FMINSDP – a code for solving optimization problems with matrix inequality constraints

Carl-Johan Thore

September 1, 2016

This document is a theoretical and practical introduction to the Matlab-code `fminsdp`, designed to find local solutions to non-linear, non-convex optimization problems (NLPs) with both scalar constraints and (small-size) matrix inequality constraints.

Notation. A matrix $\boldsymbol{A} \in \mathbb{R}^{m \times m}$ is said to be *positive semi-definite* if $\boldsymbol{y}^\mathsf{T} \boldsymbol{A} \boldsymbol{y} \geq 0$ for all $\boldsymbol{y} \in \mathbb{R}^m$. It is convenient to introduce the notation "$\boldsymbol{A} \succeq \boldsymbol{0}$" to indicate that $\boldsymbol{A}$ is positive semi-definite. A *matrix inequality* is here defined as an expression of the form

$$\boldsymbol{\mathcal{A}}(\boldsymbol{x}) \succeq \boldsymbol{0}, \tag{1}$$

where $\boldsymbol{\mathcal{A}}$ is a map from $\Omega \subset \mathbb{R}^n$ to the space $\mathbb{S}^m$ of symmetric, real-valued matrices of size $m \times m$.

1 The optimization problem

`fminsdp` attempts to find a local solution to non-linear, non-convex optimization problems of the form

$$\underset{\boldsymbol{x} \in \mathbb{R}^n}{\text{minimize}} \quad f(\boldsymbol{x})$$

$$\text{subject to} \begin{cases} \boldsymbol{A}_{eq}\boldsymbol{x} = \boldsymbol{b}_{eq} & \text{linear equality constraints} \\ \boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{b} & \text{linear inequality constraints} \\ \boldsymbol{c}_{eq}(\boldsymbol{x}) = \boldsymbol{0} & \text{nonlinear equality constraints} \\ \boldsymbol{c}(\boldsymbol{x}) \leq \boldsymbol{0} & \text{nonlinear inequality constraints} \\ \boldsymbol{l} \leq \boldsymbol{x} \leq \boldsymbol{u} & \text{box constraints} \\ \boldsymbol{\mathcal{A}}_i(\boldsymbol{x}) \succeq \boldsymbol{0}, \quad i = 1,...,q & \text{matrix inequality constraints,} \end{cases} \tag{2}$$

where $\boldsymbol{A}_{eq}$, $\boldsymbol{A}$, $\boldsymbol{b}_{eq}$, $\boldsymbol{b}$, $\boldsymbol{l}$ and $\boldsymbol{u}$ are constant matrices and vectors, respectively. The functions f, $\boldsymbol{c}$, $\boldsymbol{c}_{eq}$ and $\boldsymbol{\mathcal{A}}_i : \Omega_i \to \mathbb{S}^{m_i}$, $i = 1,...,q$, can be non-linear and are (preferably) at least twice continuously differentiable. Users familiar with `fmincon` from the Optimization Toolbox in Matlab [10] should recognize the form of problem (2) — the novelty here is the addition of q matrix inequality constraints.

`fminsdp` offers three different methods to treat problem (2): the "cholesky-method", the "ldl-method" [2] and "penlab" [5]. The first two methods works by reformulating the problem into a standard NLP which can be solved by any of the NLP-solvers interfaced by `fminsdp`. The third method relies on the external solver PENLAB[1], which treats the problem directly using an augmented Lagrangian-type algorithm [5].

---

[1]Downloadable from `http://web.mat.bham.ac.uk/kocvara/penlab/`.

2   Theoretical background

2.1   The cholesky-method

In the cholesky-method, `fminsdp` reformulates the matrix inequality constraints into scalar equality constraints using the fact that a matrix is positive semi-definite if and only if it admits a Cholesky decomposition. In other words,

Theorem 1. A symmetric matrix $\boldsymbol{A} \succeq \boldsymbol{0}$ if and only if

$$\boldsymbol{A} = \boldsymbol{L}\boldsymbol{L}^{\mathsf{T}}$$

for some lower triangular matrix $\boldsymbol{L}$.

By default in `fminsdp`, $\boldsymbol{L}$ is restricted to the space of lower triangular $m \times m$-matrices with non-negative diagonal elements, referred to as $\mathcal{L}^m$. Non-negativity of the diagonal elements is not required by the theory, but enforcement of this condition is often important for computational efficiency.

