

CasADi: a software framework for nonlinear optimization and optimal control

Joel A. E. Andersson¹ · Joris Gillis^{2,3} ·
Greg Horn⁴ · James B. Rawlings¹ · Moritz Diehl⁵

Received: 20 May 2017 / Accepted: 3 March 2018 / Published online: 11 July 2018
© Springer-Verlag GmbH Germany, part of Springer Nature and The Mathematical Programming Society 2018

Abstract We present CasADi, an open-source software framework for numerical optimization. CasADi is a general-purpose tool that can be used to model and solve optimization problems with a large degree of flexibility, larger than what is associated with popular algebraic modeling languages such as AMPL, GAMS, JuMP or Pyomo. Of special interest are problems constrained by differential equations, i.e. optimal control problems. CasADi is written in self-contained C++, but is most conveniently used via full-featured interfaces to Python, MATLAB or Octave. Since its inception in

The software reviewed as part of this submission was given the digital object identifier (DOI) <https://doi.org/10.5281/zenodo.1257968>.

✉ Joel A. E. Andersson
jandersson@wisc.edu

Joris Gillis
joris.gillis@kuleuven.be

Greg Horn
greg@kittyhawk.aero

James B. Rawlings
james.rawlings@wisc.edu

Moritz Diehl
moritz.diehl@imtek.uni-freiburg.de

- ¹ Department of Chemical and Biological Engineering, University of Wisconsin-Madison, Madison, USA
- ² MECO Research Team, Department Mechanical Engineering, KU Leuven, Leuven, Belgium
- ³ DMMS Lab, Flanders Make, Leuven, Belgium
- ⁴ Kitty Hawk, Mountain View, USA
- ⁵ Department of Microsystems Engineering IMTEK, University of Freiburg, Freiburg, Germany

late 2009, it has been used successfully for academic teaching as well as in applications from multiple fields, including process control, robotics and aerospace. This article gives an up-to-date and accessible introduction to the CasADi framework, which has undergone numerous design improvements over the last 7 years.

Keywords Optimization · Optimal control · Open source optimization software

Mathematics Subject Classification 90C99 · 93A30 · 97A01

1 Introduction

CasADi is an open-source software framework for numerical optimization, offering an alternative to conventional algebraic modeling languages such as AMPL [45], Pyomo [67] and JuMP [35]. Compared to these tools, the approach taken by CasADi, outlined in this paper, is more flexible, but also lower-level, requiring an understanding of the *expression graphs* the user is expected to construct as part of the modeling process.

This flexible approach is in particular valuable for problems constrained by differential equations, introduced in the following.

1.1 Optimal control problems

Consider the following basic optimal control problem (OCP) in ordinary differential equations (ODE):

$$\begin{aligned} & \text{minimize}_{x(\cdot), u(\cdot), p} \int_0^T L(x(t), u(t), p) dt + E(x(T), p) \\ & \text{subject to} \quad \left. \begin{aligned} \dot{x}(t) &= f(x(t), u(t), p), \\ u(t) &\in \mathcal{U}, \quad x(t) \in \mathcal{X}, \\ x(0) &\in \mathcal{X}_0, \quad x(T) \in \mathcal{X}_T, \quad p \in \mathcal{P}, \end{aligned} \right\} t \in [0, T] \end{aligned} \quad (\text{OCP})$$

where $x(t) \in \mathbb{R}^{N_x}$ is the vector of (differential) states, $u(t) \in \mathbb{R}^{N_u}$ is the vector of free control signals and $p \in \mathbb{R}^{N_p}$ is a vector of free parameters in the model. The OCP here consists of a Lagrange term (L), a Mayer term (E), as well as an ODE with initial (\mathcal{X}_0) and terminal (\mathcal{X}_T) conditions. Finally, there are admissible sets for the states (\mathcal{X}), control (\mathcal{U}) and parameters (\mathcal{P}). For simplicity, all sets can be assumed to be simple intervals.

Problems of form (OCP) can be efficiently solved with the *direct approach*, where (OCP) is transcribed into a nonlinear program (NLP):

$$\begin{aligned} & \text{minimize}_w \quad J(w) \\ & \text{subject to} \quad g(w) = 0, \quad w \in \mathcal{W}, \end{aligned} \quad (\text{NLP})$$

where $w \in \mathbb{R}^{N_w}$ is the decision variable, J the objective function and \mathcal{W} again taken to be an interval set.

Popular direct methods include *direct collocation* [94,95], which reached widespread popularity through the work of Biegler and Cuthrell [14,27], and *direct multiple shooting* by Bock and Plitt [15,16]. Both these methods exhibit good convergence properties and can be easily parallelized.

Real-world optimal control problems are often significantly more general than (OCP). Industrially relevant features include multi-stage formulations, i.e., different dynamics in different parts of the time horizon, robust optimal control, problems with integer variables, differential-algebraic equations (DAEs) instead of ODEs, and multipoint constraints such as periodicity.

1.2 Scope of CasADi

CasADi started out as a tool for algorithmic differentiation (AD) using a syntax similar to a computer-algebra system (CAS), explaining its name. While state-of-the-art AD is still a key feature of CasADi, the focus has since shifted towards optimization. In its current form, CasADi provides a set of general-purpose building blocks that drastically decreases the effort needed to implement a large set of algorithms for numerical optimal control, without sacrificing efficiency. This “toolkit design” makes CasADi suitable for teaching optimal control to graduate-level students and allows researchers and industrial practitioners to write codes, with a modest programming effort, customized to a particular application or problem structure.

1.3 Organization of the paper

The remainder of the paper is organized as follows. We start by introducing the reader to CasADi’s symbolic core in Sect. 2. The key property of the symbolic core is a state-of-the-art implementation of AD. Some unique features of CasADi are introduced, including how sparse matrix-valued atomic operations can be used in a source-code-transformation AD framework.

In Sect. 3, we show how systems of linear or nonlinear equations, as well as initial-value problems in ODEs or DAEs, can be embedded into symbolic expressions, while maintaining differentiability to arbitrary order. This automatic sensitivity analysis for ODEs and DAEs is, to the best knowledge of the authors, a unique feature of CasADi.

Section 4 outlines how certain optimization problems in canonical form can be solved with CasADi, including nonlinear programs (NLPs), linear programs (LPs) and quadratic programs (QPs), potentially with a subset of the variables confined to integer values, i.e. mixed-integer formulations. CasADi provides a common interface for formulating such problems, while delegating the actual numerical solution to a third-party solver, either free or commercial, or to an in-house solver, distributed with CasADi.

Section 5 consists of a tutorial, showing some basic use cases of CasADi. The tutorial mainly serves to illustrate the syntax, scope and usage of the tool.

Finally, Sect. 6 gives an overview of applications where CasADi has been successfully used to date before Sect. 7 wraps up the paper.

2 Symbolic framework

The core of CasADi consists of a symbolic framework that allows users to construct expressions and use these to define automatically differentiable functions. These general-purpose expressions have no notion of optimization and are best likened with expressions in e.g. MATLAB's Symbolic Toolbox or Python's SymPy package. Once the expressions have been created, they can be used to efficiently obtain new expressions for derivatives using AD or be evaluated efficiently, either in CasADi's virtual machines or by using CasADi to generate self-contained C code.

We detail the symbolic framework in the following sections, where we also provide MATLAB/Octave and Python code snippets¹ corresponding to CasADi 3.1 below in order to illustrate the functionality. For a self-contained and up-to-date walkthrough of CasADi's syntax, we recommend the user guide [6].

2.1 Syntax and usage

CasADi uses a MATLAB inspired "everything-is-a-matrix" type syntax, i.e., scalars are treated as 1-by-1 matrices and vectors as n -by-1 matrices. Furthermore, all matrices are *sparse* and stored in the compressed column format. For a symbolic framework like CasADi, working with a single sparse data type makes the tool easier to learn and maintain. Since the linear algebra operations are typically called only once, to construct symbolic expressions rather than to evaluate them numerically, the extra overhead of e.g. treating a scalar as a 1-by-1 sparse matrix is negligible.

The following code demonstrates loading CasADi into the workspace, creating two symbolic primitives $x \in \mathbb{R}^2$ and $A \in \mathbb{R}^{2 \times 2}$ and finally the creation of an expression for $e := A \sin(x)$:

<pre><i>% MATLAB/Octave</i> import casadi.* x = SX.sym('x', 2); A = SX.sym('A', 2, 2); e = A * sin(x); disp(e)</pre>	<pre><i># Python</i> from casadi import * x = SX.sym('x', 2) A = SX.sym('A', 2, 2) e = mtimes(A, sin(x)) print(e)</pre>
--	--

Output: @1=sin(x_0), @2=sin(x_1),
 [((A_0*@1) + (A_2*@2)), ((A_1*@1) + (A_3*@2))]

The output should be interpreted as the definition of two shared subexpressions, @1 := $\sin(x_0)$ and @2 := $\sin(x_1)$ followed by an expression for the resulting column

¹ One may run these snippets on a Windows, Linux or Mac platform by following the install instructions for CasADi 3.1 at <http://install31.casadi.org/>.

vector (2-by-1 matrix). The fact that CasADi expressions are allowed to contain shared subexpressions is essential to be able to solve large-scale problems and for CasADi to be able to implement AD as described in Section 2.6.

2.2 Graph representation—scalar expression type

In the code snippet above, we used CasADi's *scalar expression* type—*SX*—to construct a symbolic expression. Scalar in this context does not refer to the type itself—*SX* is a general sparse matrix type—but the fact that each nonzero element is defined by a sequence of scalar-valued operations. CasADi uses the compressed column storage (CCS) [4] format to store matrices, the same used to represent sparse matrices in MATLAB, but with the difference that in CasADi entries are allowed to be structurally nonzero but numerically zero as illustrated by the following:

<pre>% MATLAB/Octave C = SX.eye(2); C(1,1) = SX.sym('x'); C(2,1) = 0; disp(C)</pre>	<pre># Python C = SX.eye(2) C[0,0] = SX.sym('x') C[1,0] = 0 print(C)</pre>
<p>Output: $\begin{bmatrix} [x, & 00] , \\ [0, & 11] \end{bmatrix}$</p>	

Note the difference between structural zeros (denoted 00) and numerical zeros (denoted 0). The fact that symbolic matrices are always sparse in CasADi stands in contrast to e.g. MATLAB's Symbolic Math Toolbox, where expressions are always dense. Also note that CasADi has indices starting with 1 in MATLAB/Octave, but 0 in Python. In C++, CasADi also follows the index-0 convention.

When working with the *SX* type, expressions are stored as a directed acyclic graph (DAG) where each node—or *atomic operation*—is either:

- A symbolic primitive, created with *SX.sym* as above
- A constant
- A unary operation, e.g. *sin*
- A binary operation, e.g. ***, *+*

This relatively simple graph representation is designed to allow numerical evaluation with very little overhead, either in a virtual machine (Sect. 2.4) or in generated C code (Sect. 2.5). Each operation also has a chain-rule that can efficiently be expressed with the other atomic operations.

2.3 Graph representation—matrix expression type

There is a second expression type in CasADi, the *matrix expression* type—*MX*. For this type, each operation is a matrix operation; an expression such as $A + B$ where A and B are n -by- m matrices would result in a single *addition* operation, in contrast to up to mn scalar addition operations using the *SX* type. In the most general case, an *MX*

operation can have multiple matrix-valued inputs and return multiple matrix-valued outputs.

The choice to implement two different expression types in CasADi—and expose them both to the end user—is the result of a design compromise. It has proven difficult to implement an expression type that works efficiently both for e.g. the right-hand-side of an ODE, where minimal overhead is critical, and at the same time be able to represent the very general symbolic expressions that make up the NLP resulting from a direct multiple shooting discretization, which contains embedded ODE integrators.

The syntax of the `MX` type mirrors that of `SX`:

<code>% MATLAB/Octave</code>	<code># Python</code>
<code>x = MX.sym('x', 2);</code>	<code>x = MX.sym('x', 2)</code>
<code>A = MX.sym('A', 2, 2);</code>	<code>A = MX.sym('A', 2, 2)</code>
<code>e = A * sin(x);</code>	<code>e = mtimes(A, sin(x))</code>
<code>disp(e)</code>	<code>print(e)</code>

Output: `mac(A, sin(x), zeros(2x1))`

The resulting expression consists of two matrix valued symbolic primitives (for A and x , respectively), a 2-by-1 all-zero constant, a unary operation (`sin`) and a matrix multiply-accumulate operation, $\text{mac}(X_1, X_2, X_3) := X_3 + X_1 X_2$.

