

Polyhedral Approaches for Mixed-integer Convex Optimization

GERAD

Miles Lubin

February 8, 2019

Google, New York City (This work was conducted at MIT.)

Optimization

Given: Feasible region \mathcal{Z} and objective function $f : \mathbb{R}^n \rightarrow \mathbb{R}$

Find: \mathcal{Z} with lowest possible value of $f(\cdot)$ (global solution)

Optimization

Given: Feasible region and objective function $f : \mathbb{R}^n \rightarrow \mathbb{R}$

Find: x^* with lowest possible value of $f(x^*)$ (global solution)

How to solve?

- Local search (e.g., gradient descent), genetic algorithms, ...
- Approximate or reformulate (if lucky) as another optimization problem we know how to solve to a global solution

Which optimization problems can we solve to global optimality?

Which optimization problems can we solve to global optimality?

- Integer programming
- Convex optimization

In this talk, we consider **mixed-integer convex optimization**, or mixed-integer convex programming (MICP).

These are convex optimization problems with integrality constraints.



We look for global solutions or at least global bounds.

We apply extended formulations, conic duality, and branch and cut to get a state-of-the-art open-source MICP solver.

This work motivated developments that will appear in MOSEK 9 (presented at ISMP 2018).

What can you model with MICP?

It is possible to model a broad scope of problems using MILP. What can you do if allowed to model with convex constraints?



Ceria and Soares (1999) develop MICP formulations for unions of certain (e.g., bounded) convex sets.

See L., Zadik, and Vielma (IPCO 2017, arXiv 1706.05135) for much more on this.

Mixed-integer sum-of-squares optimization

- Sum-of-squares constraints (which are convex) can be used to model polynomial constraints, including trajectories of a moving object
- Control decisions with discrete aspect (e.g., minimize number of thrusts) yield integer constraints

Credit: Joey Huchette (SIAM Conference on Optimization, 2017)

We define MICP general form as:

$$\begin{aligned} \min_{z \in \mathbb{R}^o} & \quad : \\ & \quad z \in X; \\ & \quad z \in Z; \quad \text{and } z_{j+1} = f_j(z_1, \dots, z_j); \end{aligned}$$

where X is a closed, convex set, and the first o variables are constrained to take integer values.

We can always assume objective is linear, via an epigraph transformation.

When X is a polyhedron, this reduces to MILP.

- Branch & bound – recursively partition possible assignments to integer variables and use convex relaxations to prune the search tree (Gupta and Ravindran, 1985; Bonami et al., 2013)

- Branch & bound – recursively partition possible assignments to integer variables and use convex relaxations to prune the search tree (Gupta and Ravindran, 1985; Bonami et al., 2013)
- **Iterative outer approximation (OA)** – based on polyhedral relaxation of convex constraints, solve a sequence of MILP relaxations and convex subproblems (Duran and Grossmann, 1986; Leyffer, 1993; Bonami et al., 2008)

- Branch & bound – recursively partition possible assignments to integer variables and use convex relaxations to prune the search tree (Gupta and Ravindran, 1985; Bonami et al., 2013)
- **Iterative outer approximation (OA)** – based on polyhedral relaxation of convex constraints, solve a sequence of MILP relaxations and convex subproblems (Duran and Grossmann, 1986; Leyffer, 1993; Bonami et al., 2008)
- **LP-based branch & bound** – use polyhedral relaxations within a single branch & bound (B&B) search tree (Quesada and Grossmann, 1992; Bonami et al., 2008)

Developing polyhedral relaxations

Suppose:

$$X = \{x \in \mathbb{R}^n : (x, 0) \in J\}$$

Then under assumptions on J :

$$X = \{x \in \mathbb{R}^n : (x, 0) + r(x, 0) \in J, r \in \mathbb{R}^+, J \text{ convex}\}$$

Any finite subset of linearization points yields a polyhedral outer approximation.


Given a polyhedral outer approximation \tilde{X} , consider:

$$\begin{aligned} \min \quad & \\ & \in \tilde{X}; \\ & \in \mathbb{Z}; \end{aligned} \quad \text{8 2JK:}$$

The above problem is a mixed integer linear relaxation.

Algorithmic idea: Iteratively improve the relaxation until lower bound matches feasible solution.

An illustration:



What can go wrong?