     `fminsdp` makes use of the function $\mathsf{svec} : \mathbb{R}^{m \times m} \to \mathbb{R}^p$, where p denotes the number of non-zero elements in the Cholesky factor $\boldsymbol{L}$ of a given matrix; i.e., the number of elements in the sparsity pattern of $\boldsymbol{L}$, obtained by a *symbolic* Cholesky factorization of the given matrix. svec takes, column-wise, the elements of the input matrix corresponding to potential non-zeros of $\boldsymbol{L}$ and stacks them on top of each other to form a vector of length p. If the lower triangular part of $\boldsymbol{L} \in \mathcal{L}^m$ is full we have $p = m(m+1)/2$ and

$$\mathsf{svec}(\boldsymbol{A}) = (A_{11}, A_{21}, \ldots, A_{m1}, A_{22}, \ldots, A_{m2}, \ldots, A_{mm})^{\mathsf{T}};$$

i.e., the vector $\mathsf{svec}(\boldsymbol{A})$ contains the elements of the lower triangular part of $\boldsymbol{A}$.

     Now, based on Theorem 1, problem (2) is reformulated into the following problem:

$$
\begin{aligned}
& \underset{\boldsymbol{x} \in \mathbb{R}^n,\, \boldsymbol{\ell}_1 \in \mathbb{R}^{p_1},\, \ldots,\, \boldsymbol{\ell}_q \in \mathbb{R}^{p_q}}{\text{minimize}} \quad f(\boldsymbol{x}) \\
& \text{subject to} \quad
\begin{cases}
\boldsymbol{A}_{eq}\boldsymbol{x} = \boldsymbol{b}_{eq} \\
\boldsymbol{A}\boldsymbol{x} \le \boldsymbol{b} \\
\boldsymbol{c}_{eq}(\boldsymbol{x}) = \boldsymbol{0} \\
\boldsymbol{c}(\boldsymbol{x}) \le \boldsymbol{0} \\
\boldsymbol{l} \le \boldsymbol{x} \le \boldsymbol{u} \\
\tilde{\boldsymbol{l}} \le (\boldsymbol{\ell}_1, \ldots, \boldsymbol{\ell}_q) \le \tilde{\boldsymbol{u}} \\
\mathsf{svec}\left(\boldsymbol{\mathcal{A}}_i(\boldsymbol{x}) - \boldsymbol{L}_i(\boldsymbol{\ell}_i)\boldsymbol{L}_i(\boldsymbol{\ell}_i)^{\mathsf{T}}\right) = \boldsymbol{0}, \qquad i = 1, ..., q \\
\mathsf{diag}\{\boldsymbol{L}_i(\boldsymbol{\ell}_i)\} \ge \boldsymbol{0}, \qquad\qquad\qquad i = 1, ..., q,
\end{cases}
\end{aligned}
\tag{3}
$$

where the variables $\boldsymbol{\ell}_i$ defining the non-zero elements of the Cholesky factors are referred to as *auxiliary variables*, $\tilde{\boldsymbol{l}}$ are $\tilde{\boldsymbol{u}}$ are constant vectors, and `diag` returns a vector containing the diagonal elements of a matrix. This problem is in a form amenable to direct treatment by an NLP-solver, and this is the problem to which `fminsdp` attempts to find a local solution.

Note. Any local minimum of problem (3) is also a local minimum for (2), and vice versa (assuming that bounds on the auxiliary variables do not prevent this). However, problem (3) may have additional stationary points not present in (2). The author has not experienced any difficulties obviously attributed to this fact, but it cannot be ruled out that it might cause trouble on some problems.

     A key motivation behind `fminsdp` is to abstract away the exact treatment of the matrix inequality constraints so that the user only "sees" a problem of the form (2); that is, the user should not have to deal with the Cholesky factors and the auxiliary variables. Thus, when using `fminsdp` one only works with the primary variables $\boldsymbol{x}$, and user-supplied derivatives are only with respect to $\boldsymbol{x}$. The examples in Section 7 and the folder `examples` shows how this is done in practice.

2.2   The ldl-method

In the ldl-method [12, 2], `fminsdp` reformulates the matrix inequality constraints into scalar inequality constraints using the fact that a matrix is positive semi-definite if and only if the diagonal elements in an LDL-factorization of the matrix is non-negative. In other words,

Theorem 2. A symmetric $m \times m$-matrix $\boldsymbol{A} \succeq \boldsymbol{0}$ if and only if

$$d_i(\boldsymbol{A}) \geq 0, \quad i = 1, \ldots, m, \tag{4}$$

where $d_i(\boldsymbol{A})$ are the diagonal elements in an LDL-factorization of $\boldsymbol{A}$.