The choice of atomic operations for the `MX` type was made so that derivatives calculated either using the forward or reverse mode of algorithmic differentiation can be efficiently expressed using the same set of operations. The choice also takes into account that CasADi's `MX` virtual machine, cf. Sect. 2.4, supports *in-place* operations. This last fact explains why matrix multiply-accumulate was chosen as an atomic operation instead of the standard matrix multiplication; in practice, the operation performed is $X_3 := X_3 + X_1 X_2$.

A list of the most important atomic operations for the `MX` expression graph can be found in Table 1. The list also shows how the different atomic operations are interdependent under algorithmic differentiation operations. For example, reverse mode AD performed on the operation to retrieve an element of a matrix (operation 9 in the table) results in an operation to *assign* a quantity to a matrix element (operation 10 in the table). Note that the assignment operation is a two step procedure consisting of copying the existing matrix into a new variable before the actual assignment (or optionally, addition) takes place. The copy operation is typically eliminated in the virtual machine, cf. Sect. 2.4.

Some operations, e.g., the binary operation $A + B$, assume that the arguments have the same sparsity pattern (i.e., the same location of the nonzero elements). If this is not the case, CasADi inserts “projection” nodes into the expression graph during the construction.

A special type of atomic operation is a *function call* node. It consists of a call to a function object created at runtime. Importantly, there can be multiple calls to the same function object, which can keep the size of the expression graphs small. Function objects are covered in the following.

Table 1 Selected atomic operations for CasADi's MX type

	Operation	Definition or example	Interdependencies	
			AD forward	AD reverse
1	Constant	$y = 0, y = [1.2, 4]^T$ etc.		
2	Unary operation	$Y = \sin(X), Y = \sqrt{X}$ etc.		
3	Binary operation	$Y = X_1 + X_2, Y = X_1 * X_2$ etc.		4 ^a
4	Inner product	$y = \text{tr}(X_1^T X_2)$ ^b		3
5	Transpose	$Y = X^T$		
6	Multiply-accumulate	$Y = X_3 + X_1 X_2$		3
7	Reshape	(Changes dimension)		
8	Projection	(Changes sparsity)		
9	Submatrix access	$y = X(i, j)$		10
10	Submatrix assignment	$Y(i, j) = x_2$ with $Y = X_1$ elsewhere		9
11	Linear solve	$y = X_2^{-1} x_1$ or $y = X_2^{-T} x_1$	5, 6, 12, 13	5, 6, 12, 13
12	Horizontal concatenation	$Y = [x_1 \dots x_n]$		13
13	Horizontal split	$\{y_1, \dots, y_n\}$ such that $[y_1 \dots y_n] = X$		12
14	Function call	(See text)		

^a Dependency appears for e.g. $x_1 * X_2$, where x_1 is a scalar and X_2 is a matrix

^b Efficiently evaluated by elementwise multiplication and summation, $\text{sum}(X_1 * X_2)$

2.4 Function objects and virtual machines

The symbolic expressions in CasADi can be used to define *function objects*, class instances that behave like conventional functions but are created at runtime, cf. [3]. In addition to supporting numerical evaluation, CasADi function objects support symbolical evaluation, C code generation (Sect. 2.5) as well as derivative calculations (Sect. 2.6). They can be created by providing a display name and a list of input and output expressions:

```
% MATLAB/Octave          # Python
F = Function('F', {x,A}, {e});  F = Function('F', [x,A], [e])
```

which defines a function object with the display name “F” with two inputs (x and A) and one output (e), as defined in the previous code segments. Function objects can have an arbitrary number of inputs and outputs, each of which is a sparse matrix. Should an input—e.g. A above—contain structural zeros, the constructed function is understood not to depend on the corresponding part of the matrix.

The creation of a function object in CasADi essentially amounts to topologically sorting the expression graph, turning the directed acyclic graph (DAG) into an *algorithm* that can be evaluated. Unlike traditional tools for AD such as ADOL-C [58] or CppAD [24], there is no relation between the order in which expressions were created (i.e. a *tracing step*) and the order in which they appear in the sorted algorithm. Instead, CasADi uses a *depth-first search* to topologically sort the nodes of the DAG.

Given the sorted sequence of operations, CasADi implements two *register based virtual machines* (VMs), one for each graph representation. For the inputs and outputs of each operation, we assign an element (SX VM) or an interval (MX VM) from a work vector. This design contrasts to the *stack based* VM used in e.g. the AMPL Solver Library [45]. To limit the size of the work vector, the *live variable range* of each operation is analyzed and work vector elements or intervals are then reused in a last-in, first-out manner. When possible, operations are performed *in-place*, e.g., an operation for assigning the top left element of a matrix:

$$Y(i, j) = \begin{cases} x_2 & \text{if } i = 1 \text{ and } j = 1 \\ X_1(i, j) & \text{otherwise} \end{cases}$$

typically simplifies to just an element assignment, $X_1(1, 1) := x_2$, assuming that X_1 is not needed later in the algorithm.

Since MX expression graphs can contain calls to function objects, it is possible to define nested function objects. Encapsulating subexpressions used multiple times into separate function objects allows expression graphs to stay small and has implications for the memory use in the context of algorithmic differentiation, cf. Sect. 2.6.

Function objects in CasADi, as represented by the `Function` class, need not be defined by symbolic expressions as above. ODE/DAE integrators, solvers of nonlinear systems of equations and NLP solvers are examples of function objects in CasADi that are not explicitly defined by symbolic expressions. In many, but not all cases, these function objects are also automatically differentiable. We return to these classes in the following sections.

2.5 C code generation and just-in-time compilation

The VMs in CasADi are designed for high-speed and low overhead e.g. by avoiding memory allocation during numerical evaluation. In a framework such as CasADi, which is frequently used for rapid prototyping with many design iterations, fast VMs are important not only for numerical evaluation, but also for symbolic processing, which can make up a significant portion of the total solution time.

An alternative way to evaluate symbolic expressions in CasADi is to generate C code for the function objects. When compiled with the right compiler flags, the generated code can be significantly faster than CasADi's VMs. Since the generated code is self-contained C and has no dynamic memory allocation, it is suited to be deployed on embedded systems. The generated code is designed to be linked into a dynamically linked library, for static/dynamic linking (via a generated header file), to be called from the OS command-line (via a generated `main` entry point), or be called from MATLAB/Octave (via a generated `mexFunction` entry point).

The generated C code can be used for *just-in-time* compilation, which is supported either by using the system compiler or via an interface to the LLVM compiler framework with its C/C++ front-end Clang [79]. The latter is available for all platforms and is distributed with CasADi, meaning that the user does not need to have a binary compatible system compiler in order to use the generated C code.

2.6 Algorithmic differentiation

Algorithmic differentiation (AD)—also known as automatic differentiation—is a technique for efficiently calculating derivatives of functions represented as algorithms. For a function $y = f(x)$ with vector-valued x and y , the *forward mode* of AD provides a way to accurately calculate a Jacobian-times-vector product:

$$\hat{y} := \frac{\partial f}{\partial x} \hat{x}, \quad (1)$$

at a computational cost comparable to evaluating the original function $f(x)$. This definition naturally extends to a matrix-valued function $Y = F(X)$, by simply defining $x := \text{vec}(X)$ and $y := \text{vec}(Y)$, where $\text{vec}(\cdot)$ denotes stacking the columns of the matrix vertically. It also naturally generalizes further to the case when there are multiple matrix-valued inputs and outputs, which is the general case for functions in CasADi.

The *reverse mode* of AD, on the other hand, provides a way to accurately calculate a Jacobian-transposed-times-vector product:

$$\bar{x} := \left(\frac{\partial f}{\partial x} \right)^T \bar{y} \quad (2)$$

also at a computational cost comparable to evaluating the original function $f(x)$. In contrast to the forward mode, the reverse mode in general carries a larger, but often avoidable, memory overhead. This definition likewise naturally extends to the case when the function takes multiple matrix-valued inputs and multiple matrix-valued outputs.

Any implementation of AD works by breaking down a calculation into a sequence of atomic operations with known, preferably explicit, chain rules. For example, the forward mode AD rule for a matrix-matrix multiplication $Y = X_1 X_2$ is given by:

$$\hat{Y} = \hat{X}_1 X_2 + X_1 \hat{X}_2$$

and the reverse mode AD rule is given by:

$$\bar{X}_1 = \bar{Y} X_2^T; \quad \bar{X}_2 = X_1^T \bar{Y},$$

as shown in e.g. [55].

The forward and reverse modes thus offer two ways to efficiently and exactly calculate directional derivatives. Efficiently calculating the complete Jacobian, which can be large and sparse, is a considerably more difficult problem. For large problems, some heuristic that falls back on the above forward and/or reverse modes, is usually required. Higher order derivatives can be treated as special cases, or, which is the case here, by applying the AD algorithms recursively. The Hessian of a scalar-valued function is then simply calculated as the Jacobian of the gradient, preferably exploiting the symmetry of the Hessian.

CasADi implements AD using a source-code-transformation approach, which means that new symbolic expressions, using the same graph representation, are generated whenever derivatives are requested. Differentiable expressions for directional derivatives as well as large-and-sparse Jacobians and Hessians can be calculated.

The following demonstrates how to generate a new expression for the first column of a Jacobian using a Jacobian-times-vector product as well as an expression for the complete Jacobian:

<code>% MATLAB/Octave</code>	<code># Python</code>
<code>J1 = jtimes(e,x,[1;0]);</code>	<code>J1 = jtimes(e,x,[1,0])</code>
<code>J = jacobian(e,x);</code>	<code>J = jacobian(e,x)</code>

In the remainder of this section, we present the implementation of AD in CasADi, assuming that the reader is already familiar with AD. We refer to [5, Chapter 3] for a simple introduction and Griewank and Walther [60] or Naumann [106] for a more complete introduction to AD.

Directional derivatives

For a function object defined by a symbolic expression, CasADi implements the forward and reverse modes of AD by propagating symbolic seeds forward and backward through the algorithm respectively, resulting in new symbolic expressions that contain references to nodes of the expression graph of the non-differentiated function. Whenever a *function call* node is encountered, cf. Sect. 2.3, a new function object is generated for calculating directional derivatives. The generated function objects for the derivatives are *cached* in order to limit memory use. How a function class calculates directional derivatives is class-specific; e.g., an integrator node typically generates functions for directional derivatives by augmenting the ODE/DAE integration with its sensitivity equations.

If a symbolic expression consists of a large number of nodes, the evaluation of reverse mode derivatives may be costly in memory, since the intermediate results must be kept in memory and then accessed in reverse order. The CasADi user is responsible for avoiding such a memory blowup by breaking up large expressions into a hierarchy of smaller expressions, each encapsulated in a separate function object. Choosing a suitable hierarchy of function objects is equivalent to a *checkpointing strategy* [60, Chapter 12] in AD terminology, and as such comes at the price of a moderate increase in the number of floating point operations for reverse mode AD.

Calculation of complete Jacobians and Hessians

We use a graph coloring approach [50] to generate expressions for complete large and sparse Jacobians and Hessians. The idea of this approach is to reconstruct the Jacobian with a set of Jacobian-vector-products, i.e. directional derivatives. Using greedy graph coloring algorithms, we seek to find a set of vectors that is smaller than the naive choice of simply using the unit vectors, i.e. v_i being the i -th column of the identity matrix.

CasADi uses a heuristic to construct Jacobian and Hessian expressions. The heuristic uses a symmetry-exploiting *greedy, distance-2, star-coloring algorithm* [50, Algorithm 4.1], whenever it is known a priori that the resulting Jacobian is symmetric, in particular whenever a Hessian is being constructed. For asymmetric Jacobians, a *greedy, distance-2, unidirectional algorithm* [50, Algorithm 4.1] is attempted both column-wise (corresponding to forward mode AD) and row-wise (corresponding to reverse mode AD). Depending on the number of rows and columns of the Jacobian, one or the other is attempted first and the algorithm attempted second is interrupted prematurely if determined to need more colors, i.e. more directional derivatives. A factor α , by default 2, is introduced to take into account that a reverse mode AD sweep is usually slower than a forward mode AD sweep. We summarize the algorithm in Table 1.

The coloring algorithms in CasADi are used together with a *largest-first* preordering step [50, 136].

Algorithm 1 Heuristic to calculate complete Jacobians in CasADi

```

Calculate the Jacobian sparsity pattern
if  $J$  is symmetric (typically a Hessian) then
  Run a star-coloring algorithm.
else
  if The number of columns of  $J$  is fewer than  $\alpha$  times the number of rows then
    Run a column-wise unidirectional algorithm.
    Run a row-wise unidirectional algorithm, interrupting prematurely if the number of colors surpasses
    the column-wise algorithm with more than a factor  $1/\alpha$ .
  else
    Run a row-wise unidirectional algorithm.
    Run a column-wise unidirectional algorithm, interrupting prematurely if the number of colors sur-
    passes the row-wise algorithm with more than a factor  $\alpha$ .
  end if
end if

```

Jacobian sparsity pattern calculation

A priori knowledge of the sparsity pattern of the Jacobian is a precondition to be able to implement the above graph coloring approach. Obtaining this pattern for a generic expression is a nontrivial task and often turns out to be the most expensive step in the Jacobian construction.

