

An Introduction to Algorithmic Differentiation (AD)

Systems seminar

Joel Andersson

28 October 2016

1 Algorithmic differentiation (AD)

- Methods for calculating derivatives
- The forward and reverse modes
- Checkpointing
- Jacobians and Hessians
- AD software
- Summary

2 CasADi

- Recall: Optimal control
- Scope of CasADi
- Capabilities
- Summary

Table of Contents

1 Algorithmic differentiation (AD)

- Methods for calculating derivatives
 - The forward and reverse modes
 - Checkpointing
 - Jacobians and Hessians
 - AD software
 - Summary

2 CasADi

- Recall: Optimal control
- Scope of CasADi
- Capabilities
- Summary

Derivatives play a central role in nonlinear optimization – how compute them?


- By hand ← **Time consuming & error prone!**
- Symbolic differentiation
- Finite difference approximation
- Complex step differentiation (“Imaginary trick”)
- Algorithmic differentiation (AD)

Symbolic differentiation

Obtain derivatives with a computer algebra system (CAS):

- Symbolic Toolbox for MATLAB / MuPAD
- SymPy
- ...

Easy to use but often doesn't scale well with number of operations



Finite differences

Consider a function $f : \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_y}$ with Jacobian $J(x) = \frac{\partial f}{\partial x}$

$$J(x) \hat{x} \approx \frac{f(x + t \hat{x}) - f(x)}{t}$$

Pros and cons:

- + Easy to implement and relatively fast
- Poor accuracy, need to carefully choose t :
 - Small $t \Rightarrow$ *cancellation errors*
 - Large $t \Rightarrow$ *approximation errors*
- No efficient way to calculate $\hat{y}^T J(x)$

Complex step differentiation (“Imaginary trick”)

Finite differences with imaginary perturbation:

$$J(x) \hat{x} \approx \Re \left(\frac{f(x + i t \hat{x}) - f(x)}{i t} \right) = -\frac{\Im(f(x + i t \hat{x}))}{t}$$

Pros and cons:

- + Easy to implement in MATLAB/Octave, Python, (Julia?) and relatively fast
- + Good accuracy (choose t very small)
- Error prone, e.g. x'
- Restricted
- No efficient way to calculate $\hat{y}^T J(x)$

Table of Contents

- 1 Algorithmic differentiation (AD)
 - Methods for calculating derivatives
 - **The forward and reverse modes**
 - Checkpointing
 - Jacobians and Hessians
 - AD software
 - Summary
- 2 CasADi
 - Recall: Optimal control
 - Scope of CasADi
 - Capabilities
 - Summary

Algorithmic differentiation (AD) (e.g. Griewank & Walther, 2008)

Decomposable function: $y = F(x)$

- $F : \mathbb{R}^{n_0} \rightarrow \mathbb{R}^{n_K}$ sufficiently smooth
- Decompose into “atomic operations” with known differentiation rules:

```
z0 ← x
for k = 1, ..., K do
    zk ← fk ({zi}i ∈ Ik)
end for
y ← zK
return y
```

Such a decomposition is always available if F written as a computer program!

Example

$$y = \sin(\sqrt{x})$$

```
z0 ← x
z1 = √z0
z2 = sin z1
y ← z2
return y
```

- Decomposition can be with simple scalar operations ...
 - $x + y$, $x * y$, $\sin(x)$, x^y
- ...or with higher-level operations for which a chain-rule can be defined
 - x^T , $x[i] = y$, XY , e^X
 - E.g. gradient of $\det(X)$: $\det(X) X^{-T}$
 - Linear and nonlinear systems of equations (in CasADi!)
 - Initial-value problems in ODE or DAE (in CasADi!)