The original problem (2) is now replaced by

$$
\begin{aligned}
&\underset{\boldsymbol{x} \in \mathbb{R}^n}{\text{minimize}} \quad f(\boldsymbol{x}) \\
&\text{subject to} \quad
\begin{cases}
\boldsymbol{A}_{eq}\boldsymbol{x} = \boldsymbol{b}_{eq} \\
\boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{b} \\
\boldsymbol{c}_{eq}(\boldsymbol{x}) = \boldsymbol{0} \\
\boldsymbol{c}(\boldsymbol{x}) \leq \boldsymbol{0} \\
\boldsymbol{l} \leq \boldsymbol{x} \leq \boldsymbol{u} \\
d_{ij}(\boldsymbol{\mathcal{A}}_i(\boldsymbol{x})) \geq 0, \quad j = 1, \ldots, m_i, \ i = 1, \ldots, q,
\end{cases}
\end{aligned}
\tag{5}
$$

The functions $d_{ij} : \mathbb{S}^m \to \mathbb{R}$ are smooth, and concave, on the set of positive definite matrices [12], so provided $\boldsymbol{\mathcal{A}}_i(\boldsymbol{x})$ are always positive definite, (5) is a smooth NLP. In practise, even if $\boldsymbol{\mathcal{A}}_i(\boldsymbol{x})$ are all positive definite at the initial point (see section 3), this property can not be ensured throughout the optimization process without taking special measures. When using `fmincon` the line search step-length can be reduced until, due to continuity, $\boldsymbol{\mathcal{A}}_i(\boldsymbol{x})$ is positive definite. When using the NLP-solver `gcmma` the subproblems are made more conservative to reach the same effect.

Compared to the cholesky-method, the ldl-method introduces no additional variables, but requires, e.g., shortening of the search step in line-search procedures, and derivatives of the matrix constraints can be more expensive to compute.

When using NLP-solver `gcmma` the computational cost can sometimes be reduced significantly by using an active-set approach [2] where the gcmma subproblem only takes into accounts those constraints which satisfies

$$d_{ij}(\boldsymbol{\mathcal{A}}_i(\boldsymbol{x})) \leq \eta,$$

where η is a positive constant. This approach can be efficient for problems with low-rank solutions, and is activated by setting the option `eta` (see Section 6.1) to some suitable value.

2.3   Penlab

The reader is referred to [5] and references therein for a description of the algorithm implemented in PENLAB.

3   Infeasible initial points

It is recommended that the user supplies an initial point which is feasible with respect to the constraints of problem (2), in particular the matrix inequality constraints. If the latter is not possible, an option can be set (see section 6.1) so that `fminsdp` attempts to solve the following

problem in place of (2):

$$\underset{\boldsymbol{x}\in\mathbb{R}^n,\,s\in\mathbb{R}}{\text{minimize}}\quad f(\boldsymbol{x}) + cs$$

$$\text{subject to}\quad\begin{cases}\boldsymbol{A}_{eq}\boldsymbol{x} = \boldsymbol{b}_{eq} \\ \boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{b} \\ \boldsymbol{c}_{eq}(\boldsymbol{x}) = \boldsymbol{0} \\ \boldsymbol{c}(\boldsymbol{x}) \leq \boldsymbol{0} \\ \boldsymbol{l} \leq \boldsymbol{x} \leq \boldsymbol{u} \\ \boldsymbol{\mathcal{A}}_i(\boldsymbol{x}) + s\boldsymbol{I}^{m_i} \succeq \boldsymbol{0}, \qquad i = 1, ..., q \\ \underline{s} \leq s \leq \overline{s}.\end{cases} \qquad (6)$$

Here s is an auxiliary variable, $\boldsymbol{I}^{m_i}$, $i = 1, ..., q$, are identity matrices of size $m_i \times m_i$, and $\underline{s}$ and $\overline{s}$ are constants. The constant c appearing in the objective should be set to some large positive number (finding a suitable value might require experimenting a bit) such that s becomes close to zero at a solution.

Unless the user specifies something else, given $\boldsymbol{x}_0$ `fminsdp` will generate an initial value for the auxiliary variable according to

$$s_0 = \max_{i=1,...,q}\Big(-\min\{\lambda_1(\mathcal{A}_i(\boldsymbol{x}_0)), -10^{-12}\}\Big),$$

where $\lambda_1(\cdot)$ returns the smallest eigenvalue of a matrix. This guarantees that $\boldsymbol{\mathcal{A}}_i(\boldsymbol{x}_0) + s_0\boldsymbol{I}^{m_i} \succ \boldsymbol{0}$ for all i.

Note. The ldl-method requires a feasible initial point, so if you're unable to supply such a point you must set $c > 0$. The cholesky- and penlab-methods do not require a feasible initial point.

4   Problem size and some limitations

Assuming the constraint matrices are of small size or very sparse, `fminsdp` might be able to solve large-scale problems within reasonable time. The ldl-method might also work for problems involving large constraint matrices. For the cholesky-method, one can expect a fairly large number of auxiliary variables and thus many non-zero elements in the constraint Jacobian (and Hessian of the Lagrangian), resulting in much memory and CPU time being devoted to solution of linear systems by the NLP solver used to solve (3). Penlab, finally, relies on exact second-derivative information which can make problems costly to solve.