CasADi uses the *bitvector approach* [54], which is essentially an implementation of forward or reverse mode AD using a single bit as a datatype, indicating whether a component of a Jacobian-vector-product is structurally zero. The bitvector approach can be implemented efficiently using bitwise operations on an unsigned integer datatype. CasADi uses the 64-bit `unsigned long long` datatype for this, meaning that up to 64 rows or columns can be calculated in a single sparsity propagation sweep.

For the large and sparse Jacobians and Hessians encountered in CasADi, where the number of rows and columns can be in the millions, performing tens of thousands of sparsity propagation sweeps to determine the Jacobian sparsity pattern can be pro-




Fig. 1 The sparsity pattern of the Hessian of the Lagrangian of an NLP arising from direct collocation

hibitively expensive. If nothing is known about the location of the nonzero elements, probabilistic methods as shown by Griewank and Mitev [59] have been proposed. For the type of Jacobians typically encountered in CasADi, both in simulation and optimization, the nonzeros are unlikely to be encountered in random locations. A more typical sparsity pattern is one which has large regions that do not have a single nonzero entry. An example of such a structured sparsity pattern, corresponding to the Hessian of the Lagrangian of an NLP arising from direct collocation, can be seen in Fig. 1.

To exploit this *block sparse* structure, CasADi implements a *hierarchical sparsity pattern calculation algorithm* based on graph coloring. The algorithm alternates between calculating successively less crude sparsity patterns and the same graph coloring algorithms used in the previous section.

In a first step, the rows and columns of the Jacobian are divided into 64 groups of similar size. The sparsity pattern propagation algorithm, either forward or reverse, is then executed yielding a coarse sparsity pattern. For the pattern in Fig. 1, this will result in either a block diagonal or block tridiagonal sparsity pattern, depending on how the blocks are chosen. Either case is amenable for graph coloring. After either symmetry exploiting or non-symmetry exploiting graph coloring, the process is repeated for a finer sparsity pattern, using one sparsity propagation sweep for each distinct color found in the coarse pattern.

This process is then repeated for successively finer sparsity patterns, until the the actual sparsity pattern is obtained.

Example 1 Assume that the Jacobian has dimension 100,000-by-100,000 and has a (a priori unknown) nonsymmetric tridiagonal sparsity pattern.

The proposed propagation algorithm first performs one sweep which results in a block tridiagonal pattern, with block sizes no larger than $\text{ceil}(100,000/64) = 1563$. Tridiagonal patterns can trivially be colored with 3 colors, regardless of dimension, meaning that the columns can be divided into 3 groups. For each color, we make one sweep (3 in total) to find a finer sparsity pattern, which will also be tridiagonal. At this point, the blocks are no larger than $\text{ceil}(1563/64) = 25$. The graph coloring algorithm is again executed and again resulting in 3 colors. Finally, with one sweep for each color, we obtain the true sparsity pattern.

For this example, the sparsity pattern is thus recovered in 7 sweeps, which can be compared with $\text{ceil}(100,000/64) = 1563$ sweeps needed for the naive algorithm, where 64 rows or columns are calculated in each sweep.

This proposed hierarchical sparsity pattern calculation algorithm is not efficient if the nonzero entries are spread out randomly and CasADi assumes that the user takes this into account when e.g. formulating a large NLP.

2.7 Control flow handling

A common question posed by CasADi users is how to handle expressions that involve flow control such as if-statements, for-loops and while-loops. Expressions containing flow control appear naturally in a range of applications, e.g. for physical models governed by different equations for different value ranges. Being able to calculate derivatives for such models, that are at least accurate in the *almost everywhere* sense, is essential for practical numerical optimization.

The approach chosen in CasADi is to support flow control by implementing concepts from *functional programming languages* such as Haskell.

Conditionals

Conditional expressions, which include switches and if-statements, can be expressed using a dedicated `Switch` function object class in CasADi. This construction, which is defined by a set of other function objects in CasADi, is defined by a vector of function objects corresponding to the different cases as well as the default case, all with the same input–output signature:

Since CasADi calculates derivatives of function objects by generating expressions for its directional derivatives, the derivative rules for the above class—for forward mode and reverse mode—can be defined as new conditional function objects defined by the corresponding derivative functions. Note that derivatives with respect to the first argument (c above) are zero almost everywhere.

An important special case of the `Switch` construct is if-then-else operations, for which $N = 1$.

Algorithm 2 A `Switch` function object in CasADi defined by (f_0, \dots, f_N)

```

input  $(c, x_0, \dots, x_{n-1})$ 
if  $\text{floor}(c) = 0$  then
   $(y_0, \dots, y_{m-1}) = f_0(x_0, \dots, x_{n-1})$ 
else if ... then
  ...
else if  $\text{floor}(c) = N - 1$  then
   $(y_0, \dots, y_{m-1}) = f_{N-1}(x_0, \dots, x_{n-1})$ 
else
   $(y_0, \dots, y_{m-1}) = f_N(x_0, \dots, x_{n-1})$ 
end if
return  $(y_0, \dots, y_{m-1})$ 

```

Maps

Another readily differentiable concept from functional programming is a *map*. A map in this context is defined as a function being evaluated multiple times with different arguments. This evaluation can be performed serially or in parallel.

The `Map` function object class in CasADi allows users to formulate maps. This function object takes a horizontal concatenation of all inputs and returns a horizontal concatenation of all outputs:

Algorithm 3 A `Map` function object in CasADi defined by f and N

```

input  $(X_0, \dots, X_{n-1})$ 
for  $i = 0, \dots, N - 1$  do
   $(y_{i,0}, \dots, y_{i,m-1}) = f(x_{i,0}, \dots, x_{i,n-1})$ 
end for
return  $(Y_0, \dots, Y_{m-1})$ 

```

where $X_j := [x_{0,j}, \dots, x_{N-1,j}]$, $j = 0, \dots, m - 1$ and $Y_k := [y_{0,k}, \dots, y_{N-1,k}]$, $k = 0, \dots, n - 1$.

3 Implicitly defined differentiable functions

Optimization problems may contain quantities that are defined implicitly. An important example of this, and one of the motivations to write CasADi in the first place, is the direct multiple shooting method, where the NLP contains embedded solvers for initial-value problems in differential equations. We will showcase this method in Sect. 5.4. Another example is a dynamic system containing algebraic loops, which can be made explicit by embedding a root-finding solver.

In the following, we discuss how certain implicitly defined functions can be embedded into symbolic expressions, but still have their derivative and sparsity information generated automatically and efficiently.

3.1 Linear systems of equations

As shown in [55], the solution to a linear system of equations $y = X_2^{-1} x_1$ has forward and reverse mode AD rules defined by:

$$\hat{y} = X_2^{-1} (\hat{x}_1 - \hat{X}_2 y); \quad \bar{x}_1 = X_2^{-T} \bar{y}; \quad \bar{X}_2 = -\bar{x}_1 y^T$$

Apart from standard operations such as matrix multiplications, the directional derivatives can thus be expressed using a linear solve with the same linear system as the nondifferentiated expression. In the reverse mode case, a linear solve with the transposed matrix is needed.

CasADi supports linear system of equations of this form through a linear solver abstract base class, which is an oracle class able to factorize and solve dense or sparse linear systems of equations. The solve step supports optional transposing, as required by the reverse mode of AD. The linear solver class is a plugin class in CasADi, and leaves to the derived class—typically an interface to a third-party linear solver—to actually perform the factorization and solution.

The linear solver instances are embedded into CasADi's MX expression graphs using a dedicated *linear solve* node, which is similar to the more generic function call node introduced in Sect. 2.3. The linear solve node implements the above derivative AD rules, using MX's nodes for horizontal split and concatenation to be able to reuse the same factorization for multiple directional derivatives, i.e. multiple right-hand-sides.

Sparsity pattern propagation for the linear solve node was implemented by first making a block-triangular reordering of the rows and columns, exposing both unidirectional and bidirectional dependencies between the elements. The reordering is calculated only once for each sparsity pattern and then cached.

At the time of this writing, the linear solver plugins in CasADi—which may impose additional restrictions such as symmetry or positive definiteness—included CSparse [28] (sparse LU, sparse QR and sparse Cholesky decompositions), MA27 [65] (frontal method) and LAPACK [4] (dense LU and dense QR factorizations).

3.2 Nonlinear systems of equations

A more general implicitly defined function is the solution of a root-finding problem:

$$g(y, x) = 0 \Leftrightarrow y = f(x). \quad (3)$$

If regularity conditions are satisfied, in particular the existence and invertibility of the Jacobian $\frac{\partial g}{\partial y}$, the problem is well-posed and its Jacobian is given by the *implicit function theorem*:

$$\frac{\partial f}{\partial x} = - \left(\frac{\partial g}{\partial y} \right)^{-1} \frac{\partial g}{\partial x}$$

which readily gives the forward and reverse AD propagation rules of (3)

$$\hat{y} = - \left(\frac{\partial g}{\partial y} \right)^{-1} \frac{\partial g}{\partial x} \hat{x}; \quad \bar{x} = - \left(\frac{\partial g}{\partial x} \right)^T \left(\frac{\partial g}{\partial y} \right)^{-T} \bar{y}. \quad (4)$$

For the forward mode, the problem of calculating directional derivatives for the implicitly defined function $f(x)$ is reduced to the problem of calculating forward mode directional derivatives of the residual function $g(y, x)$ followed by a linear solve. For the reverse mode, the calculation involves a linear solve for the transposed system followed by a reverse mode directional derivative for the residual function.

CasADi supports the above via so-called `rootfinder` function objects. These use a generalization of (3) which includes multiple matrix-valued input parameters and a set of auxiliary outputs:

$$\begin{pmatrix} g_0(y_0, x_1, \dots, x_{n-1}) \\ g_1(y_0, x_1, \dots, x_{n-1}) - y_1 \\ \vdots \\ g_{m-1}(y_0, x_1, \dots, x_{n-1}) - y_{m-1} \end{pmatrix} = 0 \Leftrightarrow (y_0, \dots, y_{m-1}) = f(x_0, \dots, x_{n-1}). \quad (5)$$

where x_0 has been introduced as an *initial guess* for y_0 . Note that the derivative with respect x_0 is zero almost everywhere.

Like linear solvers, root-finders are implemented using a plugin-design. In the base class, rules for derivative calculation and sparsity pattern propagation are defined whereas solving the actual nonlinear system of equations is delegated to a derived class, typically in the form of an interfaced third-party tool. Root-finding objects in CasADi can be differentiated an arbitrary number of times since both its forward and reverse mode directional derivatives can be expressed by (arbitrarily differentiable) CasADi constructs, namely directional derivatives for the residual function and the linear solver operation treated in Sect. 3.1. The linear solver is typically the same as the linear solver used in a Newton method for the nonlinear system of equations. By default, the CSparse [28] plugin is used, but this can be changed by passing an option to the `rootfinder` constructor.

3.3 Initial-value problems in ODE and DAE

Certain methods for optimal control, including direct multiple shooting and direct single shooting, will result in optimization problem formulations that require solving initial-value problems (IVP) in ODEs or DAEs. Since the integrator calls appear in the constraint and/or objective functions of an NLP, we need ways to calculate first and preferably second order derivative information.

The CasADi constructs introduced until now can be used to define either explicit or implicit fixed step-step integrator schemes. For example, a Runge–Kutta 4 (RK4) scheme can be implemented with less than 10 lines of code using CasADi's `MX` type, and derivatives can be generated automatically to any order. Similarly, an implicit

fixed-step scheme, such as a collocation method, can be implemented using CasADi’s `rootfinder` functionality, described in Sect. 3.2.

