

Exercise 5: Algorithmic Differentiation

Prof. Dr. Moritz Diehl, Dimitris Kouzoupis, Andrea Zanelli, Florian Messerer, Yizhen Wang

The aim of this exercise is to gain experience with the two modes of algorithmic differentiation (AD) discussed in the class.

1. **Forward and backward algorithmic differentiation:** Consider the following discrete-time dynamical system:

$$x_{k+1} = x_k + h((1 - x_k)x_k + u_k), \quad (1)$$

with state $x_k \in \mathbb{R}$, input $u_k \in \mathbb{R}$ and $h \in \mathbb{R}_+$ a constant parameter (you can think of it as the time step of an explicit Euler integrator). We are interested in simulating the dynamics forward for N steps, starting from the initial value $x_0 = \bar{x}_0$ and computing the derivatives of the obtained states with respect to the controls:

$$\frac{\partial x_i}{\partial u_{j-1}}, \quad \forall i, j = 1, \dots, N. \quad (2)$$

- (a) Fix $\bar{x}_0 = 0.5$, $N = 50$, $h = 0.1$. Make sure to define them once only, so you can easily adapt their values later. Using CasADi, implement the function $\Phi : \mathbb{R}^N \rightarrow \mathbb{R}^N$ that maps controls to the obtained state trajectory

$$x = \Phi(u), \quad (3)$$

where $x = (x_1, \dots, x_N)$ and $u = (u_0, \dots, u_{N-1})$ denote the vector of stacked states and controls respectively. Define a CasADi function that outputs the Jacobian of x with respect to u ,

$$J(u) = \frac{\partial \Phi(u)}{\partial u}. \quad (4)$$

You will use the output of this function as a reference for your implementations in the rest of the exercise.

- (b) Implement a MATLAB/Python function `forw_AD(u, m, x0, h)` that takes as input a vector containing the values for u and a scalar m and returns the derivative $\frac{\partial x_m}{\partial u_m}$, i.e., the m -th *column* of the Jacobian, evaluated at input u , using forward AD. Check that the result provided by your implementation is equal to the corresponding entries in the output obtained with CasADi, e.g., by evaluating both at randomly generated values of u , `u_tst = rand(N,1)` resp. `u_tst = np.random.rand(N,1)`, and comparing the maximum of their elementwise absolute difference. You should be able to reach machine precision, i.e., order of magnitude around 10^{-16} .
- (c) Analogously, implement a MATLAB/Python function `back_AD` that takes as input u , a scalar m as well as the parameters. It returns the derivative $\frac{\partial x_m}{\partial u}$, i.e., the m -th *row* of the Jacobian, using backward AD. Check that the result provided by your implementation is equal to the corresponding entry in the output obtained with CasADi.
- (d) Implement now a function `J_FAD` that takes as inputs u and a scalar m and, using forward AD, computes the last m rows of the Jacobian $\frac{\partial \Phi(u)}{\partial u}$ containing the derivatives of the last m states in the simulation with respect to the all the controls. Again, validate your results against the reference output.

Note: You can just call `forw_AD` several times, but then you unnecessarily repeat some computations. Can you do better?

- (e) Analogously, implement a function `J_BAD` that takes as inputs u and a scalar m and computes the last m rows of the Jacobian $\frac{\partial \Phi(u)}{\partial u}$ using your implementation of backward AD and validate it against the reference.

Again: You can just call `back_AD` several times, but then you would unnecessarily repeat some computations. Can you do better?

- (f) Which of the two implementation do you expect to be more performant for small values of m ? Which one for high values of m ? Why?

Each forward AD pass computes a column of the Jacobian, while each backward AD pass computes a row. To compute m rows, we need m backward passes or N forward passes/ The cost of a forward pass is roughly $2 \text{cost}(\Phi)$, while the cost of a backward pass is roughly $3 \text{cost}(\Phi)$. Thus, the cost of computing m rows with the backward mode is $3m \text{cost}(\Phi)$, while the cost of computing m columns with the forward mode is $2N \text{cost}(\Phi)$. In consequence, the forward mode is more efficient for small values of m , while the backward mode is more efficient for values of close to N .

Of course, if we were interested in computing columns of the Jacobian, the story would be different.

- (g) Run your implementations for m ranging from 1 to N and measure the execution time using the MATLAB functions `tic` and `toc`/Python function `timeit.default_timer()`. For this simulation choose $h = 0.01$ and $N = 500$. Plot the obtained execution times as a function of m . Do the results validate your considerations from the previous question?

Note: You may need to adapt the value of N to obtain a reasonable execution time for the full script. Overall, you should not wait more than a few seconds. If it takes too long, lower the value of N .