Here are some additional things to note:

1. NLP-solvers ipopt, knitro, snopt, penlab, mma and gcmma must be downloaded and installed separately.

2. Penlab requires a user-supplied function for evaluating the Hessian of the Lagrangian.

3. The ldl-method can currently not make use of a user-supplied Hessian of the Lagrangian and has only been tested with NLP-solvers fmincon and gcmma.

4. Although it is possible to improve performance by implementing parts of the code as MEX-functions, `fminsdp` is a pure Matlab-code. The reason is to make installation, set up and maintenance as easy as possible. The interested user could try out the Matlab-function `profile` to identify bottlenecks in the code and implement them as MEX-functions instead.

5   Alternatives

There are many important special cases of problem (2) for which specialized solvers are available. Although this is possible, `fminsdp` is not intended to replace such solvers, which, when applicable, can be a lot faster.

The following is an incomplete list of currently available solvers:

- Problems with linear matrix inequality (LMI) constraints:
  SeDuMi, SDPT3, SDPA, LMI Lab, PENSDP, BMISolver, PENBMI, PENNON

- Problems with bilinear matrix inequality (BMI) constraints:.
  BMISolver, PENBMI, PENNON, PENLAB

- General problems of type (2):
  PENNON, PENLAB

Note that except for PENNON, and its open-source version PENLAB, there are additional constraints on the structure of the problems treated by the listed solvers.

The Matlab code YALMIP [8] provides a convenient unified interface to the solvers listed above (except BMISolver).

One could of think of other criteria, not without drawbacks of course, for positive semi-definiteness beside the ones used here that might be used to obtain a standard NLP formulation of (2). Non-negativity of the smallest eigenvalue or of the leading principle minors of a matrix are both necessary and sufficient criteria. If one is satisfied with sufficiency, diagonal dominance could also be considered.

6 Using `fminsdp`

A user of `fminsdp` is expected to provide two functions; one for evaluating the objective function (i) and one for evaluating the non-linear constraints (ii):

i. `[fval,`*`grad`*`] = objfun(x)`

ii. `[cineq,ceq,`*`cineqgrad`*`,`*`ceqgrad`*`] = nonlcon(x)`

(Here and in the following it is assumed that the reader is familiar with the Matlab syntax.) The return arguments written in italics are optional and only needed if the user wishes to provide gradients for the objective and constraints. The matrix inequality constraints are defined in the non-linear constraints function (even linear matrix inequalities). The values of the constraint matrices, vectorized using `svec`, are returned as the last elements in the output vector `ceq`; see Section 7 for an example.

In addition to objective and non-linear constraints functions, the user may optionally provide a function for evaluating the Hessian of the Lagrangian:

- `H = hessian(x,lambda)`,

where `lambda` is a struct with two fields `ineqnonlin` and `eqnonlin` containing values of the Lagrange multipliers associated with the non-linear inequality and equality constraints, respectively, and one field `sigma` which only used when running with NLP-solver Ipopt.

The calling syntax of `fminsdp` is similar to that of `fmincon`:

```
>> [x,fval,exitflag,output,lambda,grad,hessian] = ...
                     fminsdp(objfun,x0,A,b,Aeq,beq,lb,ub,nonlcon,options)
```

Unlike fmincon, at least 9 input arguments are required. If your problem has no linear constraints and no simple bounds, the corresponding input arguments can be set to empty matrices. Note that `fminsdp` does not support passing additional arguments to the objective or constraint functions through an 11-th input argument. To pass additional arguments to the user-supplied functions one should use anonymous or nested functions.

6.1   Options

A number of options are available when calling `fminsdp`. These are passed to the code in the form of a struct containing parameter-value pairs. The most convenient way to specify options is to use the function `sdpoptionset`:

```
>> options = sdpoptionset('MatrixInequalities',true);
```

   In addition to the options accepted by the function `optimset` from the Matlab Optimization Toolbox, the following options are available when calling `fminsdp`:

`MatrixInequalities`   logical scalar
Indicates whether or not the problem has any matrix inequality constraints. If not, `fminsdp` will simply call the NLP solver specified using options.NLPsolver.
Default: true

`Aind`   scalar or numeric array
Marks the beginning of each matrix constraint in the vector `ceq` returned from the non-linear constraints function. If the option `sp_pattern`, described below, is used it is only necessary to mark the beginning of the **first** matrix constraint.
Default: 1

`method`   {'cholesky', 'ldl', 'penlab'}
Select treatment of matrix inequality constraints.
Default: 'cholesky'

`NLPsolver`   {'fmincon', 'ipopt', 'snopt', 'knitro', 'mma', 'gcmma'}
Select NLP-solver (not applicable if `method`='penlab'). Interfaces are provided to fmincon, Ipopt [13], SNOPT [6], KNITRO [3], MMA/GCMMA [9], and PENLAB. NOTE: If you run Ipopt older than 3.11.0, make sure to modify the file ipopt_main.m appropriately.
Default: 'fmincon'

