

Automatic Repair of Convex Optimization Problems

Shane Barratt

Guillermo Angeris

Stephen Boyd

April 8, 2020

Abstract

Given an infeasible, unbounded, or pathological convex optimization problem, a natural question to ask is: what is the smallest change we can make to the problem's parameters such that the problem becomes solvable? In this paper, we address this question by posing it as an optimization problem involving the minimization of a convex regularization function of the parameters, subject to the constraint that the parameters result in a solvable problem. We propose a heuristic for approximately solving this problem that is based on the penalty method and leverages recently developed methods that can efficiently evaluate the derivative of the solution of a convex cone program with respect to its parameters. We illustrate our method by applying it to examples in optimal control and economics.

1 Introduction

Parametrized convex optimization. We consider parametrized convex optimization problems, which have the form

$$\begin{aligned} & \text{minimize} && f_0(x; \theta) \\ & \text{subject to} && f_i(x; \theta) \leq 0, \quad i = 1, \dots, m, \\ & && g_i(x; \theta) = 0, \quad i = 1, \dots, p, \end{aligned} \tag{1}$$

where $x \in \mathbf{R}^n$ is the optimization variable, $\theta \in \mathbf{R}^k$ is the parameter, the objective function $f_0 : \mathbf{R}^n \times \mathbf{R}^k \rightarrow \mathbf{R}$ is convex in x , the inequality constraints functions $f_i : \mathbf{R}^n \times \mathbf{R}^k \rightarrow \mathbf{R}$, $i = 1, \dots, m$, are convex in x , and the equality constraint functions $g_i : \mathbf{R}^n \times \mathbf{R}^k \rightarrow \mathbf{R}$, $i = 1, \dots, p$, are affine in x .

Solvable problems. A point x is said to be feasible if $f_i(x; \theta) \leq 0$, $i = 1, \dots, m$, and $g_i(x; \theta) = 0$, $i = 1, \dots, p$. The optimal value p^* of the problem (1) is defined as

$$p^* = \inf \{ f_0(x; \theta) \mid f_i(x) \leq 0, \quad i = 1, \dots, m, \quad h_i(x) = 0, \quad i = 1, \dots, p \}.$$

We allow p^* to take on the extended values $\pm\infty$. Roughly speaking, we say that (1) is *solvable* if p^* is finite and attainable. (We will define solvable formally below, when we canonicalize (1) into a cone program.) When the problem is unsolvable, it falls into one of three cases: it is *infeasible* if $p^* = +\infty$, *unbounded below* if $p^* = -\infty$, and *pathological* if p^* is finite, but not attainable by any x , or strong duality does not hold for (1). Unsolvable problems are often undesirable since, in many cases, there does not exist a solution.

Performance metric. The goal in this paper is to repair an unsolvable problem by adjusting the parameter θ so that it becomes solvable. We will judge the desirability of a new parameter θ by a (convex) performance metric function $r : \mathbf{R}^k \rightarrow \mathbf{R} \cup \{+\infty\}$, which we would like to be small. (Infinite values of r denote constraints on the parameter.) A simple example of r is the Euclidean distance to an initial parameter vector θ_0 , or $r(\theta) = \|\theta - \theta_0\|_2$.

Repairing a convex optimization problem. In this paper, we consider the problem of repairing a convex optimization problem, as measured by the performance metric, by solving the problem

$$\begin{aligned} & \text{minimize} && r(\theta) \\ & \text{subject to} && \text{problem (1) is solvable,} \end{aligned} \tag{2}$$

with variable θ .

Pathologies. There are various pathologies that can occur in this formulation. For example, the set of θ that lead to solvable problems could be open, meaning there might not exist a solution to (2), or the complement could have (Lebesgue) measure zero, meaning that the problem can be made solvable by essentially any perturbation. Both of these cases can be demonstrated with the following problem:

$$\begin{aligned} & \text{minimize} && 0 \\ & \text{subject to} && \theta x = 1, \end{aligned} \tag{3}$$

and regularization function $r(\theta) = \theta^2$. The set of solvable θ is $\{\theta \mid \theta \neq 0\}$, which is both open and has complement with measure zero. The optimal value of problem (2) is 0, but is not attainable by any solvable θ . The best we can hope to do in these situations is to produce a minimizing sequence.