Differentiate the algorithm!

```
z0 ← x
for k = 1, ..., K do
  zk ← fk ({zi}i ∈ Ik)
end for
y ← zK
return y
```



```
z0 ← x
dz0 / dx ← I
for k = 1, ..., K do
  zk ← fk ({zi}i ∈ Ik)
  dzk / dx ← ∑i ∈ Ik (∂fk / ∂zi) ({zi}i ∈ Ik) dzi / dx
end for
y ← zK
J ← dzK / dx
return y, J
```

Write as a system of linear equations:

$$\frac{dz}{dx} = B + L \frac{dz}{dx}, \quad J = A^T \frac{dz}{dx},$$

Write as a system of linear equations:

$$\frac{dz}{dx} = B + L \frac{dz}{dx}, \quad J = A^T \frac{dz}{dx},$$

with

$$z = \begin{pmatrix} z_0 \\ z_1 \\ \vdots \\ z_K \end{pmatrix}, \quad A = \begin{pmatrix} 0 \\ \vdots \\ 0 \\ I \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} I \\ 0 \\ \vdots \\ 0 \end{pmatrix},$$

with I and 0 of appropriate dimensions, as well as the *extended Jacobian*,

$$L = \begin{pmatrix} 0 & \cdots & \cdots & 0 \\ \frac{\partial f_1}{\partial z_0} & \ddots & & \vdots \\ \vdots & \ddots & \ddots & \vdots \\ \frac{\partial f_K}{\partial z_0} & \cdots & \frac{\partial f_K}{\partial z_{K-1}} & 0 \end{pmatrix},$$

Since $I - L$ is invertible, we can solve for J :

$$J = A^T (I - L)^{-1} B$$

- Have $J = A^\top (I - L)^{-1} B$
- Multiply J from the right: **Forward mode of AD**
 - $\hat{y} := J \hat{x} = A^\top (I - L)^{-1} B \hat{x}$
 - Cheap with *forward substitution* of lower triangular $(I - L)$
 - Computational cost: \approx cost of evaluating F
 - Small memory requirements (no storage of L needed)
- Multiply J from the left: **Reverse mode of AD**
 - $\bar{x} := J^\top \bar{y} = B^\top (I - L)^{-\top} A \bar{y}$
 - Cheap with *backward substitution* of upper triangular $(I - L)^\top$
 - Computational cost: \approx cost of evaluating F
 - If $F(x)$ is scalar, $\bar{y} = 1$ gives $\nabla_x F(x)$
 - Intermediate operations (or their linearization) must be stored
 - Remedy: Checkpointing

Table of Contents

- 1 Algorithmic differentiation (AD)
 - Methods for calculating derivatives
 - The forward and reverse modes
 - **Checkpointing**
 - Jacobians and Hessians
 - AD software
 - Summary
- 2 CasADi
 - Recall: Optimal control
 - Scope of CasADi
 - Capabilities
 - Summary

Checkpointing

- E.g.

```
y = f(x) :  
z0 ← x  
for k = 1, ..., 1000 do  
    zk ← sin(zk-1)  
end for  
y ← z1000
```

- Divide up!:

```
y = f(x) :  
z0 ← x  
for k = 1, ..., 20 do  
    zk ← g(zk-1)  
end for  
y ← z20
```

```
y = g(x) :  
z0 ← x  
for k = 1, ..., 50 do  
    zk ← sin(zk-1)  
end for  
y ← z50
```

- Memory for reverse mode decreases from $50 * 20$ to $50 + 20$
- Price: Need to reevaluate $g(x)$ – twice as many sin calls!
- Checkpointing trades memory for reevaluation!

Table of Contents

- 1 Algorithmic differentiation (AD)
 - Methods for calculating derivatives
 - The forward and reverse modes
 - Checkpointing
 - **Jacobians and Hessians**
 - AD software
 - Summary
- 2 CasADi
 - Recall: Optimal control
 - Scope of CasADi
 - Capabilities
 - Summary

Calculating complete Jacobians and Hessians

- Jacobians can be calculated by multiplying with n_{col} vectors from the right or n_{row} vectors from the left
- Worst-case: $\approx \min(n_{\text{row}}, n_{\text{col}})$ times cost of evaluating F
- Hessians can be calculated as Jacobian-of-gradient – exploit symmetry!
- *Much cheaper* if J is sparse, e.g. banded