`max_cpu_time`   positive scalar
Maximum CPU time. Applicable to NLP-solvers fmincon, ipopt, mma and gcmma.
Default: inf

`sp_pattern`   (sparse) matrix or cell array of matrices
Sparsity patterns of the matrix constraints. If this option is used, the user must provide one matrix for each matrix constraint.
Default: []

`L0`   (sparse) matrix or cell array of matrices
Initial guess for the Cholesky factors of the matrix constraints. If the user does not provide an initial guess, `fminsdp` will generate one by attempting a Cholesky factorization of the constraint matrices at the initial point. If this fails, `fminsdp` will add a multiple of the identity matrix to the constraint matrices until all of them are positive definite and use the Cholesky factorizations of these matrices as an initial guess.
Default: []

`Ldiag_low`   scalar or numeric array
Lower bound(s) on the diagonal elements of the Cholesky factors.
Default: 0

`L_low`   scalar or array of doubles
Lower bound(s) on the off-diagonal element of the Cholesky factors
Default: -inf

`L_upp`   scalar or array of doubles
Upper bound(s) on the elements of the Cholesky factors (including the diagonal elements).
Default: inf

`eta`   positive scalar
Tolerance for determining the active set when using the ldl-method with NLP-solver gcmma.
Default: inf

`c`   non-negative scalar
If $c > 0$, `fminsdp` attempts to solve problem (6) instead of (2). The user may in this case also
let the objective function be empty; i.e., the first input argument to `fminsdp` can be set to ']['.
This is useful when one wants to check feasibility of one or more matrix inequalities.
Default: 0

`s_low`   scalar
Lower bound on the auxiliary variable s in problem (6).
Default: 0

`s_upp`   scalar
Upper bound on the auxiliary variable s in problem (6).
Default: inf

`HessianCheck`   {'on', 'off'}
Simple check of Hessian of the Lagrangian against finite differences at the initial point. Assumes
you have the code DERIVEST, which must be obtained separately, on your Matlab path. This
check can be very time consuming and should preferably be carried out on small instances of a
problem.
Default: 'off'

`HessMult`   {function_handle, 'on'}
If using fmincon with `options.SubProblemAlgorithm = 'cg'`, you can work with Hessian times
vector products directly, thereby avoiding the formation of the full Hessian of the Lagrangian.
Set to a function handle or simply to 'on' if the Hessian of the Lagrangian with respect to the
primary variables is zero. Only applicable when using the cholesky-method.
Default: []

`ipopt`   struct
Options to be passed on to NLP solver Ipopt. Please refer to the Ipopt documentation for a list
of available options.
Default: []

`eigs_opts`   struct
Options passed to the Matlab function `eigs` used for eigenvalue computations.
Default: struct('isreal',true,'issym',true)

`KnitroOptionsFile`   character array
Name of an options file to be read by NLP solver KNITRO. Please refer to the KNITRO docu-
mentation for a list of available options.
Default: []

`SnoptOptionsFile`   character array
Name of an options file to be read by NLP solver SNOPT. Please refer to the SNOPT documen-
tation for a list of available options.
Default: []

`GradPattern`   numeric array
Sparsity pattern for the gradient of the objective function. Only used by NLP solver SNOPT and
only effective if you also set `options.JacobPattern` (see the fmincon documentation for details
on the latter).
Default: []

6.2   Output

The available output arguments from `fminsdp` are the same as those of fmincon:

`>> [x,fval,exitflag,output,lambda,grad,hessian] = fminsdp(...)`

The only difference is that the struct `output`, the fourth output argument, contains some additional fields (some are only available when running the cholesky-method):

`A`   cell array of matrices
Constraint matrices evaluated at the solution `x`.

`L`   cell array of matrices
Cholesky factors of the constraint matrices evaluated at the solution `x`. These are computed by assembling each $\boldsymbol{L}_i$ from the variable vector $\boldsymbol{\ell}_i$ and not by a Cholesky factorization of the corresponding constraint matrix.

`L0`   cell array of matrices
Cholesky factors of the constraint matrices evaluated at the initial point `x0`.

`nxvars`   numeric scalar
Number of primary variables.

`nLvars`   numeric scalar
Number of auxiliary variables.

`A_size`   numeric array
Size of the constraint matrices; i.e., m_i in (3).

`nMatrixConstraints`   numeric scalar
Number of matrix constraints.

`NLPsolver`   string
Selected NLP solver.

It the user has set `options.c` to some positive number in order to solve problem (6), then two additional fields are available:

`s0`   double scalar
Initial value for the auxiliary variable s.

`s`   double scalar
Value of s at the solution.

7   A tutorial example

To aid the user, a very simple tutorial example is provided here. For more details, please refer to the examples found in the `examples`-folder.