NP-hardness. Repairing a convex optimization problem is NP-hard. To show this, we reduce the 0-1 integer programming problem

$$\begin{aligned} & \text{minimize} && 0 \\ & \text{subject to} && Ax = b, \\ & && x \in \{0, 1\}^n, \end{aligned} \tag{4}$$

with variable $x \in \mathbf{R}^n$ and data $A \in \mathbf{R}^{m \times n}$ and $b \in \mathbf{R}^m$ to an instance of problem (2).

Let $r(\theta) = 0$. The convex optimization problem that we would like to be solvable be

$$\begin{aligned}
& \text{minimize} && 0 \\
& \text{subject to} && x = \theta, \\
& && \theta x = x, \\
& && Ax = b,
\end{aligned} \tag{5}$$

with variable x . Problem (2) has the same constraints as (4), since (5) is feasible if and only if $\theta_i(\theta_i - 1) = 0$, $i = 1, \dots, n$, and $A\theta = b$. Therefore, the problem of finding any feasible parameters for (1) (*i.e.*, with $r(\theta) = 0$), is at least as hard as the 0-1 integer programming problem, which is known to be NP-hard [22].

2 Cone program formulation

In practice, most convex optimization problems are solved by reformulating them as equivalent conic programs and passing the numerical data in the reformulated problem to general conic solvers such as SCS [30], Mosek [26], or Gurobi [20]. This process of canonicalization is often done automatically by packages like CVXPY [15], which generate a conic program from a high-level description of the problem.

Canonicalization. For the remainder of the paper, we will consider the canonicalized form of problem (1). The primal (P) and dual (D) form of the canonicalized convex cone program is (see, *e.g.*, [7, 8])

$$\begin{aligned}
\text{(P)} \quad & \text{minimize} && c(\theta)^T x \\
& \text{subject to} && A(\theta)x + s = b(\theta), \\
& && s \in \mathcal{K}, \\
\text{(D)} \quad & \text{minimize} && -b(\theta)^T y \\
& \text{subject to} && A(\theta)^T y + c(\theta) = 0, \\
& && y \in \mathcal{K}^*.
\end{aligned} \tag{6}$$

Here $x \in \mathbf{R}^n$ is the primal variable, $y \in \mathbf{R}^m$ is the dual variable, and $s \in \mathbf{R}^m$ is the slack variable. The functions $A : \mathbf{R}^k \rightarrow \mathbf{R}^{m \times n}$, $b : \mathbf{R}^k \rightarrow \mathbf{R}^m$, and $c : \mathbf{R}^k \rightarrow \mathbf{R}^n$ map the parameter vector θ in problem (1) to the cone program problem data A , b , and c . The set $\mathcal{K} \subseteq \mathbf{R}^m$ is a closed convex cone with associated dual cone $\mathcal{K}^* = \{y \mid y^T x \geq 0 \text{ for all } x \in \mathcal{K}\}$.

Solution. The vector (x^*, y^*, s^*) is a solution to problem (6) if

$$A(\theta)x^* + s^* = b(\theta), \quad A(\theta)^T y^* + c(\theta) = 0, \quad (s^*, y^*) \in \mathcal{K} \times \mathcal{K}^*, \quad c(\theta)^T x^* + b(\theta)^T y^* = 0. \tag{7}$$

These conditions merely state that (x^*, s^*) is primal feasible, y^* is dual feasible, and that there is zero duality gap, which implies that (x^*, y^*, s^*) is optimal by weak duality [8, §5.2.2]. Problems (1) and (6) are solvable if and only if there exists a point that satisfies (7).

Primal-dual embedding. The primal-dual embedding of problem (6) is the cone program

$$\begin{aligned}
& \text{minimize} && t \\
& \text{subject to} && \left\| \begin{bmatrix} A(\theta)x + s - b(\theta) \\ A(\theta)^T y + c(\theta) \\ c(\theta)^T x + b(\theta)^T y \end{bmatrix} \right\|_2 \leq t \\
& && s \in \mathcal{K}, \quad y \in \mathcal{K}^*,
\end{aligned} \tag{8}$$

with variables t , x , y , and s . This problem is guaranteed to be feasible since, for any $\theta \in \mathbf{R}^k$, setting $x = 0$, $y = 0$, $s = 0$, and $t = \|(b(\theta), c(\theta))\|_2$ yields a feasible point. The problem is also guaranteed to be bounded from below, since the objective is nonnegative. Taken together, this implies that problem (8) always has a solution, assuming it is not pathological.

