

AD Masterclass: Part 1

Viktor Mosenkis

nag[®]

Experts in numerical algorithms
and HPC services

The Numerical Algorithms Group

- Founded in 1970 as a co-operative project in UK
- Operates as a commercial, not-for-profit organization
 - Funded entirely by customer income
- NAG Portfolio
 - NAG Numerical Libraries
 - NAG Fortran Compiler
 - HPC services
 - Consultancy work for bespoke application development
 - Algorithmic Differentiation (AD) in collaboration with Prof. Naumann at Aachen University

AD Portfolio and AD work/consulting

■ AD Portfolio

- dco/c++ - AD tool based on operator overloading
- AD consultancy
- NAG AD Library (AD versions of NAG Library routines)

■ AD work/consulting

- Several Tier 1 & 2 investment banks have global licences for our AD software
- Similar arrangements with others across automotive
- We've helped a number of banks implement and optimise their AD codes
- We've delivered bespoke AD training for many organisations

Remarks

- Please submit your questions via the questions panel at any time during this session, these will be addressed at the end.
- A recording of this session, along with the slides will be shared with you in a day or two.

Dialogue

We want this webinar series to be interactive (even though it's hard to do)


- We want your feedback, we want to adapt material to your feedback
- Please feel free to contact us via email to ask questions at any time
- We'd love to reach out offline, discuss what's working, what to spend more time on
- For some orgs, may make sense for us to do a few bespoke sessions

Outcomes

- Recall on the problems arising when computing the derivatives by finite differences (bumping)
- Introduce Algorithmic Differentiation (AD) esp. the two AD models
 - tangent-linear and
 - adjoint model
- Demonstrate how AD can help to address the problems that arise with finite differences

Real-world examples


Figure shows the sensitivity of the amount of water flowing through the Drake passage to changes in the topography of the ocean floor. The simulation was performed with the AD-enabled MIT Global Circulation Model (MITgcm) run on a supercomputer. The ocean was meshed with 64,800 grid points.



Obtaining the gradient through finite differences took a month and a half. The adjoint AD code obtained the gradient in less than 10 minutes.

Real-world examples

AD enables **sensitivity analyses** of huge simulations, enabling shape optimization, intelligent design and **comprehensive risk studies**.



The figure shows **sensitivities of the drag coefficient** to each point on a car's surface when it moves at high speed (left) and low speed (right). The simulation was performed with **AD-enabled OpenFOAM** built on top of **dco/c++**. The normal simulation took **44s**, while the AD-enabled simulation took **273s**, to obtain the same gradient by finite differences would take roughly **5 years** (surface mesh had 5.5 million cells).

AD has been used heavily in the finance industry for many years now to compute Greeks:

- Applied to PDE and Monte Carlo codes
- Also been applied to XVA codes
- Adjoint chained backwards through calibration to get sensitivities to market instruments
- Speedups of 10x or more are common (higher the more inputs): full gradient in seconds rather than minutes. Combined with GPUs it can even give full XVA sensitivities in seconds
- Main challenge is controlling memory use

How to compute derivatives?

Typically $F : \mathbb{R}^n \rightarrow \mathbb{R}^m$, $\mathbf{y} = F(\mathbf{x})$ is not given in closed form but rather implemented in some programming language.

Writing derivative code by hand is difficult E.g.

```
1 void foo(int n, double *x, double &y){  
2     for (int i=0; i<n; i++){  
3         if (i == 0)  
4             y = sin(x[i]);  
5         else  
6             y *= x[i];  
7     }  
8 }
```

How to compute derivatives?

We first need to understand that `foo` is computing

$$y = F(\mathbf{x}) = \sin(x_0) \cdot \prod_{i=1}^{n-1} x_i.$$

We then need to differentiate this function

$$F'(x)^T = \begin{pmatrix} \cos(x_0) \cdot \prod_{i=1}^{n-1} x_i \\ \sin(x_0) \cdot \prod_{i=1, i \neq 1}^{n-1} x_i \\ \dots \\ \sin(x_0) \cdot \prod_{i=1, i \neq 1}^{n-1} x_i \end{pmatrix}$$

and then implement the corresponding derivative code.

```
1 void d_foo(int n, double *x, double *Jac){  
2     double prod = 1.0;  
3     for (int i=1; i<n; i++) prod=x[i];  
4     Jac[0] = cos(x[0])*prod;  
5     for (int i=0; i<n; i++)  
6         Jac[i] = sin(x[0]) * prod / x[i];  
7 }
```

How to compute derivatives?