Consider the following (non-sense) problem:

$$\begin{aligned}
\underset{\boldsymbol{x}\in\mathbb{R}}{\text{minimize}} \quad & x^2 \\
\text{subject to} \quad & \begin{cases} x^3 = 0 \\ x^2 \leq 0 \\ \begin{pmatrix} x & x^2 \\ x^2 & 0 \end{pmatrix} \succeq \boldsymbol{0} \end{cases}
\end{aligned} \tag{7}$$

To solve this problem using `fminsdp` we need to implement at least two functions:

1. The objective function:

```
function [fval,grad] = objfun(x)
fval = x^2;
if nargout>1
    grad = 2*x;
end
```

2. The non-linear constraints function:

```
function [cineq,ceq,cineqgrad,ceqgrad] = nonlcon(x)
cineq = x^2;
ceq = [x^3; svec([x x^2; x^2 0]))];
if nargout>2
    cineqgrad = 2*x;
    ceqgrad = [3*x^2 svec([1 2*x; 2*x 0])'];
end
```

A function for evaluating the Hessian of the Lagrangian is optional, but recommended. One such function is given here:

```
function H = hessian(x, lambda)
H = lambda.ineqnonlin*2 + lambda.eqnonlin(1)*6*x  + ...
    lambda.eqnonlin(2:end,1)'*svec([0 2; 2 0]));
```

Using the functions specified above, a simple script to solve problem (7) can now be written:

```
% Mark the beginning of the matrix inequality constraints in the vector ceq
% returned from nonlcon
options.Aind = 2;

% Specify that analytical gradients should be used
options.GradObj = 'on';
options.GradConstr = 'on';

% Specify that the function "hessian" should be used for the
% Hessian of the Lagrangian
options.Hessian = 'user-supplied'
options.HessFcn = @(x,lambda) hessian(x,lambda);

% Specify initial point
x0 = 1;

% Call fminsdp
[x,fval] = fminsdp(objfun,x0,[],[],[],[],[],[],nonlcon,options);
```

8   Practicalities

For a given problem there are usually a number of things that can be done to improve numerical performance. One can, for instance, try out various scalings or introduce new variables and constraints to reduce the degree of non-linearity. Here are some additional tips:

Exploit sparsity

In many cases, the functions $\mathcal{A}_i$, $i = 1, ..., q$, are sparse in the sense that the matrix $\mathcal{A}_i(\boldsymbol{x})$ is sparse for every $\boldsymbol{x}$. This fact can be exploited to reduce the number of auxiliary variables and nonlinear constraints for the cholesky-method.

To take advantage of sparsity, the user should compute the sparsity pattern of each matrix constraint (or an (pessimistic) estimate thereof) and supply this to `fminsdp` via the option `sp_pattern`. For the dummy problem in Section 7 we simply add two lines to the driver script:

```
sp_A = [1 1; 1 0]              % Sparsity pattern of the constraint matrix
options.sp_pattern = sp_A;

...

[x,fval] = fminsdp(@(x) objfun(x),x0,[],[],[],...
    [],[],[],@(x) nonlcon(x),options);
```

Add artificial upper and lower bounds

Even though some of the variables in your problem has no "natural" bounds on them, it is usually wise to add upper and lower bounds. These bounds should of course be loose enough to not exclude any solution of interest, but even if they don't, setting them too tight might have a negative impact on the solution process. Therefore, a bit of experimenting is recommended.

Bounds on the auxiliary variables in the cholesky-method can be specified by setting `options.L_low` and `options.L_upp`. Sometimes it is easy to derive suitable bounds *a priori* [11].

Experiment with various solvers and algorithms

`fminsdp` can use fmincon, SNOPT, KNITRO, Ipopt, mma, gcmma and PENLAB to solve problems (3). In addition, fmincon and KNITRO lets the user choose between three different algorithms (two interior-point and one sqp). It is unlikely that a single solver and/or algorithm is best suited for all types of problems so it is recommended that the user experiment with various alternatives. The author's experience is that the interior-point method of fmincon (selected by setting `options.Algorithm='interior-point'` and `options.SubProblemAlgorithm = 'ldl-factorization'`) works well in terms of the number of function evaluations, but that more efficient handling of the linear algebra makes the other solvers better suited for larger problems.