Optimal value of (8). Let $t^* : \mathbf{R}^k \rightarrow \mathbf{R}$ denote the optimal value of problem (8) as a function of θ . Notably, we have that $t^*(\theta) = 0$ if and only if problem (1) is solvable, since if $t^*(\theta) = 0$, the solution to problem (8) satisfies (7) and therefore problem (1) is solvable. On the other hand, if problem (1) is solvable, then there exists a point that satisfies (7) and is feasible for (8), so $t^*(\theta) = 0$.

Differentiability of t^* . In practice, t^* is often a differentiable function of θ . This is the case when A , b , and c are differentiable, which we will assume, and the optimal value of problem (8) is differentiable in A , b , and c . Under some technical conditions that are often satisfied in practice, the optimal value of a cone program is a differentiable function of its problem data [2]. We will assume that t^* is differentiable, and that we can efficiently compute its gradient $\nabla t^*(\theta)$ using the methods described in [2] and the chain rule.

Reformulation. In light of these observations, we can reformulate problem (2) as

$$\begin{aligned}
& \text{minimize} && r(\theta) \\
& \text{subject to} && t^*(\theta) = 0,
\end{aligned} \tag{9}$$

with variable θ . Here we have replaced the intractable constraint in problem (2) with an equivalent smooth equality constraint. Since this problem is NP-hard, we must resort to heuristics to find approximate solutions; we give one in §3.

3 Heuristic solution method

Penalty method. One simple heuristic is to use the penalty method to (approximately) solve (9). Starting from $\theta^0 \in \mathbf{R}^k$ and $\lambda^0 > 0$, at iteration ℓ , the penalty method performs the update

$$\theta^{\ell+1} = \underset{\theta}{\operatorname{argmin}} \lambda^\ell r(\theta) + t^*(\theta), \tag{10}$$

and then decreases λ^ℓ , e.g., $\lambda^{\ell+1} = (1/2)\lambda^\ell$, until $t^*(\theta^\ell) \leq \epsilon_{\text{out}}$ for some given tolerance $\epsilon_{\text{out}} > 0$.

To perform the update (10), we must solve the unconstrained optimization problem

$$\text{minimize } L(\theta, \lambda^\ell) = \lambda^\ell r(\theta) + t^*(\theta), \quad (11)$$

with variable θ . The objective is the sum of a differentiable function and a potentially nonsmooth convex function, for which there exist many efficient methods. The simplest (and often most effective) of these methods is the proximal gradient method (which stems from the proximal point method [25]; for a modern reference see [27]). The proximal gradient method consists of the iterations

$$\theta^{\ell+1} = \mathbf{prox}_{\alpha^\ell \lambda^\ell r}(\theta^\ell - \alpha^\ell \nabla t^*(\theta^\ell)),$$

where the proximal operator is defined as

$$\mathbf{prox}_{\alpha \lambda r}(\tilde{\theta}) = \underset{\theta}{\operatorname{argmin}} \alpha \lambda r(\theta) + \frac{1}{2} \|\theta - \tilde{\theta}\|_2^2.$$

Since r is convex, evaluating the proximal operator of $\alpha \lambda r$ requires solving a convex optimization problem. Indeed, for many practical choices of r , its proximal operator has a closed-form expression [31, §6].

We run the proximal gradient method until the stopping criterion

$$\|(\theta^\ell - \theta^{\ell+1})/\alpha^\ell + (g^{\ell+1} - g^\ell)\|_2 \leq \epsilon_{\text{in}},$$

is reached, where $g^\ell = \nabla t^*(\theta^\ell)$, for some given tolerance $\epsilon_{\text{in}} > 0$ [6]. We employ the adaptive step size scheme described in [6]. The full procedure is described in algorithm 3.1 below.