More advanced integrator schemes, such as backwards difference formula (BDF) methods with variable order and/or adaptive step-size, cannot be handled with this approach. Compared to a fixed-step integrator scheme, an adaptive scheme often results in fewer steps for the same accuracy and the user is relieved from choosing appropriate step-sizes. CasADi’s `integrator` functionality enables the user to embed solvers of initial-value problems in ODEs or DAEs and have derivatives to any order calculated exactly using state-of-the-art codes such as those in the SUNDIALS suite [70]. CasADi’s `integrator` objects solve problems of the following form:

$$\begin{aligned}
 & f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_z} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_r} \times \mathbb{R}^{n_s} \times \mathbb{R}^{n_v} \rightarrow \mathbb{R}^{n_x} \times \mathbb{R}^{n_z} \times \mathbb{R}^{n_q} \times \mathbb{R}^{n_r} \times \mathbb{R}^{n_s} \times \mathbb{R}^{n_p}, \\
 & (x_0, z_0, u, r_T, s_T, v) \mapsto (x(T), z(T), q(T), r(0), s(0), p(0)) \\
 & \left\{ \begin{array}{ll} \dot{x}(t) = \phi(x(t), z(t), u), & x(0) = x_0 \\ 0 = \theta(x(t), z(t), u), & z_0 \text{ initial guess for } z(0) \\ \dot{q}(t) = \psi(x(t), z(t), u), & q(0) = 0 \\ -\dot{r}(t) = \phi^*(x(t), z(t), u, r(t), s(t), v), & r(T) = r_T \\ 0 = \theta^*(x(t), z(t), u, r(t), s(t), v), & s_T \text{ initial guess for } s(T) \\ -\dot{p}(t) = \psi^*(x(t), z(t), u, r(t), s(t), v), & p(T) = 0 \end{array} \right. \quad t \in [0, T]
 \end{aligned} \tag{6}$$

The problem consists of two semi-explicit DAEs with initial and terminal constraints, respectively, and both with support for calculation of quadratures. The second DAE is coupled to the solution trajectory of the first DAE. We impose the additional requirement that $\theta^*(x, z, u, r, s, v)$, $\phi^*(x, z, u, r, s, v)$ and $\psi^*(x, z, u, r, s, v)$ are affine in r, s and v . Initial guesses for $z(0)$ and $s(T)$ are included for efficiency, robustness and to ensure solution uniqueness.

This integrator formulation has two major advantages in the context of optimization and sensitivity analysis. Firstly, it is general enough to handle most industry relevant simulation problems and the quadrature functionality allows integral terms in objective functions to be calculated efficiently using quadrature formulas. Secondly, it can be shown [5] that both forward and reverse mode directional derivatives of this problem can be calculated efficiently by solving a problem that has exactly the same structure.

By performing the differentiation repeatedly, we can calculate derivatives to any order, using an appropriate mix of forward and adjoint sensitivity analysis. In particular, we can perform forward-over-adjoint sensitivity analysis for efficient Hessian calculation. A potential drawback, which is inherent with this so-called *variational approach to sensitivity analysis*, is that when used inside a gradient-based optimization code, the calculated derivatives may not be consistent with the nondifferentiated function evaluation, due to differences in time discretization.

Like linear solvers and root-finding solvers, integrators are implemented using a plugin design in CasADi. The base class implements the rules for differentiation and sparsity pattern propagation and the derived class, which can be an interface to a third-party tool, performs the actual solution. Integrator plugins in CasADi include IDAS and CVODES from the SUNDIALS suite [70], a fixed step-size RK4 code, and an implicit Runge–Kutta code implementing Legendre or Radau collocation.

4 Optimization

Just like conventional algebraic modeling languages, CasADi combines support for modeling with support for mathematical optimization. Two classes of optimization problems are supported; nonlinear programs (NLPs) and conic optimization problems. The latter class includes both linear programs (LPs) and quadratic programs (QPs). The actual solution typically takes place in a derived class, and may use a tool distributed with CasADi or be an interface to a third-party solver. The role of CasADi is to extract information about structure and generate the required derivative information as well as to provide a common interface for all solvers.

4.1 Nonlinear programming

CasADi uses the following formulation for a nonlinear program (NLP):

$$\begin{aligned} & \underset{x, p}{\text{minimize}} && f(x, p) \\ & \text{subject to} && \underline{x} \leq x \leq \bar{x}, \quad p = \bar{p}, \quad \underline{g} \leq g(x, p) \leq \bar{g}. \end{aligned} \quad (7)$$

This is a parametric NLP where the objective function $f(x, p)$ and the constraint function $g(x, p)$ depend on the decision variable x and a known parameter p . For equality constraints, the variable bounds $[\underline{x}, \bar{x}]$ or constraint bounds $[\underline{g}, \bar{g}]$ are equal.

The solution of (7) yields a primal (x, p) and a dual $(\lambda_x, \lambda_p, \lambda_g)$ solution, where the Lagrange multipliers are chosen to be consistent with the following definition of the Lagrangian function:

$$\mathcal{L}(x, p, \lambda_x, \lambda_p, \lambda_g) := f(x, p) + \lambda_x^T x + \lambda_p^T p + \lambda_g^T g(x, p). \quad (8)$$

This formulation drops all terms that do not depend on x or p and uses the same multipliers (but with different signs) for the inequality constraints according to

$$\lambda_{\bar{x}}(x - \bar{x}) + \lambda_{\underline{x}}(\underline{x} - x) = \underbrace{(\lambda_{\bar{x}} - \lambda_{\underline{x}})^T x}_{:=\lambda_x} - \underbrace{\lambda_{\bar{x}}^T \bar{x} + \lambda_{\underline{x}}^T \underline{x}}_{\text{ignored}}$$

and equivalently for $g(x, p)$. Note that $\lambda_{\bar{x}}(i)$ and $\lambda_{\underline{x}}(i)$ cannot both be positive and that the dropped terms do not appear in the KKT conditions of (7).

In this NLP formulation, a strictly positive multiplier signals that the corresponding *upper* bound is active and vice versa. Furthermore, λ_p is the *parametric sensitivity* of the objective with respect to the parameter vector p .

Table 2 lists the available NLP solver plugins at the time of this writing. The list includes both open-source solvers that are typically distributed along with CasADi and commercial solvers that require separate installation. The table attempts to make a rough division of the plugins between solvers suitable for very large but very sparse NLPs, e.g., arising from direct collocation, and large and structured NLPs, e.g. arising from direct multiple shooting, cf. Sect. 4.3. The plugin `'sqpmethod'` corresponds

Table 2 NLP solver plugins in CasADi 3.1

Solver	Plugin name	License	Method	Problem
IPOPT [134]	<code>'ipoppt'</code>	EPL	Nonlinear interior point ^b	NLP
BONMIN [9]	<code>'bonmin'</code>	EPL	Nonlinear interior point ^b	MINLP
KNITRO [21]	<code>'knitro'</code>	Commercial	Multiple ^{b,c}	MINLP
WORHP [20]	<code>'worhp'</code>	Commercial ^a	SQP ^b	NLP
SNOPT [56]	<code>'snopt'</code>	Commercial	SQP ^c	NLP
<code>casadi::Blocksqp^d</code>	<code>'blocksqp'</code>	zlib	SQP ^c	NLP
<code>casadi::Sqpmethod</code>	<code>'sqpmethod'</code>	LGPL	SQP boilerplate code	NLP
<code>casadi::Scpgen</code>	<code>'scpgen'</code>	LGPL	SQP ^c	NLP

^a Free for academic use

^b Suitable for large and sparse NLPs

^c Suitable for structured NLPs

^d Fork of blockSQP [75]

to a “vanilla” SQP method, mainly intended to serve as a boilerplate code for users who intend to develop their own solver codes. A subset of the solvers supports mixed-integer nonlinear programming (MINLP) formulations, where a subset of the decision variables are restricted to take integer values.

4.2 Conic optimization

When the objective function $f(x, p)$ is quadratic in x and the constraint function $g(x, p)$ is linear in x , (7) can be posed as a quadratic program (QP) of the form

$$\begin{aligned} & \underset{x}{\text{minimize}} && \frac{1}{2}x^T H x + g^T x \\ & \text{subject to} && \underline{x} \leq x \leq \bar{x}, \quad \underline{a} \leq A x \leq \bar{a}. \end{aligned} \quad (9)$$

CasADi supports solving QPs, formulated either as (7) or as (9). In the former case, AD is used to reformulate the problem in form (9), automatically identifying the sparse matrices H and A as well as the vectors g , \underline{a} and \bar{a} .

As with NLP solvers, the solution of optimization problems takes place in one of CasADi’s *conic* solver plugins, listed in Table 3. Some plugins impose additional restrictions on (9); *linear* programming solvers require that the H term is zero and most interfaced QP solvers require H to be positive semi-definite. A subset of the solvers supports mixed-integer formulations. Note that mixed integer quadratic programming (MIQP) is a superset of QP and that QP is a superset of LP. Future versions of CasADi may allow more generic conic constraints such as SOCP and SDP [18].

4.3 Sparsity and structure exploitation

NLPs and QPs arising from transcription of optimal control problems are often either *sparse* or *block-sparse*. A sparse QP in this context is one where matrices—i.e. A and

Table 3 Conic solver plugins in CasADi 3.1

Solver	Plugin name	License	Method	Problem
qpOASES [42]	<code>'qpoases'</code>	LGPL	Active-set	QP ^b
CPLEX [74]	<code>'cplex'</code>	Commercial	Multiple ^d	MIQP
GUROBI [64]	<code>'gurobi'</code>	Commercial	Multiple ^d	MIQP
CLP [23]	<code>'clp'</code>	CPL	Simplex	LP
OOQP [52]	<code>'ooqp'</code>	Open-source	Interior point ^d	QP
HPMPC [48]	<code>'hmpc'</code>	LGPL	Interior point	QP ^c

^a Free for academic use

^b Supports non-convex QPs

^c Supports integer decision variables

^d For large and sparse problems

^e For structured problems

H in (9)—have few enough nonzero entries per row or column to be handled efficiently by general sparse linear algebra routines. For a sparse NLP, the same applies to the matrices that arise from the linearization of the KKT conditions, i.e. to $\frac{\partial g}{\partial x}$ and $\nabla_x^2 \mathcal{L}$ in (7) and (8). A direct collocation type OCP transcription will typically result in sparse NLPs or QPs, provided that the Jacobian of the ODE right-hand-side function (or DAE residual function) is sufficiently sparse. Several tools exist that can handle these problems efficiently, cf. Tables 2 and 3. These solvers generally rely on sparse direct linear algebra routines, e.g. from the HSL Library [65].

A direct multiple shooting type transcription, on the other hand, will typically result in NLPs or QPs that have dense sub-blocks and hence an overall block-sparse pattern, with too many nonzero entries to be handled efficiently by general sparse linear algebra routines. *Condensing* [82] in combination with a dense solver such as qpOASES [42] is known to work well for certain structured QPs, in particular when the time horizon is short and the control dimension small relative to the state dimension. Other QPs can be solved efficiently with tools such as FORCES [34], qpDUNES [46] and HPMPC [48]. The lifted Newton method [1] is a generalization of the condensing approach to NLPs and nonlinear root-finding problems. We refer to [76] for a comprehensive treatment of structured QPs and NLPs.

At the time of this writing, CasADi supported one structured QP solver, HPMPC [48], and two structured NLP solvers, Scpgen and blockSQP. Scpgen [5] is an implementation of the lifted Newton method using CasADi's AD framework and blockSQP [75], incorporated into CasADi in modified form, can handle NLPs with a block-diagonal Hessian matrices.

5 Tutorial examples

In the following, we showcase the CasADi syntax and usage paradigm through a series of tutorial examples of increasing complexity. The first two examples introduce the optimization modeling approach in CasADi, which differs from that of conventional

algebraic modeling languages such as AMPL, GAMS, JuMP or Pyomo. In Sect. 5.3, we demonstrate the automatic ODE/DAE sensitivity analysis in CasADi and finally, in Sect. 5.4, we combine the tools introduced in the previous examples in order to implement the *direct multiple shooting* method.

We will use a syntax corresponding to CasADi version 3.1 in both MATLAB/Octave and Python, side-by-side. The presentation attempts to convey a basic understanding of what modeling in CasADi entails. For a more comprehensive and up-to-date introduction to CasADi, needed to understand each line of the example scripts, we refer to the user guide [6].