Provide gradients and Hessian of the Lagrangian

Providing code that computes gradients and the Hessian of the Lagrangian is usually a good idea (if you use Ipopt you *must* provide code for evaluating gradients; SNOPT does not make use of code for evaluating the Hessian of the Lagrangian). Due to the homogeneity of `svec` one has simply

$$\frac{\partial \mathsf{svec}(\mathcal{A})}{\partial x_i} = \mathsf{svec}\left(\frac{\partial \mathcal{A}}{\partial x_i}\right),$$

and for a function $L(\boldsymbol{x}) = \boldsymbol{\lambda}^{\mathsf{T}}\mathsf{svec}(\mathcal{A}(\boldsymbol{x}))$,

$$\frac{\partial^2 L}{\partial x_i \partial x_j} = \boldsymbol{\lambda}^T \mathsf{svec}\left(\frac{\partial^2 \mathcal{A}}{\partial x_i \partial x_j}\right).$$

The user is strongly recommended to check the correctness of the derivatives by comparing against finite-difference approximations. This can be done by setting `options.DerivativeCheck='on'` and `options.HessianCheck='on'`. If necessary, the accuracy of the finite-difference approximations in `fmincon` can be increased by setting `options.FinDiffType='central'`.

If you use NLP solvers SNOPT, Ipopt or KNITRO, you can also provide sparsity patterns for the constraint Jacobian, and Ipopt and KNITRO can also exploit a sparsity pattern for the Hessian of the Lagrangian. Sparsity patterns are passed to the solvers by setting `options.JacobPattern` and `options.HessPattern` to sparse matrices.

As an alternative to deriving and implementing derivatives one might consider using automatic differentiation, which provides derivatives with accuracy to machine precision, and, at least in principle, does so completely automatically.

9   Examples

This section comprise a brief description of three problems concerned with structural optimization of (plane) trusses. Implementations of the example problem can be found in the folder `examples`.[2]

Let $n_d = 2$ be the number of spatial dimensions, N the number of nodes in the truss, and n_{fixed} the number of prescribed (to zero) displacement components. The number of displacement degrees of freedom is $n = n_d N - n_{fixed}$ and there are m potential bars in the truss. The m bar volumes are collected in a vector $\boldsymbol{x} \geq \boldsymbol{0}$ which is used to parametrize the design of the truss. Assuming (infinitesimally) small deformations and a quasi-static situation, the nodal displacement vector $\boldsymbol{u} \in \mathbb{R}^n$ satisfies the equilibrium equation

$$\boldsymbol{K}(\boldsymbol{x})\boldsymbol{u} = \boldsymbol{f}, \tag{8}$$

where $\boldsymbol{K}(\boldsymbol{x})$ is known as the stiffness matrix and $\boldsymbol{f} \in \mathbb{R}^n$ contains forces applied to the nodes. The stiffness matrix is given by

$$\boldsymbol{K}(\boldsymbol{x}) = \sum_{i=1}^{m} x_i E_i \boldsymbol{b}_i \boldsymbol{b}_i^{\mathsf{T}},$$

where x_i is the volume of the i:th bar, E_i the Young's modulus, and the vectors $\boldsymbol{b}_i$ depend on the geometry of the undeformed truss. Clearly, $\boldsymbol{K}(\boldsymbol{x})$ is positive semi-definite and symmetric. Assuming rigid body motions are prevented, by appropriate support conditions, and $\boldsymbol{x} > \boldsymbol{0}$, it is even positive definite.

1. The first problem in the `examples`-folder is to minimize the volume of a truss subject to the equilibrium condition (8) and an upper bound c on the so-called compliance $\boldsymbol{f}^{\mathsf{T}}\boldsymbol{u}$. This can be formulated as a problem involving a single LMI:

$$\begin{aligned} &\underset{\boldsymbol{x} \in \mathbb{R}^m}{\text{minimize}} \quad \sum_{i=1}^{m} x_i \\ &\text{subject to} \quad \begin{cases} \begin{pmatrix} c & \boldsymbol{f}^{\mathsf{T}} \\ \boldsymbol{f} & \boldsymbol{K}(\boldsymbol{x}) \end{pmatrix} \succeq \boldsymbol{0} \\ x_i \geq 0, \quad i = 1, \ldots, m. \end{cases} \end{aligned} \tag{9}$$

2. A drawback of problem (9) is that optimized structures may be unstable in the sense that they are prone to global buckling. One way to avoid many such designs is to impose an additional constraint [7] requiring that

$$\boldsymbol{K}(\boldsymbol{x}) + \boldsymbol{G}(\boldsymbol{u}(\boldsymbol{x}), \boldsymbol{x}) \succeq \boldsymbol{0}, \tag{10}$$

where the geometrix stiffness matrix

$$\boldsymbol{G}(\boldsymbol{u}(\boldsymbol{x}), \boldsymbol{x}) = \sum_{i=1}^{m} x_i E_i \boldsymbol{b}_i^{\mathsf{T}} \boldsymbol{u}(\boldsymbol{x}) \boldsymbol{\gamma}_i \boldsymbol{\gamma}_i^{\mathsf{T}},$$

---

[2]Note that there are often many different ways to formulate such problems [4, 1], and formulations other than those given here may certainly be better suited to one's needs. Here, however, we are specifically interested in solving problems with matrix inequality constraints.

in which $\boldsymbol{u}(\boldsymbol{x})$ denotes a solution to (8) and $\boldsymbol{\gamma}_i$ depend on the geometry of the undeformed truss.