Writing derivative code by hand is **very error prone job** and leads to software engineering problem of **maintaining two sources** as any changes to function `foo` must be ported to `d_foo`.

That is why practitioners tend to use more automatic approaches to compute derivatives

- Finite differences (also known as bumping)
- Algorithmic (or Automatic) Differentiation (AD)

Finite Difference

Allows to approximate the derivative of
 $F : \mathbb{R}^n \rightarrow \mathbb{R}^m, \quad \mathbf{y} = F(\mathbf{x}).$

Example (Forward finite difference Error: $O(h)$)

$$\tilde{F}(\mathbf{x}, \dot{\mathbf{x}}) = F'(\mathbf{x})\dot{\mathbf{x}} \approx \frac{F(\mathbf{x} + h\dot{\mathbf{x}}) - F(\mathbf{x})}{h}$$

Example (Centred finite difference Error: $O(h^2)$)

$$\overset{\circ}{F}(\mathbf{x}, \dot{\mathbf{x}}) = F'(\mathbf{x})\dot{\mathbf{x}} \approx \frac{F(\mathbf{x} + h\dot{\mathbf{x}}) - F(\mathbf{x} - h\dot{\mathbf{x}})}{2h}$$

Automatic approach, no need to maintain two sources.


Finite Difference: Problems

The finite differences approach suffer from two problems

- Accuracy:
How to find a proper perturbation parameter h .
- Complexity when computing the full Jacobian of a function

Finite Difference: Example Oscillating function

```
1 template <typename T>
2 T foo(T &x){
3     return sin(exp(3*x));
4 }
```




Derivative at $x = 5.5$

- with centred finite difference: -35439200.1768122
- exact: -38105558.7965109

Finite Difference: Example different code branches

```
1 template <typename T>
2 T foo(T &x){
3     if (x < 100*epsilon)
4         return x;
5     else
6         return sqrt(x);
7 }
```




Derivative at $x = 0$

- with centered finite difference: 4096.5
- exact: 1

Finite Difference: Example highly non-linear function

```
1 template <typename T>
2 T foo(T &x){
3     return sqrt(x);
4 }
```



■ with centered finite difference:

- NAN for $h > 10^{-10}$
- 50140.1439388931 for $h \approx 10^{-11}$
- 50000.0138777977 for $h \approx 10^{-13}$
- 50000.0000016598 for $h \approx 10^{-15}$ $(10^{-7} \cdot (|x| + 1) \cdot \sqrt{\epsilon})$

But for $x = 10000$ similar choice for h is failing to provide good results

Example: Dense Jacobian with forward finite difference

Consider that `void foo(int n, double *x, double &y)` implements the function $F : \mathbf{R}^n \rightarrow \mathbf{R}$, $y = F(\mathbf{x})$. The following code demonstrates the computation of approximated Jacobian (gradient) of the function F .

```
1  ...
2  foo(n, x, y);
3  ...
4  for (int i = 0; i<n; i++) {
5      x[i] += h;
6      foo(n, x, yp);
7      J[i] = (yp - y)/h;
8      x[i] -= h;
9  }
```

The function `foo` is executed $n + 1$ times

Example: Dense Jacobian with centred finite difference

Consider that `void foo(int n, double *x, double &y)` implements the function $F : \mathbb{R}^n \rightarrow \mathbb{R}$, $y = F(\mathbf{x})$. The following code demonstrates the computation of approximated Jacobian (gradient) of the function F .

```
1  ...
2  for (int i = 0; i<n; i++) {
3      tmp = x[i]
4      x[i] = tmp + h;
5      foo(n, x, yp);
6      x[i] = tmp - h;
7      foo(n, x, ym);
8      J[i] = (yp - ym)/(2*h);
9      x[i] = tmp;
10 }
```