Algorithm 3.1 *Finding the closest solvable convex optimization problem.*

given regularization function r , initial parameter θ^0 , penalty λ^0 , step size α^0 , iterations n_{iter} , outer tolerance ϵ_{out} , inner tolerance ϵ_{in} .

for $l = 1, \dots, n_{\text{iter}}$

1. *Compute t^* .* Compute $t^*(\theta^\ell)$.

2. *Compute gradient of t^* .* Let $g^\ell = \nabla t^*(\theta^\ell)$.

3. *Compute the gradient step.* Let $\theta^{\ell+1/2} = \theta^\ell - \alpha^\ell g^\ell$.

4. *Compute the proximal operator.* Let $\theta^{\text{tent}} = \mathbf{prox}_{\alpha^\ell \lambda^\ell r}(\theta^{\ell+1/2})$.

5. **if** $L(\theta^{\text{tent}}, \lambda^\ell) < L(\theta^\ell, \lambda^\ell)$,

Increase step size and accept update. $\alpha^{\ell+1} = (1.2)\alpha^\ell$, $\theta^{\ell+1} = \theta^{\text{tent}}$.

Inner stopping criterion. if $\|(\theta^\ell - \theta^{\ell+1})/\alpha^\ell + (g^{\ell+1} - g^\ell)\|_2 \leq \epsilon_{\text{in}}$, then set $\lambda^{\ell+1} = (1/2)\lambda^\ell$.

6. **else** *Decrease step size and reject update.* $\alpha^{\ell+1} = (1/2)\alpha^\ell$, $\theta^{\ell+1} = \theta^\ell$.

7. *Outer stopping criterion.* if $t^*(\theta^{\ell+1}) \leq \epsilon_{\text{out}}$, quit.

end for

Implementation. We have implemented algorithm 3.1 in Python, which is available online at

<https://github.com/cvxgrp/cvxpyrepair>

The interface is the `repair` method, which, given a parametrized CVXPY problem [15] and a convex regularization function, uses algorithm 3.1 to find the parameters that approximately minimize that regularization function and result in a solvable CVXPY problem. We use SCS [30] to solve cone programs and `diffcp` [2] to compute the gradient of cone programs. We require the CVXPY problem to be a disciplined parametrized program (DPP), so that the mapping from parameters to (A, b, c) is affine, and hence differentiable [1].

Until this point, we have assumed that the optimal value of problem (8) is differentiable in A , b , and c . However, in our implementation, we do not require this to be the case. So long as it is differentiable almost everywhere, it is reasonable to apply the proximal gradient method to (11). At non-differentiable points, we instead compute a heuristic quantity. For example, a source of non-differentiability is the singularity of a particular matrix; in this case, `diffcp` computes a least-squares approximation of the gradient [2, §3].

4 Examples

4.1 Spacecraft landing

We consider the problem of landing a spacecraft with a gimbaled thruster. The dynamics are

$$m\ddot{x}(t) = f(t) - mge_3,$$

where $m > 0$ is the spacecraft mass, $x(t) \in \mathbf{R}^3$ is the spacecraft position, $f(t) \in \mathbf{R}^3$ is the force applied by the thruster, $g > 0$ is the gravitational acceleration, and $e_3 = (0, 0, 1)$.

Our goal, given some initial position $x^{\text{init}} \in \mathbf{R}^3$ and velocity $v^{\text{init}} \in \mathbf{R}^3$, is to land the spacecraft at zero position and velocity at some touchdown time $T > 0$, *i.e.*, $x(T) = 0$ and $\dot{x}(T) = 0$.

We have a total available fuel M^{fuel} and a thrust limit F^{max} . This results in the constraints

$$\int_0^T \gamma \|f(t)\|_2 dt \leq M^{\text{fuel}}, \quad \|f(t)\|_2 \leq F^{\text{max}}, \quad 0 \leq t \leq T,$$

where γ is the fuel consumption coefficient. We also have a gimbal constraint

$$f_3(t) \geq \alpha \|(f_1(t), f_2(t))\|_2,$$

where α is equal to the tangent of the maximum gimbal angle.

We discretize the thrust profile, position, and velocity at intervals of length h , or

$$f_k = f((k-1)h), \quad x_k = x((k-1)h), \quad v_k = \dot{x}((k-1)h), \quad k = 1, \dots, H,$$

where $H = T/h + 1$.