Classical approaches form polyhedral outer approximations by a finite collection of “gradient linearizations.” But a good polyhedral outer approximation might need **too many linear constraints**.

Recall the ℓ_1 ball:

$$B_1 = \{x \in \mathbb{R}^n : \|x\|_1 \leq 1\} = \left\{ x \in \mathbb{R}^n : \sum_{j=1}^n |x_j| \leq 1 \right\}$$

$$B_1 = \begin{pmatrix} f & 2R & : & j & j & 1 \\ g & 2R & : & & & \end{pmatrix} \quad X \quad \begin{pmatrix} 1 \\ 2f \\ 1 \end{pmatrix} + 1g$$

$= 1$

Standard LP extended formulation:

$$B_1 = \begin{pmatrix} & 2R & : & 9 & 2R & : \\ & & & & & \end{pmatrix} \quad X \quad \begin{pmatrix} \\ \\ 1 \end{pmatrix}$$


$= 1$

The new formulation has 2 variables and 2 + 1 constraints versus variables and 2 constraints.

The same happens with smooth constraints

Hijazi et al. (2014) consider the set:

$$B_2 = \left\{ x \in \mathbb{R}^2 : \|x\|_2 \leq 1, \frac{x_1}{2} \leq \frac{x_2}{4} \right\}$$



How to fix it?

Consider the equivalent extended formulation ($B_2 = \text{proj } \hat{B}_2$):

$$\hat{B}_2 = \left\{ (x, y) \in \mathbb{R}^2 : \begin{array}{l} x \geq -1 \\ x \leq \frac{1}{4} \\ y \geq \frac{1}{2}x^2 \\ y \leq 8 - 2x \end{array} \right\}$$

A polyhedral outer approximation in the space of (x, y) needs only 2 hyperplanes to exclude all integer points.

Say a user writes down a constraint:

$$f(x) + g(x) \leq c$$

If f and g are convex, then extended formulation is obtained by introducing $x_1; x_2$ such that:

$$x_1 + x_2 \leq c; \quad x_1 \leq f(x); \quad x_2 \leq g(x):$$

But $f + g$ convex doesn't imply f and g convex (e.g., $f(x) = \frac{x^2}{1}$ and $g(x) = 2 - \frac{x^2}{2}$).

Say a user writes down a constraint:

$$f_1(x) + f_2(x) \leq c$$

If f_1 and f_2 are convex, then extended formulation is obtained by introducing λ_1, λ_2 such that:

$$\lambda_1 + \lambda_2 \leq c; \lambda_1 f_1(x); \lambda_2 f_2(x):$$

But $f_1 + f_2$ convex doesn't imply f_1 and f_2 convex (e.g., $f_1(x) = \frac{x^2}{1}$, $f_2(x) = \frac{x^2}{2}$ and $f_3(x) = 2 \frac{x^2}{2}$).

So if we have access to the algebraic representation, in principle need to solve a subproblem of convexity detection before constructing extended formulation.

Convexity detection is hard.

Say a user writes down a constraint:

$$f(x) + g(y):$$

If f and g are convex, then extended formulation is obtained by introducing $z_1; z_2$ such that:

$$z_1 + z_2; \quad z_1 \leq f(x); \quad z_2 \leq g(y):$$

But $f + g$ convex doesn't imply f and g convex (e.g., $f(x) = \frac{x^2}{1}$ and $g(y) = 2 \frac{y^2}{2}$).

So if we have access to the algebraic representation, in principle need to solve a subproblem of convexity detection before constructing extended formulation.

Convexity detection is hard. (But some solvers try.)

Conic optimization solves the problem

$$\min_{x \in K} c^T x$$

A set $K \subseteq \mathbb{R}^n$ is a **closed convex cone** if it is closed and contains all conic combinations of its points, i.e.,

$$\theta_1 x_1 + \theta_2 x_2 \in K \quad \forall \theta_1, \theta_2 \geq 0, \theta_1 x_1 + \theta_2 x_2 \in K:$$

An example

The constraint

$$\log \sum_{i=1}^n \exp(x_i) \leq 1$$

is convex, not separable. Converting it to conic form exposes the summation structure.