5.1 An unconstrained optimization problem

Let us start out by finding the minimum of Rosenbrock's banana-valley function:

$$\underset{x, y}{\text{minimize}} \quad x^2 + 100 \left(y - (1 - x)^2 \right)^2 \quad (10)$$

By inspection, we can see that its unique solution is (0, 1). The problem can be formulated and solved with CasADi as follows:

<code>% MATLAB/Octave</code>	<code># Python</code>
<code>import casadi.*</code>	<code>from casadi import *</code>
<code>% Symbolic representation</code>	<code># Symbolic representation</code>
<code>x = SX.sym('x');</code>	<code>x = SX.sym('x')</code>
<code>y = SX.sym('y');</code>	<code>y = SX.sym('y')</code>
<code>z = y - (1 - x)^2;</code>	<code>z = y - (1 - x)**2</code>
<code>f = x^2 + 100 * z^2;</code>	<code>f = x**2 + 100 * z**2</code>
<code>P = struct('x', [x; y], 'f', f);</code>	<code>P = dict(x=vertcat(x, y), f = f)</code>
<code>% Create solver instance</code>	<code># Create solver instance</code>
<code>F = nlpsol('F', 'ipopt', P);</code>	<code>F = nlpsol('F', 'ipopt', P)</code>
<code>% Solve the problem</code>	<code># Solve the problem</code>
<code>r = F('x0', [2.5 3.0])</code>	<code>r = F(x0 = [2.5, 3.0])</code>
<code>disp(r.x)</code>	<code>print(r['x'])</code>

The solution consists of three parts. Firstly, constructing a symbolic representation of the problem in the form of a MATLAB/Octave `struct` or a Python `dict`. A naming scheme, consistent with (7), is used for the different fields. The variable `z` is an intermediate expression; more complex models will contain a large number of such expressions. Secondly, a solver instance is created, here using IPOPT. During this step, the AD framework is invoked to generate a set of solver-specific functions for numerical evaluation, here corresponding to the cost function, its gradient, and its Hessian. Finally, the solver instance is evaluated numerically in order to obtain

the optimal solution. We pass the initial guess (2.5, 3.0) as an input argument. Other inputs are left at their default values, e.g. the bounds on x are $\underline{x} = -\infty$ and $\bar{x} = \infty$.

The above script converges to the optimal solution in 26 iterations. The total solution time on a MacBook Pro is in the order of 0.02 s.

5.2 Nonlinear programming example

Let us reformulate (10) as a constrained optimization problem, introducing a decision variable corresponding to z above:

$$\begin{aligned} & \text{minimize} && x^2 + 100z^2 \\ & && x, y, z \\ & \text{subject to} && z + (1 - x)^2 - y = 0. \end{aligned} \quad (11)$$

The problem can be formulated and solved with CasADi as follows:

<i>% MATLAB/Octave</i>	<i># Python</i>
import casadi.*	from casadi import *
<i>% Formulate the NLP</i>	<i># Formulate the NLP</i>
x = SX.sym('x');	x = SX.sym('x')
y = SX.sym('y');	y = SX.sym('y')
z = SX.sym('z');	z = SX.sym('z')
f = x^2+100*z^2;	f = x**2+100*z**2
g = z+(1-x)^2-y;	g = z+(1-x)**2-y
P = struct('f',f,'g',g,...	P = dict(f = f,g = g,\
'x',[x;y;z]);	x = vertcat(x,y,z))
<i>% Create solver instance</i>	<i># Create solver instance</i>
F = nlpsol('F','ipopt',P);	F = nlpsol('F','ipopt',P)
<i>% Solve the problem</i>	<i># Solve the problem</i>
r = F('x0',[2.5 3.0 0.75],...	r = F(x0 = [2.5,3.0,0.75],\
'ubg',0,'lbg',0);	ubg = 0, lbg = 0)
disp (r.x)	print (r['x'])

We impose the equality constraint by setting the upper and lower bound of g to 0 and use the (2.5, 3.0, 0.75) as an initial value, consistent with the initial guess for the unconstrained formulation. The above script converges to the optimal solution in 10 iterations taking around 0.01 s.

Notice how *lifting* the optimization problem to a higher dimension like this resulted in faster local convergence of IPOPT's Newton-type method. This behavior can often be observed for structurally complex nonlinear problems as discussed in e.g. [1]. This faster local convergence is one of the advantages of the direct multiple shooting method, which we will return to in Sect. 5.4.

5.3 Automatic sensitivity analysis example

We now shift our attention to simulation and sensitivity analysis using CasADi's `integrator` objects introduced in Sect. 3.3. Consider the following initial-value problem in ODE corresponding to a Van der Pol oscillator:

$$\begin{cases} \dot{x}_1 = (1 - x_2^2) x_1 - x_2 + p, & x_1(0) = 0 \\ \dot{x}_2 = x_1, & x_2(0) = 1 \end{cases} \quad (12)$$

With p fixed to 0.1, we wish to solve for $x_f := x(1)$. This can be solved as follows:

<i>% MATLAB/Octave</i>	<i># Python</i>
<code>import casadi.*</code>	<code>from casadi import *</code>
<i>% Formulate the ODE</i>	<i># Formulate the ODE</i>
<code>x = SX.sym('x',2);</code>	<code>x = SX.sym('x',2)</code>
<code>p = SX.sym('p');</code>	<code>p = SX.sym('p')</code>
<code>z = 1-x(2)^2;</code>	<code>z = 1-x[1]**2</code>
<code>f = [z*x(1)-x(2)+p;x(1)];</code>	<code>f = vertcat(z*x[0]-x[1]+p,\</code>
<code>dae = struct('x',x,'p',p,...</code>	<code> </code>
<code> 'ode',f);</code>	<code> x[0])</code>
	<code>dae = dict(x = x,p = p,ode = f)</code>
<i>% Create integrator</i>	<i># Create integrator</i>
<code>op = struct('t0',0,'tf',1);</code>	<code>op = dict(t0 = 0,tf = 1)</code>
<code>F = integrator('F',...</code>	<code>F = integrator('F',\</code>
<code> 'cvodes',dae,op);</code>	<code> 'cvodes',dae,op)</code>
<i>% Integrate</i>	<i># Integrate</i>
<code>r = F('x0',[0,1],'p',0.1);</code>	<code>r = F(x0 = [0,1],p = 0.1)</code>
<code>disp(r.xf)</code>	<code>print(r['xf'])</code>

As for the optimization examples above, the solution consists of three parts; construction of a symbolic representation of the problem, creating a solver instance, and evaluating this solver instance in order to obtain the solution. In the scripts above, we used CVODES from the SUNDIALS suite [70] to solve the initial value problem, which implements a variable step-size, variable-order backward differentiation formula (BDF) method.

Since F in the above scripts is a differentiable CasADi function, as described in Sect. 3.3, we can automatically generate derivative information to any order. For example, the Jacobian of $x(1)$ with respect to $x(0)$ can be calculated as follows:

```

% Create Jacobian function      # Create Jacobian function
D = F.factory('D',...         D = F.factory('D',\
{'x0','p'},{'jac:xf:x0'});   ['x0','p'],['jac:xf:x0'])

% Sensitivity analysis         # Sensitivity analysis
r = D('x0',[0,1],'p',0.1);   r = D(x0 = [0,1],p = 0.1)
disp(r.jac_xf_x0)             print(r['jac_xf_x0'])
    
```

The automatic sensitivity analysis is often invoked indirectly, when using a gradient-based optimization solver, as the next example shows.

5.4 The direct multiple shooting method

By combining the nonlinear programming example in Sect. 5.2 with the embeddable integrator in Sect. 5.3, we can implement the direct multiple shooting method by Bock and Plitt [15, 16]. We will consider a simple OCP with the same IVP as in (12), but reformulated as a DAE and with p replaced by a time-varying control u :

$$\text{minimize}_{x(\cdot), z(\cdot), u(\cdot)} \int_0^T (x_1(t)^2 + x_2(t)^2 + u(t)^2) dt \tag{13}$$

$$\text{subject to} \begin{cases} \dot{x}_1(t) = z(t)x_1(t) - x_2(t) + u(t) \\ \dot{x}_2(t) = x_1(t) \\ 0 = x_2(t)^2 + z(t) - 1 \\ -1.0 \leq u(t) \leq 1.0, \quad x_1(t) \geq -0.25 \end{cases} \quad t \in [0, T] \tag{14}$$

$$x_1(0) = 0, \quad x_2(0) = 1 \tag{15}$$

where $x(\cdot) \in \mathbb{R}^2$ is the (differential) state, $z(\cdot) \in \mathbb{R}$ is the algebraic variable and $u(\cdot) \in \mathbb{R}$ is the control. We let $T = 10$.

Our goal is to transcribe the OCP (15) to a problem of form (7). In the direct approach, the first step in this process is a parameterization of the control trajectory. For simplicity, we assume a uniformly spaced, piecewise constant control trajectory:

$$u(t) := u_k \quad \text{for } t \in [t_k, t_{k+1}), \quad k = 0, \dots, N - 1 \quad \text{with } t_k := kT/N.$$

With the control fixed over one interval, we can use an integrator to reformulate the problem from continuous time to discrete time. This can be done as in Sect. 5.3. Since we now have a DAE, we introduce an algebraic variable z and the corresponding algebraic equation g and use IDAS instead of CVODES. We also introduce a quadrature for calculating the contributions to the cost function.

<pre><i>% MATLAB/Octave</i> import casadi.* <i>% Formulate the DAE</i> x = SX.sym('x',2); z = SX.sym('z'); u = SX.sym('u'); f = [z*x(1)-x(2)+u;x(1)]; g = x(2)^2+z-1; h = x(1)^2+x(2)^2+u^2; dae = struct('x',x,'p',u,... 'ode',f,... 'z',z,'alg',g,'quad',h); <i>% Create solver instance</i> T = 10; <i>% end time</i> N = 20; <i>% discretization</i> op = struct('t0',0,'tf',T/N); F = integrator('F',... 'idas',dae,op);</pre>	<pre><i># Python</i> from casadi import * <i># Formulate the DAE</i> x = SX.sym('x',2) z = SX.sym('z') u = SX.sym('u') f = vertcat(z*x[0]-x[1]+u,\ x[0]) g = x[1]**2+z-1 h = x[0]**2+x[1]**2+u**2 dae = dict(x = x,p = u,ode = f,\ z = z,alg = g,quad = h) <i># Create solver instance</i> T = 10. <i># end time</i> N = 20 <i># discretization</i> op = dict(t0=0,tf=T/N) F = integrator('F',\ 'idas',dae,op)</pre>
---	---

Our next step is to construct a symbolic representation of the NLP. We will use the following formulation:

$$\begin{aligned} & \text{minimize } J(w) \\ & \text{subject to } G(w) = 0, \quad \underline{w} \leq w \leq \bar{w} \end{aligned} \quad (16)$$

For this we start with an empty NLP and add a decision variable corresponding to the initial conditions:

<pre><i>% Empty NLP</i> w={}; lbw=[]; ubw=[]; G={}; J=0;</pre>	<pre><i># Empty NLP</i> w = []; lbw = []; ubw = [] G = []; J = 0</pre>
<pre><i>% Initial conditions</i> Xk = MX.sym('X0',2); w{end+1} = Xk; lbw = [lbw;0;1]; ubw = [ubw;0;1];</pre>	<pre><i># Initial conditions</i> Xk = MX.sym('X0',2) w+ = [Xk] lbw+ = [0,1] ubw+ = [0,1]</pre>

Inside a for loop, we introduce decision variables corresponding to each control interval and the state at the end of each interval:

<pre> for k=1:N % Local control Uname = ['U' num2str(k-1)]; Uk = MX.sym(Uname); w{end+1} = Uk; lbw = [lbw; -1]; ubw = [ubw; 1]; % Call integrator Fk = F('x0',Xk,'p',Uk); J = J+Fk.qf; % New local state Xname = ['X' num2str(k)]; Xk = MX.sym(Xname,2); w{end+1} = Xk; lbw = [lbw; -.25; -inf]; ubw = [ubw; inf; inf]; % Continuity constraint G{end+1} = Fk.xf-Xk; end </pre>	<pre> for k in range(1,N+1): # Local control Uname = 'U'+str(k-1) Uk = MX.sym(Uname) w+ = [Uk] lbw+ = [-1] ubw+ = [1] # Call integrator Fk = F(x0 = Xk,p = Uk) J+ = Fk['qf'] # New local state Xname = 'X'+str(k) Xk = MX.sym(Xname,2) w+ = [Xk] lbw+ = [-.25,-inf] ubw+ = [inf, inf] # Continuity constraint G+ = [Fk['xf']-Xk] </pre>
--	--

With symbolic expressions for (16), we can use CasADi's fork of blockSQP [75] to solve this block structured NLP, as in Sect. 5.2:

<pre> % Create NLP solver nlp = struct('f',J,... 'g',vertcat(G{:}),... 'x',vertcat(w{:})); S = nlpsol('S',... 'blocksqp',nlp); % Solve NLP r = S('lbx',lbw,'ubx',ubw,... 'x0',0,'lbg',0,'ubg',0); disp(r.x); </pre>	<pre> # Create NLP solver nlp = dict(f = J,\ g = vertcat(*G),\ x = vertcat(*w)) S = nlpsol('S',\ 'blocksqp',nlp) # Solve NLP r = S(lbx = lbw,ubx = ubw,\ x0 = 0,lbg = 0,ubg = 0) print(r['x']) </pre>
---	---

5.5 Further examples

More examples on CasADi usage can be found in CasADi's example collection. These examples include other OCP methods such as direct single shooting, direct collocation as well as two indirect methods and a dynamic programming method for comparison.