To find if there exists a thrust profile to land the spacecraft, we solve the problem

$$\begin{aligned}
& \text{minimize} && 0 \\
& \text{subject to} && x_{k+1} = x_k + (h/2)(v_{k+1} + v_k), \quad k = 1, \dots, H, \\
& && mv_{k+1} = f_k - hmg e_3, \quad k = 1, \dots, H, \\
& && \|f_k\|_2 \leq F^{\max}, \quad k = 1, \dots, H, \\
& && \sum_{k=1}^H h\gamma \|f_k\|_2 \leq M^{\text{fuel}}, \\
& && (f_k)_3 \geq \alpha \|((f_k)_1, (f_k)_2)\|_2, \\
& && x_1 = x^{\text{init}}, \quad v_1 = v^{\text{init}}, \\
& && x_H = 0, \quad v_H = 0,
\end{aligned} \tag{12}$$

with variables x , v , and f . This problem is a parametrized convex optimization problem, with parameter

$$\theta = (m, M^{\text{fuel}}, F^{\max}, \alpha).$$

Suppose that we are given a parameter vector θ_0 for which it is impossible to find a feasible thrust profile, *i.e.*, problem (12) is infeasible. Suppose, in addition, that we are allowed to change the spacecraft's parameters in a limited way. We seek to find the smallest changes to the mass and constraints on the fuel and thrust limit that guarantees the feasibility of problem (12). We can (approximately) do this with algorithm 3.1.

Numerical example. We consider a numerical example with data

$$T = 10, \quad h = 1, \quad g = 9.8, \quad x^{\text{init}} = (10, 10, 50), \quad v^{\text{init}} = (10, -10, -10), \quad \gamma = 1,$$

and initial parameters

$$m_0 = 12, \quad M_0^{\text{fuel}} = 200, \quad F_0^{\max} = 50, \quad \alpha_0 = 0.5.$$

The initial parameters are infeasible, *i.e.*, there is no possible thrust profile which allows the spacecraft to land in time, so we use algorithm 3.1 to modify the design parameters in order to have a feasible landing thrust profile. We use the performance metric

$$r(\theta) = \begin{cases} \frac{|m-m_0|}{m_0} + \frac{|M^{\text{fuel}}-M_0^{\text{fuel}}|}{M_0^{\text{fuel}}} + \frac{|F^{\max}-F_0^{\max}|}{F_0^{\max}} + \frac{|\alpha-\alpha_0|}{\alpha_0} & m \geq 9, \\ +\infty & \text{otherwise,} \end{cases}$$

which constrains the mass to be greater than or equal to 9, and penalizes the percentage change in each of the parameters. The resulting feasible design has the parameters

$$m = 9.03, \quad M^{\text{fuel}} = 271.35, \quad F^{\max} = 67.16, \quad \alpha = 0.5,$$

and $r(\theta) = 0.948$.

4.2 Arbitrage

Consider an event (*e.g.*, horse race, sports game, or a financial market over a short time period) with m possible outcomes and n possible wagers on the outcome. The return matrix is $R \in \mathbf{R}^{m \times n}$, where R_{ij} is the return in dollars for the outcome i and wager j per dollar bet. A betting strategy is a vector $w \in \mathbf{R}_+^n$, where w_i is the amount that we bet on the i th wager. If we use a betting strategy w and outcome i occurs, then the return is $(Rw)_i$ dollars.

We say that there is an *arbitrage opportunity* in this event if there exists a betting strategy $w \in \mathbf{R}_+^n$ that is guaranteed to have nonnegative return for each outcome, and positive return in at least one outcome. We can check whether there exists an arbitrage opportunity by solving the convex optimization problem

$$\begin{aligned} & \text{maximize} && \mathbf{1}^T R w \\ & \text{subject to} && R w \geq 0, \\ & && w \geq 0, \end{aligned} \tag{13}$$

with variable w . If this problem is unbounded above, then there is an arbitrage opportunity.