An example

The constraint

$$\log \sum_{i=1}^n \exp(x_i) \leq \alpha$$

is convex, not separable. Converting it to conic form exposes the summation structure.

$$\sum_{i=1}^n \exp(x_i) \leq \exp(\alpha)$$

$$\sum_{i=1}^n \exp(x_i) \leq 1$$

$$g : \sum_{i=1}^n \exp(x_i) \leq 1 \text{ and } g \in E;$$

where $E = \text{cone}(\{ \sum_{i=1}^n \exp(x_i) \leq 1 \})$.

In general...

- “Lectures on Modern Convex Optimization” by Ben-Tal and Nemirovski is the canonical source for how to write a problem in conic form.
- Disciplined convex programming (DCP) provides a systematic and easy-to-use approach to translating algebraic expressions into conic form (ž„ Š, ž„ Še< , žêã · ÖÜ)
- If not convinced, read more at “Polyhedral Approximation in Mixed-integer Convex Optimization”, Math Programming B, 2017.

Which cones are needed?

- Nonnegative orthant $R_+ = \{f \in \mathbb{R}^n : f \geq 0\}$
- Second-order cone $L^{1+} = \{(x, y) \in \mathbb{R}^{1+n} : \|x\| \leq y\}$
- Exponential cone $E = \{(x, y, z) \in \mathbb{R}^3 : y > 0, z \geq \exp(-x/y)\}$
- PSD cone $S_+ = \{f \in \mathbb{S}^n : f \succeq 0\}$
- Power cone $W = \{(x, y, z) \in \mathbb{R}^3 : x \geq 0, y \geq 0, z \geq 0, x^2 + y^2 \geq z^2\}$

We categorized the 333 mixed-integer convex instances in the MINLPLIB2 library according to conic representability. Excluding R_+ , the following cones are needed:

L only	E only	L and E	W only	Total
217	107	7	2	333

- $\hat{A} \hat{Y} \hat{a} \hat{y}$, $L+E$ -representable
 - 145 variables, 1268 constraints (1158 linear)
 - Using $\hat{z}eL^* \hat{S}$ as MILP solver and JS=rhY as convex solver, iterative OA solved in 6 iterations (< 10 hours)
 - Optimal solution is 21380.20 (previous best bound was 1735.06, and best known solution was 21516.83)
- $\hat{U} \hat{y}$ and $\hat{U} \hat{y}$, L -representable, unsolved since 2001
 - Solved in less than 24h by converting to MISOCP and then using 5 $\hat{u} \hat{e} \hat{E}$

Mosek 9 adds support for the exponential and power cones (first commercial solver).

Ç öý ³ê-ý âêý. Ù -êâ ýÛË³. ý Êýâö
Êýâö Ýã³. ùý. ã ö³Á

Recall:

- Iterative OA: Solve a sequence of MILP relaxations
- LP-based B&B: Use a single B&B tree with relaxations computed by solving LPs

How to generate cuts in the conic case?

A set $K \subseteq \mathbb{R}^n$ is a **closed convex cone** if it is closed and contains all conic combinations of its points, i.e.,

$$\lambda_1 x_1 + \lambda_2 x_2 \in K \quad \forall \lambda_1, \lambda_2 \geq 0, \lambda_1 x_1 + \lambda_2 x_2 \in K: \quad (1)$$

The **dual cone** of a set is defined as:

$$K^* = \{ y \in \mathbb{R}^n : \langle y, x \rangle \geq 0, \forall x \in K \}$$

If K is a closed convex cone, then:

$$K^{**} = K, \quad \forall y \in K^* : \langle y, x \rangle \geq 0, \forall x \in K :$$

Any finite subset $Z \subseteq K$ yields a polyhedral outer approximation of K . That is:

$$\hat{K} := \bigcap_{R \in \mathcal{R}} R \quad ; \quad \hat{K} \supseteq K \quad ; \quad \hat{K} \supseteq Z$$

\hat{K} is a polyhedron that contains K .

We say that an element $R \in \mathcal{R}$ corresponds to a K cut.

- The second-order cone L and PSD cone P are **self-dual**
- The exponential cone E is self dual, but E is easily described
- K cuts can be validated by checking membership in dual cone
- K cuts can be safely rounded by projection onto boundary of dual cone

Compare with gradient inequalities:

$$X = f \in \mathbb{R} : (\cdot) \quad \mathbf{0}g:$$

Then \quad is a gradient inequality iff there exists $\theta \in \mathbb{R}$ such that:

$$\begin{aligned} &= r \quad (\theta); \\ &= r \quad (\theta) \quad \theta \quad (\theta): \end{aligned}$$

This is much harder to check.