Adding (10) to (9) leads to the following problem involving both a linear and a non-linear matrix inequality[3]:

$$
\begin{aligned}
\underset{\boldsymbol{x} \in \mathbb{R}^m}{\text{minimize}} \quad & \sum_{i=1}^{m} x_i \\
\text{subject to} \quad & \begin{cases} \begin{pmatrix} c & \boldsymbol{f}^\mathsf{T} \\ \boldsymbol{f} & \boldsymbol{K}(\boldsymbol{x}) \end{pmatrix} \succeq \boldsymbol{0} \\ \boldsymbol{K}(\boldsymbol{x}) + \boldsymbol{G}(\boldsymbol{u}(\boldsymbol{x}), \boldsymbol{x}) \succeq \boldsymbol{0} \\ x_i \geq \epsilon, \quad i = 1, \dots, m, \end{cases}
\end{aligned} \tag{11}
$$

where ϵ is a small positive number introduced to avoid singularity of the stiffness matrix. Kočvara [7] showed that a solution to this problem corresponds to a truss that will not exhibit global, linear buckling for loads of the form $\tau \boldsymbol{f}$, $\tau \in [0, 1)$.

3. The third example problem is an alternative formulation of problem (11) obtained by treating the displacements as explicit variables in the optimization problem:

$$
\begin{aligned}
\underset{\boldsymbol{x} \in \mathbb{R}^m, \boldsymbol{u} \in \mathbb{R}^n}{\text{minimize}} \quad & \sum_{i=1}^{m} x_i \\
\text{subject to} \quad & \begin{cases} \boldsymbol{f}^\mathsf{T} \boldsymbol{u} \leq c \\ \boldsymbol{K}(\boldsymbol{x}) \boldsymbol{u} = \boldsymbol{f} \\ \boldsymbol{K}(\boldsymbol{x}) + \boldsymbol{G}(\boldsymbol{u}, \boldsymbol{x}) \succeq \boldsymbol{0} \\ x_i \geq \epsilon, \quad i = 1, \dots, m. \end{cases}
\end{aligned}
$$

In this formulation, the upper bound on the compliance is a linear constraint, the equilibrium equation a set of bilinear equality constraints, and the stability constraint is a BMI.

To see how these problems can be solved with `fminsdp`, and what optimized designs might look like, please check out the Matlab codes found in the `examples`-folder.

---

[3]In practice it is perhaps better to replace the LMI in (11) by the non-linear constraint $\boldsymbol{f}^\mathsf{T} \boldsymbol{u}(\boldsymbol{x}) \leq c$ but we keep it to illustrate solution of a problem with more than one matrix inequality.

References

[1] Bendsøe M, Ben-Tal A, Zowe J (1994) Optimization methods for truss geometry and topology design. Structural Optimization 7:141–159

[2] Bogani C, Kocvara M, Stingl M (2009) A new approach to the solution of the VTS problem with vibration and buckling constraints. In: 8th World Congress on Structural and Multidisciplinary Optimization

[3] Byrd R, Nocedal J, Waltz R (2006) KNITRO: An integrated package for nonlinear optimization. In: Large-Scale Nonlinear Optimization, Springer, pp 35–59

[4] Christensen P, Klarbring A (2009) An Introduction to Structural Optimization. Springer

[5] Fiala J, Kocvara M, Stingl M (2013) Penlab - a solver for nonlinear semidefinite programming. arXiv:13115240 URL https://arxiv.org/abs/1311.5240

[6] Gill P, Murray W, Saunders M (2002) SNOPT: an SQP algorithm for large-scale constrained optimization. SIAM Journal on Optimization 12:979–1006

[7] Kočvara M (2002) On the modeling and solving of the truss design problem with global stability constraints. Structural and Multidisciplinary Optimization 23:189–203

[8] Löfberg J (2004) YALMIP : A toolbox for modeling and optimization in MATLAB. In: Proceedings of the CACSD Conference, Taipei, Taiwan, URL http://users.isy.liu.se/johanl/yalmip

[9] Svanberg K (2007) MMA and GCMMA, versions September 2007. URL http://www.math.kth.se/\char'~krille/gcmma07.pdf

[10] The MathWorks Inc (2011) Optimization Toolbox - User's Guide. Natick, MA, 2000

[11] Thore CJ, Holmberg E, Klarbring A (2015) Large-scale robust topology optimization under load-uncertainty. In: 11th World Congress on Structural and Multidisciplinary Optimisation

[12] Vanderbei R, Benson H (2000) On formulating semidefinite programming problems as smooth convex nonlinear optimization problems. Tech. rep., Center for Discrete Mathematics & Theoretical Computer Science

[13] Wächter A, Biegler L (2006) On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. Mathematical programming 106:25–57