Exploiting sparsity: illustration

$$\bullet J = \begin{bmatrix} * & & & \\ & * & & \\ & & * & \\ & & & * \end{bmatrix}$$

$$\Rightarrow \hat{x} = [1, 1, 1, 1]$$

$$\bullet J = \begin{bmatrix} * & & & \\ * & * & & \\ * & & * & \\ * & & & * \end{bmatrix}$$

$$\Rightarrow \hat{x}_1 = [0, 1, 1, 1], \hat{x}_2 = [1, 0, 0, 0]$$

$$\bullet J = \begin{bmatrix} * & * & * & * \\ * & & & \\ * & & & \\ * & & & \end{bmatrix}$$

$$\Rightarrow \hat{x}_1 = [1, 0, 0, 0], \bar{y}_2 = [1, 0, 0, 0]$$

Finding a small set of vectors

- NP hard combinatorial problem!
- Graph coloring techniques usually work well (cf. Gebremedhin et al, 2005)

Table of Contents

- 1 Algorithmic differentiation (AD)
 - Methods for calculating derivatives
 - The forward and reverse modes
 - Checkpointing
 - Jacobians and Hessians
 - **AD software**
 - Summary
- 2 CasADi
 - Recall: Optimal control
 - Scope of CasADi
 - Capabilities
 - Summary

- Tools are generally divided into:
 - Operator-overloading (OO): Calculate derivatives "piggyback" during numerical evaluation
 - Source-code-transformation (SCT): Generate new source code for derivatives
- Generic tools to differentiate "black-box" code
 - Language-specific: www.autodiff.org, Wikipedia
 - ADOL-C (OO), ADIC (SCT), CppAD (OO) for C/C++
 - ADIFOR (SCT), TAPENADE (SCT) for FORTRAN
- AD implemented for domain specific languages (DMS)
 - CasADi (SCT), AMPL (OO), GAMS (OO), JuliaOpt (OO)

Table of Contents

- 1 Algorithmic differentiation (AD)
 - Methods for calculating derivatives
 - The forward and reverse modes
 - Checkpointing
 - Jacobians and Hessians
 - AD software
 - **Summary**

- 2 CasADi
 - Recall: Optimal control
 - Scope of CasADi
 - Capabilities
 - Summary

Summary: AD


- Jacobian-times-vector products cheap using *forward mode AD*
- Vector-times-Jacobian products cheap using *reverse mode AD*
 - *Checkpointing* can avoid excessive memory use
- Complete Jacobians and Hessians: depends on sparsity pattern
 - Worse case: $\approx \min(n_{\text{row}}, n_{\text{col}})$ times cost of evaluating F
 - Good heuristics exist for complete sparse Jacobians and Hessians
- Software exists for many languages and domain specific languages

Table of Contents

- 1 Algorithmic differentiation (AD)
 - Methods for calculating derivatives
 - The forward and reverse modes
 - Checkpointing
 - Jacobians and Hessians
 - AD software
 - Summary

- 2 CasADi
 - Recall: Optimal control
 - Scope of CasADi
 - Capabilities
 - Summary

Simplified optimal control problem (OCP)¹



$$\underset{x(\cdot), u(\cdot)}{\text{minimize}} \int_0^T L(x(t), u(t)) dt + E(x(T))$$

subject to

$$x(0) = \underline{x}_0,$$

(Initial condition)

$$\dot{x}(t) = f(x(t), u(t)), \quad t \in [0, T],$$

(System model)

$$\underline{x} \leq x(t) \leq \bar{x}, \quad t \in [0, T],$$

(Path constraints)

$$\underline{u} \leq u(t) \leq \bar{u}, \quad t \in [0, T],$$

(Control bounds)


$$\underline{x}_f \leq x(T) \leq \bar{x}_f$$

(Terminal constraints)

¹A.k.a. “dynamic optimization problem”

Model predictive control (MPC)

- Solve a sequence of OCPs to determine the optimal control action:



- Give first control move u_0 back to real-world system. Move horizon.
- Result: **Feedback law** $u_0(x_0)$. Allows to react to disturbances and modeling errors.