6 Applications of CasADi

Since the first release of CasADi in 2011, the tool has been used to teach optimal control in graduate level courses, to solve optimization problems in science and engineering as well as to implement new algorithms and software. An overview of applied research using CasADi, as of early 2017, is presented in the following.

6.1 Optimization and simulation in science and engineering

In energy research, applications include the exploitation [26,77,101] and transport [130] of fossil fuels, power-to-gas systems [19], control of combined-cycle [78] and steam [10] power plants, solar thermal power plants [113], production models [85] and control [61] of classical wind-turbines, design and control of airborne wind energy systems [38,62,73,83,87,103], MEMS energy harvesters [80], design of geothermal heat pumps [135], thermal control of buildings [29,104], electrical grids [40,117], electrical grid balancing [39], and price arbitrage on the energy market [30].

In the automotive industry, applications include design and operation of drivetrains [13,107,108], electrical power systems [120], control of combustion engines [119], research towards self-driving cars [12,69,93,125], traffic control [118], and operation of a driving simulator [129].

In the process industries, applications include control [71,86,90], optimal experimental design [88,105], and parameter estimation [72] of (bio-)chemical reactors.

In robotics research, applications include control of agricultural robots [124], remote sensing of icebergs with UAVs [7,68], time-optimal control of robots [132], motion templates for robot-human interaction [133], motion planning of robotic systems with contacts [49,98], and multi-objective control of complex robots [84].

Further assorted applications include estimation in systems biology [121,123], biomechanics [11], optimal control of bodily processes [17,53], signal processing [116], and machine learning [112].

6.2 Nonstandard optimization problem formulations

Most applications above deal with either system design, parameter estimation, model predictive control (MPC) or moving horizon estimation [51] (MHE) formulations. For some of these problems, it is the nonstandard formulation of the optimization problem that poses the main challenge.

Some applications transcend the classical subdivisions; dual control combines control and learning [41,66,122], codesign of optimal trajectory and reference follower [57], multi-objective design [99,126], the use of different transcription methods on different parts of the system statespace [2].

Concerning robustness, formulations include the use of scenario trees [86], ellipsoidal calculus [89], spline-relaxations [127], stochastic control using linearizations [57,115], using sigma-points [110], and using polynomial chaos expansion [109,110]. In [57], stochastic optimal control was efficiently implemented by embedding a discrete periodic Lyapunov solver in the CasADi expression graphs.

In MPC research, formulations include multi-level iterations [47], offset-free design [111], Lyapunov-based MPC [36,37], multi-objective MPC [100], distributed MPC [102,128], time-optimal path following [32] and tube following [33].

Further formulations include hybrid OCP [44], treatment of systems with invariants [63,114], an improved Gauss-Newton method for OCP [131].

6.3 Software packages using CasADi

Software packages that rely on CasADi for algorithmic differentiation and optimization include the `JModelica.org` package for simulation and optimization [8,97], the `Greybox` tool for constructing thermal building models [31], the `do-mpc` environment for efficient testing and implementation of robust nonlinear MPC [91,92], `mpc-tools-casadi` for nonlinear MPC [96], the `casiopeia` toolbox for parameter estimation and optimum experimental design [22], the `RTC-Tools 2` package for control of hydraulic networks, the `omgtools` package for real-time motion planning in the presence of moving obstacles, the `Pomodoro` toolbox for multi-objective optimal control [25], the `spline` toolbox for robust optimal control [127], and a `MATLAB` optimal control toolbox [81].

7 Discussion and outlook

Since the release of CasADi 3.0 in early 2016, the scope and syntax can be considered mature and no more major non-backwards compatible changes are foreseen. Current development focusses on making existing features more efficient, by addressing speed and memory bottlenecks, and adding new functionality.

An area of special interest are mixed integer optimal control problems (MIOCPs). Problems of this class appear naturally across engineering fields, e.g. in the form of discrete actuators in model predictive control (MPC) formulations. At the time of this writing, CasADi included support for mixed-integer QP and NLP problems, as explained in Sect. 4, but MIOCPs are largely unexplored.

Another ongoing development is to enable automatic sensitivity analysis for NLPs. Assuming the optimal NLP solution meets regularity criteria as described in e.g. [43], directional derivatives of an NLP solver object are guaranteed to exist and can be calculated using information that can be extracted from the expression graphs. Parametric sensitivity information is useful in a range of applications, including the estimation of covariances in parameter estimation.

Acknowledgements The authors thank for the generous support that has made this work possible. In particular: the K.U. Leuven Research Council via CoE EF/05/006 Optimization in Engineering (OPTEC); the Flemish Government via FWO; the Belgian State via Science Policy programming (IAP VII, DYSCO); the European Union via HDMP (223854), EMBOCON (248940), HIGHWIND (259166), TEMPO (607957), AWESCO (642682); the Helmholtz Association via vICERP; the German Federal Ministry for Economic Affairs and Energy (BMWi) via projects eco4wind and DyConPV; the German Research Foundation (DFG) via Research Unit FOR 2401; Flanders Make via MBSE4M, Drivetrain Co-design, Conceptdesign. We also thank our industrial partners, including GE Global Research and Johnson Controls International Inc. Finally, we thank the reviewers for valuable comments that helped to improve the final manuscript.

A Summary of main features of CasADi, version 3.4.4

In the following is a summary of the main features of CasADi 3.4.4, which was released in May 2018, after the acceptance of this paper. Newly added features, not present in CasADi 3.1 covered so far, are marked with footnotes. More details about these features can be found on the CasADi webpage, <http://casadi.org>, and in the paper *Sensitivity Analysis for Nonlinear Programming in CasADi*, which is currently in preparation.²

Symbolic framework with algorithmic differentiation (AD)

A state-of-the-art implementation of algorithmic differentiation (AD), implemented within a symbolic framework, forms the backbone of CasADi. Users construct directed acyclic expression graphs using an *everything-is-a-sparse-matrix* syntax and expressions for derivatives are generated automatically using AD via *source-code-transformation*. CasADi implements the forward and reverse modes of AD and uses a graph coloring approach to construct large-and-sparse Jacobians and Hessians. Generated expressions are encapsulated in function objects that can be evaluated, numerically or symbolically, in virtual machines (VMs).

Compared to similar frameworks, CasADi scales well to higher dimensions, and offers a rich set of differentiable operations. Supported operations include common matrix-valued operations, serial or parallel function calls, (non)linear systems of equations, initial-value problems in ordinary differential equations (ODE) or differential-algebraic equations (DAE) and spline-based lookup tables. External code can be embedded with derivative information either user-provided or approximated by finite differences.³

Core self-containment, auto-generated front-ends via SWIG

The symbolic core of CasADi is written in modern C++, with no external dependencies. While C++ offers great interoperability with other tools, high performance and multi-platform support, it lacks the interactivity and ease-of-use associated with scripting languages such as Python or MATLAB/Octave. CasADi was therefore designed to allow front-ends to be generated automatically using the open-source tool SWIG. At the time of this writing, Python, MATLAB and Octave were supported through full-featured and documented front-ends. The tool has also been successfully used from JAVA and Haskell.

License and availability

CasADi's source code is hosted on Github and released under GNU Lesser General Public License (LGPL) on <http://casadi.org>. The relatively permissive LGPL allows

² A paper draft is available at http://www.optimization-online.org/DB_HTML/2018/05/6642.html.

³ Support for finite differences, with automatic step-size selection, was added in CasADi 3.3.

CasADi to be used royalty-free in commercial and academic software. The code is built and tested on `travis-ci`, with full-featured binaries available for common Linux, Mac and Windows systems. In addition, the Python interface is available from `pip`. CasADi can also be run from a demo server.⁴

C code generation, just-in-time compilation

A large subset of expressions can be exported as self-contained C code without memory allocation. This is useful for embedded applications or to speed up computations using just-in-time compilation (JIT). The C code can be compiled into shared libraries or be called directly from either MATLAB/Octave via a generated MEX interface or from the command line.

Plugin infrastructure

The core of CasADi supports a number of standard problems in numerical optimization, including initial-value problems in ODE or DAE, linear and nonlinear systems of equations, nonlinear programs (NLPs) and quadratic programs (QPs). The user specifies such problems in a uniform way and the solution is delegated to a solver *plugin*, loaded as a dynamically linked library (DLL) at runtime. Solver plugins include solvers that are distributed with CasADi and interfaces to third-party software packages.

Linear systems of equations

Linear systems of equations can be embedded into symbolic expressions via “backslash” nodes. Derivatives of such operations are calculated via chain rules for linear system solves. Supported plugins include LDLT and QR⁵ as well as interfaces to CSPARSE and LAPACK.

Nonlinear systems of equations

Nonlinear systems of equations can be formulated and solved by defining `rootfinder` instances in CasADi. Derivatives of such objects are calculated analytically using the implicit function theorem and the chain rule for linear system solves. Supported plugins include standard Newton methods and KINSOL from the SUNDIALS suite.

⁴ Cf. <http://live.casadi.org> (Python) and <http://live-octave.casadi.org> (Octave).

⁵ Two sparse direct linear solvers based on LDLT (for symmetric positive definite systems) and QR (for general sparse matrices), with support for C code generation, were added in CasADi 3.3.

Initial-value problems in ODE/DAE with automatic sensitivity analysis

Initial-value problems in ODE or DAE can be calculated using explicit or implicit Runge–Kutta methods or interfaces to IDAS/CVODES from the SUNDIALS suite. Derivatives of arbitrary order are calculated using automatically generated forward and adjoint sensitivity equations.

Quadratic programming

QPs can be formulated either using a traditional syntax, by explicitly providing the linear and quadratic terms, or using a syntax which mirrors that of NLP solvers. Solvers for quadratic programming include a primal-dual active-set method⁶ and interfaces to CPLEX, GUROBI, HPMPC, OOQP and qpOASES. A subset of the solvers support mixed-integer QPs.

Nonlinear programming with automatic sensitivity analysis, Optistack

NLPs can be solved using block structure or general sparsity exploiting sequential quadratic programming (SQP) or interfaces to IPOPT/BONMIN, BlockSQP, WORHP, KNITRO and SNOPT. Solution sensitivities can be calculated automatically by applying the implicit function theorem to the first order optimality conditions.⁷ A subset of the solvers support mixed-integer NLPs.

Optistack, a simple but powerful abstraction layer, simplifies the formulation and solution of NLPs. It manages the creation and optimal-value retrieval of decision variables, allows a mathematical notation to specify constraints, and may identify problematic constraints when a solver reports infeasibility.

References

1. Albersmeyer, J., Diehl, M.: The lifted Newton method and its application in optimization. *SIAM J. Optim.* **20**(3), 1655–1684 (2010)
2. Albert, A., Imsland, L., Haugen, J.: Numerical optimal control mixing collocation with single shooting: a case study. *IFAC-PapersOnLine* **49**(7), 290–295 (2016)
3. Alexandrescu, A.: *Modern C++ Design*. Addison-Wesley, Reading (2001)
4. Anderson, E., et al.: *LAPACK Users' Guide*, 3rd edn. SIAM, Philadelphia (1999)
5. Andersson, J.: *A General-Purpose Software Framework for Dynamic Optimization*. Ph.D. thesis, Arenberg Doctoral School, KU Leuven (2013)
6. Andersson, J., Kozma, A., Gillis, J., Diehl, M.: CasADi User Guide. <http://guide.casadi.org>. Accessed 2 June 2018
7. Andersson, L.E., Scibilia, F., Imsland, L.: An estimation-forecast set-up for iceberg drift prediction. *Cold Reg. Sci. Technol.* **131**, 88–107 (2016)
8. Axelsson, M., Magnusson, F., Henningsson, T.: A framework for nonlinear model predictive control in jmodelica.org. In: 11th International Modelica Conference, 118, pp. 301–310. Linköping University Electronic Press (2015)

⁶ A new primal-dual active-set method for quadratic programming was added in CasADi 3.4.

⁷ Support for automatic sensitivity analysis for NLPs was added in CasADi 3.4.