The function `foo` is executed $2 \cdot n$ times

Algorithmic Differentiation: Tangent-Linear Model

$$F : \mathbb{R}^n \rightarrow \mathbb{R}^m, \quad \mathbf{y} = F(\mathbf{x})$$

■ Tangent-Linear Model (TLM) \dot{F} (forward mode)

$$\dot{F}(\mathbf{x}, \dot{\mathbf{x}}) = \mathbf{F}'(\mathbf{x}) \cdot \dot{\mathbf{x}} \in \mathbb{R}^{m \times n}$$

$F'(\mathbf{x})$ at $O(n) \cdot \text{Cost}(F)$

exact derivatives

$\frac{\text{Cost}(\dot{F})}{\text{Cost}(F)} \approx 2$

Jacobian with the Tangent-Linear Model

Consider

$$F : \mathbb{R}^n \rightarrow \mathbb{R}, \quad \mathbf{y} = F(\mathbf{x})$$

and $[\mathbf{y}, \mathbf{dy}] = \text{foo_dot}(\mathbf{x}, \mathbf{dx})$ contains the code that implements the tangent-linear model of F , where \mathbf{x} corresponds to x , $\dot{\mathbf{x}}$ to dx , \mathbf{y} to y and $\dot{\mathbf{y}}$ to dy . Then the Jacobian (F') of F can be computed as follows

$$dx = 0.0$$

for $i = 0; \quad i < n; \quad i + +$ do

$$dx_i = 1.0;$$

$[\mathbf{y}, \mathbf{dy}] = \text{foo_dot}(\mathbf{x}, \mathbf{dx})$

$$\frac{\partial \mathbf{y}}{\partial x_i} = \mathbf{dy}$$

$$dx_i = 0.0$$

end for

The tangent-linear model of the function `foo` is executed n times

Algorithmic Differentiation: Adjoint Model

$$F : \mathbb{R}^n \rightarrow \mathbb{R}^m, \quad \mathbf{y} = F(\mathbf{x})$$

■ Adjoint Model (ADM) \bar{F} (reverse mode)

$$\bar{F}(\mathbf{x}, \bar{\mathbf{y}}) = \underset{\in \mathbb{R}^m}{\bar{\mathbf{y}}} \cdot \underset{\in \mathbb{R}^{m \times n}}{F'(\mathbf{x})} = F'(\mathbf{x})^T \cdot \bar{\mathbf{y}} = \bar{\mathbf{x}}$$

□ $F'(\mathbf{x})$ at $O(m) \cdot \text{Cost}(F)$

□ exact derivatives

□ $\frac{\text{Cost}(\bar{F})}{\text{Cost}(F)} < 30$

Jacobian with the Adjoint Model

Consider

$$F : \mathbb{R}^n \rightarrow \mathbb{R}, \quad \mathbf{y} = F(\mathbf{x})$$

and $[y, dx] = \text{foo_bar}(x, dy)$ contains the code that implements the adjoint model of F , where \mathbf{x} corresponds to x , \bar{x} to dx , \mathbf{y} to y and \bar{y} to dy . Then the Jacobian (F') of F can be computed as follows

$$dy = 1.0;$$

$$[y, dx] = \text{foo_bar}(x, dy)$$

for $i = 0$; $i < n$; $i++$ do

$$\frac{\partial y}{\partial x_i} = dx_i$$

end for

The adjoint model of the function `foo` is executed only once. Cost for bad implementation of `foo_bar` may be much higher than 30. The challenge is to make a good implementation.

Live demo: Race!

Summary

Finite difference suffers from two main problems

- Accuracy (due to approximation)
- $O(n)$ complexity for computing dense Jacobian

Both problems can be addressed by AD. Tangent-linear as well as the adjoint model compute **exact derivatives** with machine accuracy.

With the adjoint model the complexity of Jacobian computation is **independent from the number of inputs of the function ($O(m)$)**.