Primal

$$\min \quad : \\ \mathcal{L} K$$

Dual

$$\max \quad : \\ + \quad = 0 \\ \mathcal{L} K$$

Reduces to LP duality when $K = \mathbb{R}_+^o$.

Conic duality can fail

The following MISOCP instance cannot be solved by the iterative OA algorithm:

$$\begin{aligned} \min \quad & : \\ & ; ; \\ & \quad \quad \quad 2, \\ & \quad \quad \quad = 0; \\ & \quad \quad \quad 0; \\ & \quad \quad \quad 2 f_0; 1g: \end{aligned}$$

Conic duality can fail

The following MISOCP instance cannot be solved by the iterative OA algorithm:

$$\begin{aligned} \min \quad & \begin{matrix} ; \\ ; \\ ; \end{matrix} & : \\ & & 2, \\ & & = 0; \\ & & 0; \\ & & 2 f_0; 1g: \end{aligned}$$

Why not?

- The optimal solution is zero
- The dual of the root node relaxation is infeasible
- polyhedral OA is unbounded below

$$\begin{aligned} \min \quad & \\ & 2K \\ & 2Z \qquad \qquad 8 \ 2JK \end{aligned}$$

(When applying B&B, we also assume known lower and upper bounds on the integer-constrained variables.)

Certificate K cuts

Suppose we solve:

$$\begin{aligned} \min \quad & 2K \\ & () \\ & 8 \leq JK \quad (\text{assuming }) \end{aligned}$$

and obtain an optimal objective value and dual solution (ignoring multipliers on inequalities). Then

$$\begin{aligned} = \min \quad & : \\ & () \quad 0 \\ & 8 \leq JK \end{aligned}$$

Iterative OA uses certificate cuts with integer variables fixed to the optimal solution returned from the MILP relaxation. () finite convergence)

LP-based B&B uses conic certificate cuts at multiple points in the tree. Must use at least whenever a potential new incumbent integer feasible solution is found.

Recall $2K \geq 0$. Does it matter how we scale ϵ ?

LP solvers almost never guarantee exact feasibility. Suppose instead that we solve,

$$\begin{aligned} \min \quad & \epsilon \\ \text{subject to} \quad & (x, y) \in \mathcal{K} \end{aligned}$$

This has optimal value ϵ —!

Recall $2K \delta > 0$. Does it matter how we scale δ ?

LP solvers almost never guarantee exact feasibility. Suppose instead that we solve,

$$\min_{x \in \mathcal{K}} \quad c^T x$$

$$\delta \cdot 2JK$$

This has optimal value $\delta \cdot 2JK$ —!

So choose δ to make $\delta \cdot 2JK$ on the scale of the optimal objective, e.g., 0.01 .

Separation K cuts

Given δ such that $\delta \in K$, any hyperplane that separates from K can be shifted (if needed) to become a K cut.

Example: If $\emptyset \neq S_+$ then $\lambda_{\min}(A) < 0$. Let v be an eigenvector corresponding to the smallest eigenvalue of A . Then:

$$v^T A v = \lambda_{\min}(A) < 0; \quad (2)$$

so $\mathcal{Z}(S_+)$ corresponds to a K cut that separates \emptyset from S_+ .

Example: If $\bar{z} \notin S_+$ then $\lambda_{\min}(\bar{z}) < 0$. Let \bar{z} be an eigenvector corresponding to the smallest eigenvalue of \bar{z} . Then:

$$\bar{z}^T \bar{z} = \lambda_{\min}(\bar{z}) < 0; \quad (2)$$

so $\bar{z}^T (S_+)$ corresponds to a K cut that separates \bar{z} from S_+ .

Boyd & Vandenberghe (2004) uses these cuts for MISDP. Gally et al., 2017 uses separation cuts for MISOCP.

