g}_{\mathrm{k}}(\mathrm{x}) \leq 0$
- Algorithms still need gradients $\nabla g_{k}$ !
- Faster convergence by "remembering" previous steps
- Steepest-descent tends to "zig-zag" in narrow valleys
- "Momentum" terms \& conjugate gradients - simple "memory"
- Fancier: estimate second derivative "Hessian matrix" from sequence of $\nabla \mathrm{f}$ changes: BFGS algorithm
- Lots of refinements \& competing algorithms ...
- try out multiple (pre-packaged) algorithms on your problem!


This image is in the public domain.

## Some parting advice:

Often, the main trick is finding the right mathematical formulation of your problem - i.e. what function, what constraints, what parameters? - which lets you exploit the best algorithms.
...but if you have many (> 10) parameters, always use an analytical gradient (not finite differences!) ... computed efficiently in reverse mode

## Engineering/physical optimization

## Design parameters $\mathbf{p}$ :

geometry, materials, forces, unknowns.


Physical model(s):
Solid mechanics, chemical reactions, heat transport, electromagnetism, acoustics, fluid flow...
e.g. linear model:

$$
A(\mathbf{p}) x=b(\mathbf{p})
$$

Model solution $\mathbf{x}(\mathbf{p})$ :
Forces, displacements, concentrations, temperatures, electric/magnetic fields, pressures, velocities, ...


Design objective $\mathrm{f}(\mathrm{x}(\mathbf{p}))$
Strength, speed, power, efficiency, dissipation loss, match to experiment, ...

Maximize/minimize $f(x(p))$ using gradient $\nabla_{p} f$
$\ldots \nabla_{\mathrm{p}} \mathrm{f}$ computed by reverse-mode/"adjoint" ${ }^{\mathrm{p}}$ methods

## Example: "Topology optimization" of a chair

...optimizing every voxel to support weight with minimal material (either voxel "density" or a "level-set" function)

© Source unknown. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use.


Optimization of Bone Chair
by Lothar Harmeim \& Opel GmbH
© Joris Laarman. All rights reserved. This content is excluded from our Creative Commons license. For more information, see https://ocw.mit.edu/help/faq-fair-use.

## Adjoint differentiation

## (yet another example of left-to-right/reverse-mode differentiation)

Example: gradient of scalar $f(x(p))$ where $A(p) x=$ b, i.e. $f\left(A(p)^{-1} b\right)$

- $d f=f^{\prime}(x) d x=f^{\prime}(x) d\left(A^{-1}\right) b=-f^{\prime}(x) A^{-1} d A A^{-1} b$

$$
\begin{array}{llll}
\text { row } & \overrightarrow{3} & \text { row } & =x \\
\text { vec } & \stackrel{0}{7} \\
\stackrel{\rightharpoonup}{x} & \text { vec } & & \text { "adjoint" } \\
\text { solution } v^{\top}
\end{array}
$$

- "Adjoint method:" Just multiply left-to-right! $\quad d f=-\left(f^{\prime}(x) A^{-1}\right) d A x$
- i.e. solve "adjoint equation" $A^{\top} v=f^{\prime}(x)^{\top}$ for $v \quad$ ("adjoint" meaning "transpose")
- ...then $d f=v^{\top} d A x$
- For any given parameter $p \square, \partial f / \partial p \square=v^{\top} \partial A / \partial p \square x$ (\& usually $\partial A / \partial p \square$ is very sparse)
- i.e. Takes only two solves to get both $f$ and $\nabla f$
- Solve $A x=b$ once to get $f(x)$, then solve one more time with $A^{\top}$ for $v$
- ... then all derivatives $\partial f / \partial p \square$ are just some cheap dot products


## Don't use right-to-left "forward-mode" derivatives with lots of parameters!

```
f/\partialp}\square=-\mp@subsup{f}{}{\prime}(x)(\mp@subsup{A}{}{-1}(\partialA/\partialp\squarex))=\mathrm{ one solve per parameter p}\square
    row = vector
    vector
                        (different rhs)
```

Right-to-left (a.k.a. forward mode) better when 1 input \& many outputs. Left-to-right (a.k.a. backward mode, adjoint, backpropagation) better when 1 output \& many inputs
(Note: Using dual numbers is forward mode. Most AD uses the term "forward" if it is forward mode. e.g. ForwardDiff.j in Julia is forward mode. jax.jacfwd in Python is forward mode.)

## Don't use finite differences with lots of parameters!

$$
\partial f / \partial p \square \approx[f(p+\varepsilon e \square)-f(p)] / \varepsilon \quad(e \square=\text { unit vector, } \varepsilon=\text { small number })
$$

$=$ requires one solve $x(p+\varepsilon e \square)$ for each parameter $p \square$
... even worse if you use fancier finite-difference approximations

## Adjoint differentiation with nonlinear equations

Example: gradient of scalar $f(x(p))$ where $x(p) \in \mathbb{R}^{n}$ solves $g(p, x)=0 \in \mathbb{R}^{n}$

- $g(p, x)=0 \Rightarrow d g=\partial g / \partial p d p+\partial g / \partial x d x=0 \Rightarrow d x=-(\partial g / \partial x)^{-1} \partial g / \partial p d p$
[ a.k.a. "implicit-function theorem"]

Jacobian, matrix
= inverse Jacobian,
also used in Newton solver for x !

- $d f=f^{\prime}(x) d x=-\left(f^{\prime}(x)(\partial g / \partial x)^{-1}\right) \partial g / \partial p d p$
= "adjoint"

$$
\quad \Rightarrow \text { adjoint equation: }(\partial \mathrm{g} / \partial \mathrm{x})^{\top} v=\mathrm{f}^{\prime}(x)^{\top}
$$

- i.e. Takes only two solves to get both $f$ and $\nabla f$
- one nonlinear solve for $\mathbf{x}$, and one linear solve for v !
- ... then all derivatives $\partial \mathrm{f} / \partial \mathrm{p} \square$ are just some cheap dot products


## You need to understand adjoint methods even if you use AD

- Helps understand when to use forward vs. reverse mode!
- Many physical models call large software packages written over decades in various languages, and cannot be differentiated automatically by AD
- You often just need to supply a "vector-Jacobian product" $y^{\top} d x$ for physics, or even just part of the physics, and then AD will differentiate the rest and apply the chain rule for you
- Often models involve approximate calculations, but AD tools don't know this \& spend extra effort trying to differentiate the error in your approximation
- If you solve for $x$ by an iterative method (e.g. Newton), it is inefficient for AD to backpropagate through the iteration ... instead, you want take derivative of the underlying equation $g(p, x)=0$
- For discretized physics (e.g. a finite-element methods), it is often more efficient (and sufficiently accurate) to apply adjoint method to continuous physics ("differentiate-then-discretize")

MIT OpenCourseWare
https://ocw.mit.edu

## 18.S096 Matrix Calculus for Machine Learning and Beyond

Independent Activities Period (IAP) 2023

For information about citing these materials or our Terms of Use, visit: https://ocw.mit.edu/terms.