Slightly more realistic optimal control problem

Optimal control problem in differential-algebraic equations:

$$\begin{aligned} & \text{minimize} && \int_0^T L(x(t), z(t), u(t), p, t) dt + E(x(T), p) \\ & \text{subject to} && \begin{cases} \dot{x}(t) = f(x(t), z(t), u(t), p, t) \\ 0 = g(x(t), z(t), u(t), p, t) \end{cases} \quad t \in [0, T] \\ & && \underline{x} \leq x(t) \leq \bar{x}, \quad t \in (0, T) \\ & && \underline{u} \leq u(t) \leq \bar{u}, \quad t \in [0, T] \\ & && x(0) = \bar{x}_0, \quad \underline{x}_f \leq x(T) \leq \bar{x}_f, \quad \underline{p} \leq p \leq \bar{p}, \end{aligned}$$

$x(\cdot) \in \mathbb{R}^{n_x}$	Differential states	$L(\dots)$	Lagrange cost term
$z(\cdot) \in \mathbb{R}^{n_z}$	Algebraic variables	$E(\dots)$	Mayer cost term
$u(\cdot) \in \mathbb{R}^{n_u}$	Controls	$f(\dots)$	Differential equations
$p \in \mathbb{R}^{n_p}$	Parameters	$g(\dots)$	Algebraic constraints

Table of Contents

- 1 Algorithmic differentiation (AD)
 - Methods for calculating derivatives
 - The forward and reverse modes
 - Checkpointing
 - Jacobians and Hessians
 - AD software
 - Summary

- 2 CasADi
 - Recall: Optimal control
 - **Scope of CasADi**
 - Capabilities
 - Summary

CasADi

- Started as an implementation of AD using CAS-like syntax
- Current scope: Numerical optimization general (or just AD!)
- In particular: Facilitates the solution of optimal control problems (OCPs)
 - *Facilitates*, not actually *solves* the OCPs
 - Write state-of-the-art OCP algorithms with very little code!
- Available for C++, Python, MATLAB/Octave
- Free & open-source (LGPL), also for commercial use
- Project started in December 2009, now (almost) at version 3.0
- Main developers: Joel Andersson and Joris Gillis

Search or type a command | Explore | Gist | Blog | Help | jgillis | Fork

PUBLIC casadi / casadi | Pull Request | Unwatch | Unstar | Fork

Code | Network | Pull Requests (0) | Issues (120) | **Wiki** | Graphs | Settings

Home | Pages | Wiki History | Git Access

Home

New Page | Edit Page | Page History

CasADi

Welcome to the CasADi wiki!

CasADi is a symbolic framework for [automatic differentiation](#) and numeric optimization. Using the syntax of computer algebra systems, it implements automatic differentiation in forward and adjoint modes by means of a hybrid symbolic/numeric approach. The main purpose of the tool is to be a low-level tool for quick, yet highly efficient implementation of algorithms for numerical optimization. Of particular interest is dynamic optimization, using either a collocation approach, or a shooting-based approach using embedded ODE/DAE-integrators. In either case, CasADi relieves the user from the work of efficiently calculating the relevant derivative or ODE/DAE sensitivity information to an arbitrary degree, as needed by the NLP solver. This together with full-featured [Python](#) and [Octave](#) front ends, as well as back ends to state-of-the-art codes such as [Sundials](#) (CVODES, IDAS and KINSOL), [IPOPT](#) and [KNITRO](#), drastically reduces the effort of implementing the methods compared to a pure C/C++/Fortran approach.

Every feature of CasADi (with very few exceptions) is available in C++, Python and Octave, with little to no difference in performance, so the user has the possibility of working completely in C++, Python or Octave or mixing the languages. We recommend new users to try out the Python version first, since it allows interactivity and is more stable and better documented than the Octave front-end.

CasADi is an open-source tool, written in self-contained C++ code, depending only on the Standard Template Library. It is developed by Joel Andersson and Joris Gillis at the [Optimization in Engineering Center, OPTeC](#) of the K.U. Leuven under supervision of [Moritz Diehl](#). CasADi is distributed under the [LGPL](#) license, meaning the code can be used royalty-free even in commercial applications.