9. Bonmin. <https://projects.coin-or.org/Bonmin>. Accessed 1 Feb 2017
10. Belkhir, F., Cabo, D.K., Feigner, F., Frey, G.: Optimal startup control of a steam power plant using the Jmodelica Platform. *IFAC-PapersOnLine* **48**(1), 204–209 (2015)
11. Belousov, B., Neumann, G., Rothkopf, C.A., Peters, J.R.: Catching heuristics are optimal control policies. In: Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, pp. 1426–1434. Curran Associates, Red Hook, NY (2016)
12. Berntorp, K., Magnusson, F.: Hierarchical predictive control for ground-vehicle maneuvering. In: *American Control Conference (ACC)*, pp. 2771–2776 (2015)
13. Berx, K., Gadeyne, K., Dhadamus, M., Pipeleers, G., Pinte, G.: Model-based gearbox synthesis. In: *Mechatronics Forum International Conference*, pp. 599–605 (2014)
14. Biegler, L.: Solution of dynamic optimization problems by successive quadratic programming and orthogonal collocation. *Comput. Chem. Eng.* **8**, 243–248 (1984)
15. Bock, H., Eich, E., Schlöder, J.: Numerical solution of constrained least squares boundary value problems in differential-algebraic equations. In: Strehmel, K. (ed.) *Numerical Treatment of Differential Equations*. Teubner, Leipzig (1988)
16. Bock, H., Plitt, K.: A multiple shooting algorithm for direct solution of optimal control problems. In: *IFAC World Congress*, pp. 242–247. Pergamon Press (1984)
17. Boiroux, D., Hagdrup, M., Mahmoudi, Z., Madsen, H., Bagterp, J., et al.: An ensemble nonlinear model predictive control algorithm in an artificial pancreas for people with type 1 diabetes. In: *European Control Conference (ECC)*, pp. 2115–2120 (2016)
18. Boyd, S., Vandenberghe, L.: *Convex Optimization*. Cambridge University Press, Cambridge (2004)
19. Bremer, J., Rätze, K.H., Sundmacher, K.: CO₂ methanation: optimal start-up control of a fixed-bed reactor for power-to-gas applications. *AIChE J.* **63**(1), 23–31 (2017)
20. Büskens, C., Wassel, D.: *Modeling and Optimization in Space Engineering*, Chapter. The ESA NLP Solver WORHP. Springer, Berlin (2012)
21. Byrd, R.H., Nocedal, J., Waltz, R.A.: KNITRO: an integrated package for nonlinear optimization. In: Pillo, G., Roma, M. (eds.) *Large Scale Nonlinear Optimization*, pp. 35–59. Springer, Berlin (2006)
22. casiopeia. <https://github.com/adbuerger/casiopeia>. Accessed 19 May 2017
23. Clp. <https://projects.coin-or.org/Clp>. Accessed 1 Feb 2017
24. CppAD. <http://www.coin-or.org/CppAD>. Accessed 11 May 2012
25. Cabianca, L.: *Advanced techniques for robust dynamic optimization of chemical reactors*. Master's thesis, Politecnico Milano (2014)
26. Codas, A., Jahanshahi, E., Foss, B.: A two-layer structure for stabilization and optimization of an oil gathering network. *IFAC-PapersOnLine* **49**(7), 931–936 (2016)
27. Cuthrell, J., Biegler, L.: Simultaneous optimization and solution methods for batch reactor profiles. *Comput. Chem. Eng.* **13**(1/2), 49–62 (1989)
28. Davis, T.: *Direct Methods for Sparse Linear Systems*. SIAM, Philadelphia (2006)
29. De Coninck, R., Helsen, L.: Practical implementation and evaluation of model predictive control for an office building in Brussels. *Energy Build.* **111**, 290–298 (2016)
30. De Coninck, R., Helsen, L.: Quantification of flexibility in buildings by cost curves—methodology and application. *Appl. Energy* **162**, 653–665 (2016)
31. De Coninck, R., Magnusson, F., Åkesson, J., Helsen, L.: Toolbox for development and validation of grey-box building models for forecasting and control. *J. Build. Perform. Simul.* **9**(3), 288–303 (2016)
32. Debrouwere, F., Loock, W.V., Pipeleers, G., Diehl, M., Schutter, J.D.: Time-optimal path following for robots with object collision avoidance using Lagrangian duality. In: *2012 Benelux Meeting on Systems and Control* (2013)
33. Debrouwere, F., Van Loock, W., Pipeleers, G., Swevers, J.: Time-optimal tube following for robotic manipulators. In: *2014 IEEE 13th International Workshop on Advanced Motion Control (AMC)*, pp. 392–397 (2014)
34. Domahidi, A., Zraggen, A., Zeilinger, M., Morari, M., Jones, C.: Efficient interior point methods for multistage problems arising in receding horizon control. In: *IEEE Conference on Decision and Control (CDC)*, pp. 668–674. Maui (2012)
35. Dunning, I., Huchette, J., Lubin, M.: JuMP: a modeling language for mathematical optimization. [arXiv:1508.01982](https://arxiv.org/abs/1508.01982) [math.OC] (2015)
36. Durand, H., Ellis, M., Christofides, P.D.: Economic model predictive control designs for input rate-of-change constraint handling and guaranteed economic performance. *Comput. Chem. Eng.* **92**, 18–36 (2016)

37. Ellis, M., Christofides, P.D.: On closed-loop economic performance under Lyapunov-based economic model predictive control. In: American Control Conference (ACC), pp. 1778–1783 (2016)
38. Erhard, M., Horn, G., Diehl, M.: A quaternion-based model for optimal control of an airborne wind energy system. *ZAMM J. Appl. Math. Mech. Z. Angew. Math. Mech.* **97**(1), 7–24 (2017)
39. Ersdal, A.M., Fabozzi, D., Imsland, L., Thornhill, N.F.: Model predictive control for power system frequency control taking into account imbalance uncertainty. *IFAC Proc. Vol.* **47**(3), 981–986 (2014)
40. Ersdal, A.M., Imsland, L., Uhlen, K.: Model predictive load-frequency control. *IEEE Trans. Power Syst.* **31**(1), 777–785 (2016)
41. Feng, X., Houska, B.: Real-time algorithm for self-reflective model predictive control. [arXiv:1611.02408](https://arxiv.org/abs/1611.02408) (2016)
42. Ferreau, H.: qpOASES—an open-source implementation of the online active set strategy for fast model predictive control. In: Workshop on Nonlinear Model Based Control—Software and Applications, pp. 29–30 (2007)
43. Fiacco, A.V., Ishizuka, Y.: Sensitivity and stability analysis for nonlinear programming. *Ann. Oper. Res.* **27**, 215–235 (1990)
44. Fouquet, M., Guéguen, H., Faille, D., Dumur, D.: Hybrid dynamic optimization of power plants using sum-up rounding and adaptive mesh refinement. In: IEEE Conference on Control Applications, pp. 316–321 (2014)
45. Fourer, R., Gay, D.M., Kernighan, B.W.: *AMPL: a modeling language for mathematical programming*, 2nd edn. Thomson, Stamford (2003)
46. Frasch, J.V., Vukov, M., Ferreau, H., Diehl, M.: A dual Newton strategy for the efficient solution of sparse quadratic programs arising in SQP-based nonlinear MPC (2013). *Optimization Online* 3972
47. Frasch, J.V., Wirsching, L., Sager, S., Bock, H.G.: Mixed-level iteration schemes for nonlinear model predictive control. *IFAC Proc. Vol.* **45**(17), 138–144 (2012)
48. Frison, G., Sorensen, H., Dammann, B., Jorgensen, J.: High-performance small-scale solvers for linear model predictive control. In: European Control Conference (ECC), pp. 128–133 (2014)
49. Gabiccini, M., Artoni, A., Pannocchia, G., Gillis, J.: A computational framework for environment-aware robotic manipulation planning. In: Bicchi, A., Burgard, W. (eds.) *Robotics Research. Proceedings in Advanced Robotics*, vol 3. Springer, Cham (2018)
50. Gebremedhin, A.H., Manne, F., Pothén, A.: What color is your Jacobian? Graph coloring for computing derivatives. *SIAM Rev.* **47**, 629–705 (2005)
51. Geebelen, K., Wagner, A., Gros, S., Swevers, J., Diehl, M.: Moving horizon estimation with a huber penalty function for robust pose estimation of tethered airplanes. In: American Control Conference (ACC), pp. 6169–6174 (2013)
52. Gertz, E., Wright, S.: Object-oriented software for quadratic programming. *ACM Trans. Math. Softw.* **29**(1), 58–81 (2003)
53. Gesenhues, J., Hein, M., Habigt, M., Mechelinck, M., Albin, T., Abel, D.: Nonlinear object-oriented modeling based optimal control of the heart. In: European Control Conference (ECC), pp. 2108–2114 (2016)
54. Giering, R., Kaminski, T.: Automatic sparsity detection implemented as a source-to-source transformation. In: Alexandrov, V.N., van Albada, G.D., Sloot, P.M.A., Dongarra, J. (eds.) *Lecture Notes in Computer Science*, vol. 3994, pp. 591–598. Springer, Berlin, Heidelberg (2006)
55. Giles, M.: Collected matrix derivative results for forward and reverse mode algorithmic differentiation. In: Bischof, C.H., Bücker, H.M., Hovland, P., Naumann, U., Utke, J. (eds.) *Advances in Automatic Differentiation*, pp. 35–44. Springer, Berlin (2008)
56. Gill, P., Murray, W., Saunders, M.: SNOPT: an SQP algorithm for large-scale constrained optimization. *SIAM Rev.* **47**(1), 99–131 (2005)
57. Gillis, J.: Practical methods for approximate robust periodic optimal control of nonlinear mechanical systems. Ph.D. thesis, Arenberg Doctoral School, KU Leuven (2015)
58. Griewank, A., Juedes, D., Mitev, H., Utke, J., Vogel, O., Walther, A.: ADOL-C: a package for the automatic differentiation of algorithms written in C, C++. Tech. rep., Technical University of Dresden, Institute of Scientific Computing and Institute of Geometry (1999). Updated version of the paper published in *ACM Transactions on Mathematical Software*, vol. 22, pp. 131–167 (1996)
59. Griewank, A., Mitev, C.: Detecting Jacobian sparsity patterns by Bayesian probing. *Math. Program.* **93**(1), 1–25 (2002)
60. Griewank, A., Walther, A.: *Evaluating Derivatives*, 2nd edn. SIAM, Philadelphia (2008)

61. Gros, S.: A distributed algorithm for NMPC-based wind farm control. In: IEEE Conference on Decision and Control (CDC), pp. 4844–4849 (2014)
62. Gros, S., Diehl, M.: Modeling of airborne wind energy systems in natural coordinates. In: Ahrens, U., Diehl, M., Schmehl, R. (eds.) *Airborne Wind Energy*. Springer, Berlin (2013)
63. Gros, S., Zanon, M., Diehl, M.: Baumgarte stabilisation over the SO (3) rotation group for control. In: IEEE Conference on Decision and Control (CDC), pp. 620–625 (2015)
64. Gurobi Optimizer Reference Manual Version 7.0. <https://www.gurobi.com>. Accessed 2 June 2018
65. HSL. A collection of Fortran codes for large scale scientific computation. <http://www.hsl.rl.ac.uk>. Accessed 2 June 2018
66. Hanssen, K.G., Foss, B.: Scenario based implicit dual model predictive control. *IFAC-PapersOnLine* **48**(23), 416–421 (2015)
67. Hart, W., Watson, J.P., Woodruff, D.: Pyomo: modeling and solving mathematical programs in Python. *Math. Program. Comput.* **3**(3), 1–42 (2011)
68. Haugen, J., Imsland, L.: Monitoring moving objects using aerial mobile sensors. *IEEE Trans. Control Syst. Technol.* **24**(2), 475–486 (2016)
69. Herrmann, S., Utschick, W.: Availability and interpretability of optimal control for criticality estimation in vehicle active safety. In: Design, Automation & Test in Europe Conference & Exhibition (DATE), pp. 415–420 (2016)
70. Hindmarsh, A., et al.: SUNDIALS: suite of nonlinear and differential/algebraic equation solvers. *ACM Trans. Math. Softw.* **31**, 363–396 (2005)
71. Holmqvist, A., Magnusson, F.: Open-loop optimal control of batch chromatographic separation processes using direct collocation. *J. Process Control* **46**, 55–74 (2016)
72. Holmqvist, A., et al.: Dynamic parameter estimation of atomic layer deposition kinetics applied to in situ quartz crystal microbalance diagnostics. *Chem. Eng. Sci.* **111**, 15–33 (2014)
73. Horn, G., Diehl, M.: Numerical trajectory optimization for airborne wind energy systems described by high fidelity aircraft models. In: Ahrens, U., Diehl, M., Schmehl, R. (eds.) *Airborne Wind Energy*, pp. 205–218. Springer, Berlin, Heidelberg (2013)
74. IBM Corp.: IBM ILOG CPLEX V12.1, User's Manual for CPLEX (2009)
75. Janka, D., Kirches, C., Sager, S., Wchter, A.: An SR1/BFGS SQP algorithm for nonconvex nonlinear programs with block-diagonal Hessian matrix. *Math. Program. Comput.* **8**(4), 435–459 (2016)
76. Kirches, C.: Fast numerical methods for mixed-integer nonlinear model-predictive control. Ph.D. thesis, Ruprecht-Karls-Universität Heidelberg (2010)
77. Krishnamoorthy, D., Foss, B., Skogestad, S.: Real-time optimization under uncertainty applied to a gas lifted well network. *Processes* **4**(4), 52 (2016)
78. Larsson, P.O., Casella, F., Magnusson, F., Andersson, J., Diehl, M., Akesson, J.: A framework for nonlinear model-predictive control using object-oriented modeling with a case study in power plant start-up. In: IEEE Conference on Computer Aided Control System Design (CACSD), pp. 346–351 (2013)
79. Lattner, C., Adve, V.: LLVM: a compilation framework for lifelong program analysis & transformation. In: International Symposium on Code Generation and Optimization (2004)
80. Le, T.T., Truong, B.D., Jost, F., Le, C.P., Halvorsen, E., Sager, S.: Synthesis of optimal controls and numerical optimization for the vibration-based energy harvesters. [arXiv:1608.08885](https://arxiv.org/abs/1608.08885) (2016)
81. Leek, V.: An Optimal Control Toolbox for MATLAB Based on CasADi. Master's thesis, Linköping University (2016)
82. Leineweber, D.: Efficient reduced SQP methods for the optimization of chemical processes described by large sparse DAE models, *Fortschritt-Berichte VDI Reihe 3, Verfahrenstechnik*, vol. 613. VDI Verlag, Düsseldorf (1999)
83. Licitra, G., Sieberling, S., Engelen, S., Williams, P., Ruiterkamp, R., Diehl, M.: Optimal control for minimizing power consumption during holding patterns for airborne wind energy pumping system. In: European Control Conference (ECC), pp. 1574–1579 (2016)
84. Liu, M., Tan, Y., Padois, V.: Generalized hierarchical control. *Auton. Robots* **40**, 17–31 (2016)
85. Lopes, V.V., et al.: On the use of markov chain models for the analysis of wind power time-series. In: Conference on Environment and Electrical Engineering (EEEIC), pp. 770–775 (2012)
86. Lucia, S., Andersson, J.A., Brandt, H., Bouaswaig, A., Diehl, M., Engell, S.: Efficient robust economic nonlinear model predictive control of an industrial batch reactor. *IFAC Proc. Vol.* **47**(3), 11093–11098 (2014)