Suppose that we are the event organizer (*e.g.*, sports book director, bookie, or financial exchange) and we wish to design the return matrix R such that there are arbitrage opportunities and that some performance metric r is small. We can tackle this problem by finding the nearest solvable convex optimization problem to problem (13) using algorithm 3.1.

Numerical example. We consider a horse race with $n = 3$ horses and $m = 5$ outcomes. The initial return matrix is

$$R_0 = \begin{bmatrix} 0.05 & 1.74 & -0.88 \\ 0.08 & 0.45 & -1.02 \\ 0.18 & -0.31 & 1.29 \\ 0.9 & -1.17 & 0.27 \\ -0.93 & 0.17 & 2.39 \end{bmatrix},$$

for which there is an arbitrage opportunity in the direction

$$w = (0.71, 0.62, 0.33).$$

We consider the regularization function $r(R) = \|(R - R_0)/R_0\|_1$, where $/$ is meant element-wise. After running algorithm 3.1, the arbitrage-free return matrix is

$$R_{\text{final}} = \begin{bmatrix} 0.05 & 1.71 & -0.9 \\ 0.08 & 0.42 & -1.09 \\ 0.18 & -0.31 & 1.27 \\ 0.81 & -1.22 & 0.27 \\ -0.97 & 0.17 & 2.37 \end{bmatrix}, \quad R_{\text{final}} - R_0 = \begin{bmatrix} 0 & -0.03 & -0.02 \\ 0 & -0.04 & -0.08 \\ 0 & 0 & -0.02 \\ -0.09 & -0.05 & 0 \\ -0.04 & 0 & -0.03 \end{bmatrix},$$

and $r(R_{\text{final}}) = 0.142$.

5 Related work

Irreducible infeasible subsets. An irreducible infeasible subset (IIS) of an optimization problem is a subset of the constraints that are by themselves infeasible, but any proper subset of them is feasible. Carver was the first to mention the concept of an IIS in 1921 [9]. Van Loon [38] and Greenberg [18] were among the first to recognize IIS in linear programs (LPs), but did not propose practical methods for finding them; Chinneck and Dravnieks [14] were the first to provide practical methods for finding IIS in LPs. Once IISs have been identified, various methods can be used to remove constraints until feasibility is obtained. Some notable examples of methods to prune constraints include the deletion filter [14], the additive method [36, 37], the elastic filter [14], the reciprocal filter [11], and the sensitivity filter [14]. For an excellent survey of IISs and how to use them to repair infeasible optimization problems, as well as applications, see [13, §6] and the many references therein. IISs have also been applied to convex quadratic inequalities in [28], to convex analytic inequalities in [29], and to semidefinite systems in [23].

Maximum feasible subsystem. The maximum feasible subsystem (MAX FS) of a set of constraints is the largest subset of the constraints that is feasible, introduced by Amaldi in 1999 [4]. MAX FS has some equivalent formulations, including the minimum unsatisfied linear relation problem [3], and the minimum cardinality IIS set-covering problem [10]. (All of these problems are NP-hard; see, *e.g.*, [35].) The problem of finding the MAX FS for linear constraints can be formulated as a mixed-integer linear program [19]. Various heuristics also exist for this problem, *e.g.*, based on IISs [10] and branch-and-cut [33]. For a survey of MAX FS problems and solution methods, see, *e.g.*, [13, §7] and [32, §1].

Optimal repair. IIS and MAX FS, the two aforementioned methods, both try to make the optimization problem feasible by removing constraints. In many cases, however, the goal is not to remove constraints, but to adjust parameters in the problem so that it becomes solvable [13, §8]. Also, removing or adjusting the constraints to obtain feasibility might inadvertently make the problem unbounded, so we have to assure that the problem becomes both feasible and bounded.

Problem (2) is often tractable when, in its conic representation (6), A is constant, and b and c are affine functions of θ . (This is not the case in any of our examples.) In this case, problem (2) can be expressed as a single convex problem, as noted by Roodman in 1979 [34]. In the case where the cones are products of the nonnegative reals, the resulting problem is immediately convex, while the more general case requires some care (see appendix A). Similar ideas have been applied to the more difficult problem of adjusting the coefficient matrix in LPs [5], and to repairing polynomial optimization problems [16] (for more examples see [13, §8] and the references therein). To the best of our knowledge, our paper is the first to consider automatic repair of convex optimization problems, allowing any parameter to be adjusted and guaranteeing both feasibility and boundedness.