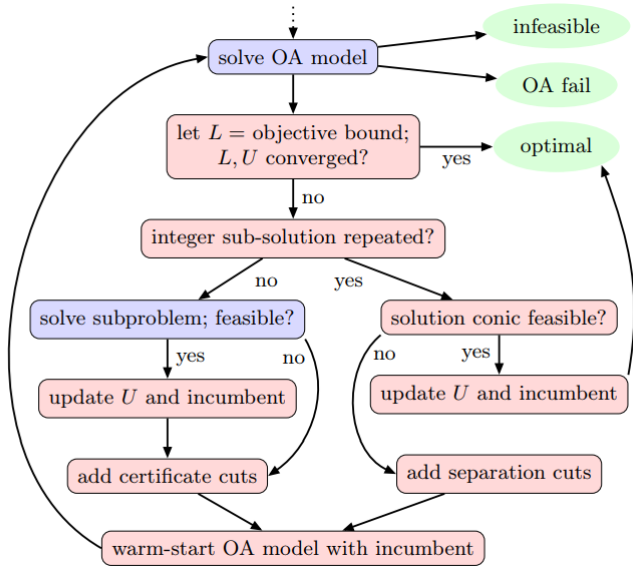
Natural extensions to the K cuts framework include:

1. Obtaining multiple K cuts from a single one by decomposing it by products of cones or into extreme rays
2. Second-order cone outer approximation of the PSD cone
3. Lifting of the second-order cone (Vielma et al., 2016)

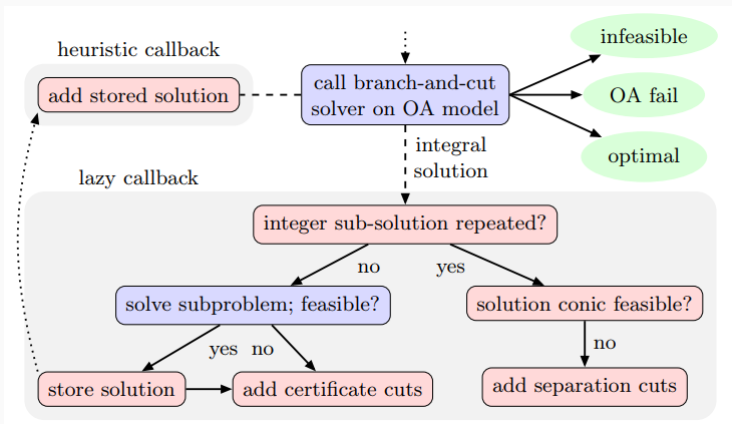
Pajarito

- An open-source MICP solver released in 2016
- Written in Julia, accessible through JuMP and via CBF files
- Can use any conic and MILP solver available to JuMP
- Implements iterative OA and LP-based B&B using callbacks

Note: For convex MINLP see Pavito. This functionality was split off last year from Pajarito.



The iterative OA algorithm in Pajarito.



The LP-based B&B algorithm in Pajarito. We call it MIP-solver-driven (MSD) because the MIP solver drives the algorithm.


Numerical experiments

- Selection of 120 MISOCP instances from CBLIB library
- Solutions with large feasibility violations marked as 'excluded'
- 'conv' means converged, 'error' means failed to converge, and 'limit' means timeout
- Time summaries presented as shifted geometric means
 $\left(\sum_{i=1}^n (t_i + \tau) \right)^{1/n}$ with $\tau = 10$ seconds


Status summary

		statuses				time (s)
solver		conv	limit	error	excluded	all conv
open source	Bonmin-BB	34	44	11	31	38.0
	Bonmin-OA	25	53	29	13	64.2
	Bonmin-OA-D	30	48	29	13	15.1
	Iter-GLPK	56	60	3	1	2.0
	Iter-CBC	78	30	3	9	1.6
restricted	SCIP	74	35	8	3	3.2
	CPLEX	90	16	5	9	0.9
	Iter-CPLEX	86	26	0	8	0.4
	MSD-CPLEX	97	20	2	1	0.4

Performance profiles



(a) Open source Bonmin (instance-wise best of 3) and Pajarito iterative solvers.



(b) CPLEX MISOCP and Pajarito MSD solvers.

Test of algorithmic variants on 95 MICP instances with a variety of cones:

- Experimental design: PSD cone and exponential cone
- Portfolio problems with risk constraints: exponential cone (entropy risk), second-order cone (norm risk), PSD cone (robust norm risk)
- Retrofit-synthesis of process networks: exponential cone, from MINLPLIB2
- Representative MISOCP instances from CBLIB

Effect of certificate cut scaling

		statuses			
		conv	limit	error	excluded
Iter	off	63	1	28	3
	on	69	1	22	3
MSD	off	60	0	30	5
	on	67	0	26	2

Significant improvement in reliability. Effect on solve time (not shown) is hard to tell.