- Decomposes algorithms into a sequence of either **scalar** or **sparse matrix-valued atomic operations**
- New symbolic expressions generated for derivatives: “Source code transformation” approach
 - Forward mode: Jacobian-times-vector products
 - Reverse mode: vector-times-Jacobian products
 - Sparse Jacobians and Hessians via:
 - Automatic detection of sparsity pattern (nontrivial!)
 - Graph coloring techniques to exploit sparsity & symmetry
 - Arbitrary order
- Supports **high-level operations**: matrix-operations, implicit functions, calls to DAE integrators

Table of Contents

- 1 Algorithmic differentiation (AD)
 - Methods for calculating derivatives
 - The forward and reverse modes
 - Checkpointing
 - Jacobians and Hessians
 - AD software
 - Summary

- 2 CasADi
 - Recall: Optimal control
 - Scope of CasADi
 - **Capabilities**
 - Summary

Quadratic programs (QPs)

- User needs to write the problem in the following standard form:

$$\begin{array}{ll} \underset{x}{\text{minimize}} & f(x) \\ \text{subject to} & \underline{g} \leq g(x) \leq \bar{g} \\ & \underline{x} \leq x \leq \bar{x} \end{array}$$

where $f(x)$ is a convex quadratic function and $g(x)$ is a linear function.

- QP solvers available: qpOASES, OOQP, CPLEX, GUROBI
- Solver “plugins” can be added post-installation
- CasADi automatically generates matrix sparsities
- GUROBI & CPLEX: a subset of x can be integer-valued!

Nonlinear programs (NLPs)

- User needs to write the problem in the following standard form:

$$\begin{array}{ll} \underset{x}{\text{minimize}} & f(x) \\ \text{subject to} & \underline{g} \leq g(x) \leq \bar{g} \\ & \underline{x} \leq x \leq \bar{x} \end{array}$$

where $f(x)$ and $g(x)$ twice continuously differentiable functions

- NLP solvers available: IPOPT, SNOPT, KNITRO, WORHP, blockSQP, Bonmin, CasADi's own
- Solver “plugins” can be added post-installation
- CasADi automatically generates derivative information
- Bonmin & KNITRO: a subset of x can be integer-valued!

Integrators

- Solves initial-value problems in ordinary or differential-algebraic equations (ODE/DAE)
- Given a **DAE with fixed initial values** coupled to another (linear) **DAE with fixed terminal value** (both with quadratures):

$$\left\{ \begin{array}{l} \dot{x} = f_x(x, z, p, t) \\ 0 = f_z(x, z, p, t) \\ \dot{q} = f_q(x, z, p, t) \end{array} \right. \quad t \in [0, T] \quad \left\{ \begin{array}{l} x(0) = x_0 \\ q(0) = 0 \\ \tilde{x}(T) = \tilde{x}_0 \\ \tilde{q}(T) = 0 \end{array} \right.$$
$$\left\{ \begin{array}{l} -\dot{\tilde{x}} = \tilde{f}_x(\tilde{x}, \tilde{z}, \tilde{p}, x, z, p, t) \\ 0 = \tilde{f}_z(\tilde{x}, \tilde{z}, \tilde{p}, x, z, p, t) \\ -\dot{\tilde{q}} = \tilde{f}_q(\tilde{x}, \tilde{z}, \tilde{p}, x, z, p, t) \end{array} \right.$$

- An *integrator in CasADi* is a mapping from $\{x_0, p, \tilde{x}_0, \tilde{p}\}$ to $\{x(T), q(T), \tilde{x}(0), \tilde{q}(0)\}$
- Enables **automatic forward and adjoint sensitivity analysis**

Table of Contents

- 1 Algorithmic differentiation (AD)
 - Methods for calculating derivatives
 - The forward and reverse modes
 - Checkpointing
 - Jacobians and Hessians
 - AD software
 - Summary

- 2 CasADi
 - Recall: Optimal control
 - Scope of CasADi
 - Capabilities
 - **Summary**

Summary: CasADi

- Open-source framework for numerical optimization
- Central feature I: general-purpose implementation of AD
- Central feature II: solve standard problems conveniently
 - QPs
 - NLPs
 - Initial-value problems in ODE/DAE
 - Everything you need for optimal control!
- Video tutorial: <http://tutorial.casadi.org/>