Software. Various software packages have been developed for IIS, MAX FS, and optimal repair, and are part of most numerical solvers. One of the earliest is ANALYZE by Greenberg [17], which is an interactive system for analyzing and modifying linear programming models; more recent software packages include PERUSE [24] and MProbe [12]. Implementations of simple optimal repair methods for the righthand side of constraints in LPs can be found in the solvers MOSEK [26, §14.2], Gurobi [20], and CPLEX [21, §34].

Acknowledgments

S. Barratt is supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1656518.

References

- [1] A. Agrawal, B. Amos, S. Barratt, S. Boyd, S. Diamond, and J. Z. Kolter. Differentiable convex optimization layers. In *Advances in Neural Information Processing Systems*, pages 9558–9570, 2019.
- [2] A. Agrawal, S. Barratt, S. Boyd, E. Busseti, and W. Moursi. Differentiating through a cone program. *Journal of Applied and Numerical Optimization*, 1(2):107–115, 2019.
- [3] E. Amaldi. *From finding maximum feasible subsystems of linear systems to feedforward neural network design*. PhD thesis, Citeseer, 1994.
- [4] E. Amaldi, M. Pfetsch, and L. Trotter. Some structural and algorithmic properties of the maximum feasible subsystem problem. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 45–59. Springer, 1999.
- [5] P. Amaral, J. Júdice, and H. Serali. A reformulation–linearization–convexification algorithm for optimal correction of an inconsistent system of linear constraints. *Computers & Operations Research*, 35(5):1494–1509, 2008.
- [6] S. Barratt and S. Boyd. Least squares auto-tuning. *arXiv preprint arXiv:1904.05460*, 2019.
- [7] A. Ben-Tal and A. Nemirovski. *Lectures on modern convex optimization: analysis, algorithms, and engineering applications*, volume 2. SIAM, 2001.
- [8] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.
- [9] W. Carver. Systems of linear inequalities. *Annals of Mathematics*, pages 212–220, 1922.
- [10] J. Chinneck. An effective polynomial-time heuristic for the minimum-cardinality IIS set-covering problem. *Annals of Mathematics and Artificial Intelligence*, 17(1):127–144, 1996.

- [11] J. Chinneck. Finding a useful subset of constraints for analysis in an infeasible linear program. *INFORMS Journal on Computing*, 9(2):164–174, 1997.
- [12] J. Chinneck. Analyzing mathematical programs using MProbe. *Annals of Operations Research*, 104(1-4):33–48, 2001.
- [13] J. Chinneck. *Feasibility and Infeasibility in Optimization:: Algorithms and Computational Methods*, volume 118. Springer Science & Business Media, 2007.
- [14] J. Chinneck and E. Dravnieks. Locating minimal infeasible constraint sets in linear programs. *ORSA Journal on Computing*, 3(2):157–168, 1991.
- [15] S. Diamond and S. Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.
- [16] C. Gambella, J. Marecek, and M. Mevissen. Projections onto the set of feasible inputs and the set of feasible solutions. In *Allerton Conference on Communication, Control, and Computing*, pages 937–943. IEEE, 2019.
- [17] H. Greenberg. ANALYZE: A computer-assisted analysis system for linear programming models. *Operations Research Letters*, 6(5):249–255, 1987.
- [18] H. Greenberg. Computer-assisted analysis for diagnosing infeasible or unbounded linear programs. In *Computation Mathematical Programming*, pages 79–97. Springer, 1987.
- [19] H. Greenberg and F. Murphy. Approaches to diagnosing infeasible linear programs. *ORSA Journal on Computing*, 3(3):253–261, 1991.
- [20] GUROBI Optimization. Gurobi optimizer reference manual. 2019.
- [21] IBM. IBM ILOG CPLEX optimization studio CPLEX users manual, 2016.
- [22] R. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103.
- [23] K. Kellner, M. Pfetsch, and T. Theobald. Irreducible infeasible subsystems of semidefinite systems. *Journal of Optimization Theory and Applications*, 181(3):727–742, 2019.
- [24] W. Kurator and R. O’Neill. PERUSE: An interactive system for mathematical programs. *ACM Transactions on Mathematical Software*, 6(4):489–509, 1980.
- [25] B. Martinet. Brève communication. régularisation d’inéquations variationnelles par approximations successives. *ESAIM: Mathematical Modelling and Numerical Analysis-Modélisation Mathématique et Analyse Numérique*, 4(R3):154–158, 1970.
- [26] MOSEK Aps. MOSEK optimizer API for Python. <https://docs.mosek.com>, January 2020.