87. Lucia, S., Engell, S.: Control of towing kites under uncertainty using robust economic nonlinear model predictive control. In: *European Control Conference (ECC)*, pp. 1158–1163 (2014)
88. Lucia, S., Paulen, R.: Robust nonlinear model predictive control with reduction of uncertainty via robust optimal experiment design. *IFAC Proc. Vol.* **47**(3), 1904–1909 (2014)
89. Lucia, S., Paulen, R., Engell, S.: Multi-stage nonlinear model predictive control with verified robust constraint satisfaction. In: *IEEE Conference on Decision and Control (CDC)*, pp. 2816–2821 (2014)
90. Lucia, S., Schliemann-Bullinger, M., Findeisen, R., Bullinger, E.: A set-based optimal control approach for pharmacokinetic/pharmacodynamic drug dosage design. *IFAC-PapersOnLine* **49**(7), 797–802 (2016)
91. Lucia, S., Tătulea-Codrean, A., Schoppmeyer, C., Engell, S.: An environment for the efficient testing and implementation of robust nmPC. In: *IEEE Conference on Control Applications*, pp. 1843–1848 (2014)
92. Lucia, S., Tătulea-Codrean, A., Schoppmeyer, C., Engell, S.: Rapid development of modular and sustainable nonlinear model predictive control solutions. *Control Eng. Pract.* **60**, 51–62 (2017)
93. Lundahl, K., Lee, C.F., Frisk, E., Nielsen, L.: Path-dependent rollover prevention for critical truck maneuvers. In: *Symposium of the International Association for Vehicle System Dynamics (IAVSD 2015)*, p. 317. CRC Press (2016)
94. Lynn, L.L., Parkin, E.S., Zahradnik, R.L.: Near-optimal control by trajectory approximations. *I&EC Fundam.* **9**(1), 58–63 (1970)
95. Lynn, L.L., Zahradnik, R.L.: The use of orthogonal polynomials in the near-optimal control of distributed systems by trajectory approximation. *Int. J. Control* **12**(6), 1079–1087 (1970)
96. mpc-tools-casadi. <https://bitbucket.org/rawlings-group/mpc-tools-casadi>. Accessed 19 May 2017
97. Magnusson, F., Åkesson, J.: Dynamic optimization in jmodelica.org. *Processes* **3**(2), 471–496 (2015)
98. Marcucci, T., Gabicchini, M., Artoni, A.: A two-stage trajectory optimization strategy for articulated bodies with unscheduled contact sequences. *IEEE Robot. Autom. Lett.* **2**(1), 104–111 (2017)
99. Maree, J., Imsland, L.: Multi-objective predictive control for non steady-state operation. In: *European Control Conference (ECC)*, pp. 1541–1546 (2013)
100. Maree, J., Imsland, L.: On multi-objective economic predictive control for cyclic process operation. *J. Process Control* **24**(8), 1328–1336 (2014)
101. Maree, J.P., Imsland, L.: Optimal control strategies for oil production under gas coning conditions. In: *IEEE Conference on Control Applications*, pp. 572–578 (2014)
102. Martí, R., et al.: An efficient distributed algorithm for multi-stage robust nonlinear predictive control. In: *European Control Conference (ECC)*, pp. 2664–2669 (2015)
103. Merz, M., Johansen, T.A.: Feasibility study of a circularly towed cable-body system for uav applications. In: *Conference on Unmanned Aircraft Systems (ICUAS)*, pp. 1182–1191 (2016)
104. Minko, T., Wisniewski, R., Bendtsen, J.D., Izadi-Zamanabadi, R.: Non-linear model predictive supervisory controller for building, air handling unit with recuperator and refrigeration system with heat waste recovery. In: *IEEE Conference on Control Applications*, pp. 1274–1281 (2016)
105. Mukkula, A.R.G., Paulen, R.: Model-based design of optimal experiments for nonlinear systems in the context of guaranteed parameter estimation. *Comput. Chem. Eng.* **99**, 198–213 (2017)
106. Naumann, U.: *The Art of Differentiating Computer Programs: An Introduction to Algorithmic Differentiation*. No. 24 in *Software, Environments, and Tools*. SIAM, Philadelphia (2012)
107. Nezhadali, V., Eriksson, L.: A framework for modeling and optimal control of automatic transmission systems. *IFAC-PapersOnLine* **48**(15), 285–291 (2015)
108. Nezhadali, V., Eriksson, L.: Optimal control of engine controlled gearshift for a diesel-electric powertrain with backlash. *IFAC-PapersOnLine* **49**(11), 762–768 (2016)
109. Nimmegeers, P., Telen, D., Beetens, M., Logist, F., Van Impe, J.: Parametric uncertainty propagation for robust dynamic optimization of biological networks. In: *American Control Conference (ACC)*, pp. 6929–6934 (2016)
110. Nimmegeers, P., Telen, D., Logist, F., Van Impe, J.: Dynamic optimization of biological networks under parametric uncertainty. *BMC Syst. Biol.* **10**(1), 86 (2016)
111. Pannocchia, G., Gabicchini, M., Artoni, A.: Offset-free MPC explained: novelties, subtleties, and applications. *IFAC-PapersOnLine* **48**(23), 342–351 (2015)
112. Patil, S., Kahn, G., Laskey, M., Schulman, J., Goldberg, K., Abbeel, P.: Scaling up gaussian belief space planning through covariance-free trajectory optimization and automatic differentiation. In: *Algorithmic Foundations of Robotics XI*, pp. 515–533. Springer, Berlin (2015)

113. Poland, J., Stadler, K.S.: Stochastic optimal planning of solar thermal power. In: Les Antibes J. (ed.) IEEE Conference on Control Applications, pp. 593–598 (2014)
114. Reiter, A., Müller, A., Gattringer, H.: Inverse kinematics in minimum-time trajectory planning for kinematically redundant manipulators. In: Annual Conference of the IEEE Industrial Electronics Society (IECON), pp. 6873–6878 (2016)
115. Rostampour, V., Esfahani, P.M., Keviczky, T.: Stochastic nonlinear model predictive control of an uncertain batch polymerization reactor. *IFAC-PapersOnLine* **48**(23), 540–545 (2015)
116. Scholz, T., Raischel, F., Lopes, V.V., Lehle, B., Wächter, M., Peinke, J., Lind, P.G.: Parameter-free resolution of the superposition of stochastic signals. *Phys. Lett. A* **381**(4), 194–206 (2017)
117. Scott, P., Thiébaux, S.: Distributed multi-period optimal power flow for demand response in micro-grids. In: ACM Conference on Future Energy Systems, pp. 17–26 (2015)
118. Sirmatel, I.I., Geroliminis, N.: Model predictive control of large-scale urban networks via perimeter control and route guidance actuation. In: IEEE Conference on Decision and Control (CDC), pp. 6765–6770 (2016)
119. Sivertsson, M., Eriksson, L.: Optimal stationary control of diesel engines using periodic control. *J. Automob. Eng.* **231**(4), 457–475 (2017)
120. Skjong, E., et al.: Management of harmonic propagation in a marine vessel by use of optimization. In: IEEE Transportation Electrification Conference and Expo (ITEC), pp. 1–8 (2015)
121. St John, P.C., Doyle, F.J.: Estimating confidence intervals in predicted responses for oscillatory biological models. *BMC Syst. Biol.* **7**(1), 71 (2013)
122. Thangavel, S., Lucia, S., Paulen, R., Engell, S.: Towards dual robust nonlinear model predictive control: A multi-stage approach. In: American Control Conference (ACC), pp. 428–433 (2015)
123. Trägårdh, M., et al.: Input estimation for drug discovery using optimal control and markov chain Monte Carlo approaches. *J. Pharmacokinet. Pharmacodyn.* **43**(2), 207–221 (2016)
124. Utstumo, T., Berge, T.W., Gravadahl, J.T.: Non-linear model predictive control for constrained robot navigation in row crops. In: IEEE Conference on Industrial Technology (ICIT), pp. 357–362 (2015)
125. van Duijkeren, N., Keviczky, T., Nilsson, P., Laine, L.: Real-time NMPC for semi-automated highway driving of long heavy vehicle combinations. *IFAC-PapersOnLine* **48**(23), 39–46 (2015)
126. Vallerio, M., Vercammen, D., Van Impe, J., Logist, F.: Interactive NBI and (E)NNC methods for the progressive exploration of the criteria space in multi-objective optimization and optimal control. *Comput. Chem. Eng.* **82**, 186–201 (2015)
127. Van Loock, W., Pipeleers, G., Swevers, J.: B-spline parameterized optimal motion trajectories for robotic systems with guaranteed constraint satisfaction. *Mech. Sci.* **6**(2), 163–171 (2015)
128. Van Parys, R., Pipeleers, G.: Online distributed motion planning for multi-vehicle systems. In: European Control Conference (ECC), pp. 1580–1585 (2016)
129. Venrooij, J., et al.: Comparison between filter- and optimization-based motion cueing in the Daimler driving simulator. In: Driving Simulation Conference (2016)
130. Verheyeweghen, A., Jäschke, J.: Health-aware operation of a subsea gas compression system under uncertainty. In: Foundations of Computer Aided Process Operations/Chemical Process Control (2017)
131. Verschueren, R., van Duijkeren, N., Quiryne, R., Diehl, M.: Exploiting convexity in direct optimal control: a sequential convex quadratic programming method. In: IEEE Conference on Decision and Control (CDC), pp. 1099–1104 (2016)
132. Verschueren, R., van Duijkeren, N., Swevers, J., Diehl, M.: Time-optimal motion planning for n-dof robot manipulators using a path-parametric system reformulation. In: American Control Conference (ACC), pp. 2092–2097 (2016)
133. Vochten, M., De Laet, T., De Schutter, J.: Generalizing demonstrated motions and adaptive motion generation using an invariant rigid body trajectory representation. In: IEEE Conference on Robotics and Automation (ICRA), pp. 234–241 (2016)
134. Wächter, A., Biegler, L.: On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. *Math. Program.* **106**(1), 25–57 (2006)
135. Walraven, D., et al.: Optimum configuration of shell-and-tube heat exchangers for the use in low-temperature organic rankine cycles. *Energy Convers. Manag.* **83**, 177–187 (2014)
136. Welsh, D.J.A., Powell, M.B.: An upper bound for the chromatic number of a graph and its application to timetabling problems. *Comput. J.* **10**, 85–86 (1967)