Outlook: AD Masterclass Part 2

In the next part we will

- Learn how AD works (basic approach behind AD),
- Learn how to apply operator overloading AD tool to a code
- Learn how to write the driver computing the Jacobian for both
 - tangent-linear and
 - adjoint model
- Understand how to choose the best model (tangent-linear or adjoint) for our problem (also Masterclass Part 3)
- Understand the caveats of using AD especially the adjoint model.

You will see a survey on your screen after exiting from this session.

We would appreciate your feedback.

We are no moving on the Q&A Session

References

-  A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, *Automatic differentiation in machine learning: a survey*, arXiv preprint arXiv:1502.05767 (2015).
-  Felix Gremse, Benjamin Theek, Sijumon Kunjachan, Wiltrud Lederle, Alessa Pardo, Stefan Barth, Twan Lammers, Uwe Naumann, and Fabian Kiessling, *Absorption reconstruction improves biodistribution assessment of fluorescent nanoprobes using hybrid fluorescence-mediated tomography*, *Theranostics* 4 (2014), 960–971.

References (cont.)

-  Mathias Luers, Max Sagebaum, Sebastian Mann, Jan Backhaus, David Grossmann, and Nicolas Gauger, *Adjoint-based volumetric shape optimization of turbine blades*, AIAA 2018-3638 (2018).
-  Uwe Naumann, Jean Utke, Carl Wunsch, Patrick Heimbach, Chris Hill, Derya Ozyurt, Carl Wunschyy, Mike Fagan, Nathan Tallent, and Michelle Strout, *Adjoint Code by Source Transformation with OpenAD/F*, Proceedings / European Conference on Computational Fluid Dynamics : Egmond aan Zee, the Netherlands, 5 - 8 September 2006 / ECCOMAS, European Community on Computational Methods in Applied Sciences. Ed.: P. Wesseling ... (Delft), TU Delft, 2006, Datentraeger: 1 CD-ROM, p. 19 S.

References (cont.)

-  Andreas Püttmann, Sebastian Schnittert, Uwe Naumann, and Eric von Lieres, *Fast and accurate parameter sensitivities for the general rate model of column liquid chromatography*, Computers & Chemical Engineering 56 (2013), 46 – 57.
-  Florian Rauser, Jan Riehme, Klaus Leppkes, Peter Korn, and Uwe Naumann, *On the use of discrete adjoints in goal error estimation for shallow water equations*, Procedia Computer Science 1 (2010), no. 1, 107 – 115, ICCS 2010.
-  Antoine Savine, *Modern computational finance: Aad and parallel simulations*, Wiley, 2018.

References (cont.)

 Markus Towara, Michel Schanen, and Uwe Naumann,
MPI-Parallel Discrete Adjoint OpenFOAM, Computational
science at the gates of nature : International Conference on
Computational Science (ICCS 2015) ; Reykjavík, Iceland, 1 - 3
June 2015 / [organised by: Háskólinn í Reykjavík ...]. Ed.:
Slawomir Koziel ... - Pt. 1 (Red Hook, NY), Procedia computer
science, vol. 51, International Conference On Computational
Science, Reykjavík (Iceland), 1 Jun 2015 - 3 Jun 2015, Curran,
Jun 2015, pp. 19–28.

References (cont.)

-  J. Ungermann, J. Blank, J. Lotz, K. Leppkes, L. Hoffmann, T. Guggenmoser, M. Kaufmann, P. Preusse, U. Naumann, and M. Riese, *A 3-D tomographic retrieval approach with advection compensation for the air-borne limb-imager GLORIA*, Atmospheric measurement techniques 4 (2011), 2509 – 2529, Record converted from VDB: 12.11.2012.
-  Andrey Vlasenko, Peter Korn, Jan Riehme, and Uwe Naumann, *Estimation of Data Assimilation Error: A Shallow-Water Model Study*, Monthly Weather Review 142 (2014), no. 7, 2502–2520.

References (cont.)

-  V. Vassiliadis, J. Riehme, J. Deussen, K. Parasyris, C. D. Antonopoulos, N. Bellas, S. Lalíš, and U. Naumann, *Towards automatic significance analysis for approximate computing*, 2016 IEEE/ACM International Symposium on Code Generation and Optimization (CGO), 2016, pp. 182–193.