- [27] Y. Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1):125–161, 2013.
- [28] W. Obuchowska. Infeasibility analysis for systems of quadratic convex inequalities. *European Journal of Operational Research*, 107(3):633–643, 1998.
- [29] W. Obuchowska. On infeasibility of systems of convex analytic inequalities. *Journal of Mathematical Analysis and Applications*, 234(1):223–245, 1999.
- [30] B. O’Donoghue, E. Chu, N. Parikh, and S. Boyd. Conic optimization via operator splitting and homogeneous self-dual embedding. *Journal of Optimization Theory and Applications*, 169(3):1042–1068, 2016.
- [31] N. Parikh and S. Boyd. Proximal algorithms. *Foundations and Trends® in Optimization*, 1(3):127–239, 2014.
- [32] M. Pfetsch. *The maximum feasible subsystem problem and vertex-facet incidences of polyhedra*. PhD thesis, 2003.
- [33] M. Pfetsch. Branch-and-cut for the maximum feasible subsystem problem. *SIAM Journal on Optimization*, 19(1):21–38, 2008.
- [34] G. Roodman. Note—post-infeasibility analysis in linear programming. *Management Science*, 25(9):916–922, 1979.
- [35] J. Sankaran. A note on resolving infeasibility in linear programs by constraint relaxation. *Operations Research Letters*, 13(1):19–20, 1993.
- [36] M. Tamiz, S. Mardle, and D. Jones. Resolving inconsistency in infeasible linear programmes. Technical report, Technical Report, School of Mathematical Studies, University of Portsmouth, UK, 1995.
- [37] M. Tamiz, S. Mardle, and D. Jones. Detecting IIS in infeasible linear programmes using techniques from goal programming. *Computers & Operations Research*, 23(2):113–119, 1996.
- [38] J. N. M. Van Loon. Irreducibly inconsistent systems of linear inequalities. *European Journal of Operational Research*, 8(3):283–288, 1981.

A Convex formulation

In the case that A is a constant while b and c are affine functions of θ , we can write (9) as an equivalent convex optimization problem. In the linear case (*i.e.*, when $\mathcal{K} = \mathbf{R}_+^n$), we can simply drop the strong duality requirement (which always holds in this case) and express (9) as

$$\begin{aligned} & \text{minimize} && r(\theta) \\ & \text{subject to} && Ax + s = b(\theta) \\ & && A^T y + c(\theta) = 0 \\ & && s \in \mathcal{K}, \quad y \in \mathcal{K}^*. \end{aligned}$$

For more general cones \mathcal{K} (such as, *e.g.*, the second order cone), a sufficient condition for strong duality is that there exist a feasible point in the interior of the cone. We can write this as, for example,

$$\begin{aligned} & \text{minimize} && r(\theta) \\ & \text{subject to} && Ax + s = b(\theta) \\ & && A^T y + c(\theta) = 0 \\ & && s \in \mathbf{int} \mathcal{K}, \quad y \in \mathcal{K}^*. \end{aligned} \tag{14}$$

(We could similarly constrain $y \in \mathbf{int} \mathcal{K}^*$ and $s \in \mathcal{K}$.)

In general, optimizing over open constraint sets is challenging and these problems may not even have an optimal point, but, in practice (and for well-enough behaved r , *e.g.*, r continuous) we can approximate the true optimal value of (6) by approximating the open set $\mathbf{int} \mathcal{K}$ as a sequence of closed sets $\mathcal{K}_\varepsilon \subseteq \mathbf{int} \mathcal{K}$ such that $\mathcal{K}_\varepsilon \rightarrow \mathbf{int} \mathcal{K}$ as $\varepsilon \downarrow 0$.