



A rewriting system for convex optimization problems

Akshay Agrawal^a, Robin Verschueren^c, Steven Diamond^a and Stephen Boyd^b

^aDepartment of Computer Science, Stanford University, Stanford, CA, USA; ^bDepartment of Electrical Engineering, Stanford University, Stanford, CA, USA; ^cInstitut für Mikrosystemtechnik, Albert-Ludwigs-Universität Freiburg, Germany

ABSTRACT

We describe a modular rewriting system for translating optimization problems written in a domain-specific language (DSL) to forms compatible with low-level solver interfaces. Translation is facilitated by reductions, which accept a category of problems and transform instances of that category to equivalent instances of another category. Our system proceeds in two key phases: analysis, in which we attempt to find a suitable solver for a supplied problem, and canonicalization, in which we rewrite the problem in the selected solver's standard form. We implement the described system in version 1.0 of CVXPY, a DSL for mathematical and especially convex optimization. By treating reductions as first-class objects, our method makes it easy to match problems to solvers well-suited for them and to support solvers with a wide variety of standard forms.

ARTICLE HISTORY

Received 9 September 2017
Accepted 23 October 2017

KEYWORDS

Convex optimization;
domain-specific languages;
rewriting systems;
reductions

1. Introduction

Mathematical optimization centres on the optimization problem. Every optimization problem has three attributes: a variable whose value is to be assigned, constraints that the variable must satisfy and a real-valued objective function that measures the displeasure or cost incurred by any particular assignment to the variable. To solve an optimization problem is to find a numerical assignment to the variable that minimizes the objective function among all choices that satisfy the constraints.

Unfortunately, most optimization problems cannot be solved efficiently (Boyd & Vandenberghe, 2004, §1.4). There are, however, classes of optimization problems that can be solved in polynomial time. An important such class is the one composed of *convex* optimization problems – problems where the objective function is convex and where the constraints are described by a set of equality constraints with affine functions and inequality constraints with convex functions (Boyd & Vandenberghe, 2004; Nesterov & Nemirovski, 1994).

Modern convex optimization has its origin in linear programming, which traces back to the late 1940s, after the Second World War (Dantzig, 1963, §2). Since then, convex optimization has been extended to include a much wider variety of problems, and has found application in machine learning (Hastie, Tibshirani, & Friedman, 2009), control

(Boyd, El Ghaoui, Feron, & Balakrishnan, 1994) and computer science (Bertsekas, 1991; Goemans & Williamson, 1995, Parrilo, 2003), to name just a few of the fields touched by it. To accommodate the applications of convex optimization, researchers and practitioners have over the years authored many software packages. We distinguish between two types of software packages: domain-specific languages (DSLs), which streamline the process of specifying optimization problems, and low-level numerical solvers, which furnish solutions to problem instances.

1.1. Domain-specific languages

A DSL is a language that is designed for a particular application domain (Mernik, Heering, & Sloane, 2005); familiar examples include MATLAB and SQL. DSLs for convex optimization are languages designed for specifying convex optimization problems in natural, human-readable forms, and they obtain solutions to problems on their users' behalf by invoking numerical solvers; popular ones include Yalmip (Löfberg, 2004), CVX (Grant & Boyd, 2014), Convex.jl (Udell et al., 2014) and CVXPY (Diamond & Boyd, 2016a). These DSLs do support some non-convex regimes (e.g. combinatorial optimization), and there also exist DSLs for non-linear optimization (see Boyd & Vandenberghe, 2004, §1.4, for a definition), including GAMS (Brook, Kendrick, & Meeraus, 1988), AMPL (Fourer, Gay, & Kernighan, 1990) and JuMP (Dunning, Huchette, & Lubin, 2017); here, however, we limit our discussion to convexity.

We present below a toy example of an optimization problem written in CVXPY (version 1.0), a Python-embedded DSL (see Hudak, 1996, for background on embedded DSLs). The choice of language here is not particularly important, as the code would look similar if translated to any of the other aforementioned DSLs.

```
from cvxpy import *

alice = Variable()
bob = Variable()

objective = Minimize(maximum(alice + bob + 2, -alice - bob))
constraints = [alice <= 0, bob == -0.5]
toy = Problem(objective, constraints)
opt = toy.solve()
```

The CVXPY problem `toy` has two scalar optimization variables, `alice` and `bob`. Every `Variable` object has stored in its `value` field a numeric value, which is unspecified upon creation; `alice` and `bob` can hold floating point values. The objective is to minimize a piecewise-affine function of `alice` and `bob`, where the function is represented with the `max atom`. Atoms are mathematical functions like `square` and `exp` that operate on CVXPY expressions. CVXPY implements as library functions dozens of atoms for users to use in constructing problems. The arguments to the `max atom` are `Expression` objects, which encode mathematical expressions. `Constraint` objects are created by linking two expressions with a relational operator (`<=`, `>=`, or `==`). In the second-to-last line, the CVXPY problem `toy` is constructed, but not solved. Finally, an invocation to `toy`'s `solve` method solves the problem. A side-effect is that the `value` fields of the

optimization variables present in the problem (`alice` and `bob`) are assigned values that minimize the objective while satisfying the constraints, and the return value of such a `solve` is the value of the objective function evaluated at the variable values. After invoking `solve` above, we find that `alice.value == -0.5`, `bob.value == -0.5` and `opt == 1.0`. These values satisfy the two constraints, and among all such assignments, yield the smallest value of the objective function.

The `solve` method cannot solve all problems: Whether or not a problem can be solved depends on the objective, constraints and variables present in the problem. In particular, these entities must be constructed in such a way that CVXPY can detect that their assemblage is in fact a convex problem. Recognising convexity in general can be difficult. Many tricks aid in this task, but their application is sometimes guided by an intuition that is difficult to codify. Nonetheless, there do exist sufficient, but not necessary, rulesets for algorithmically detecting convexity. The ruleset employed by CVXPY, CVX, Convex.jl, and Yalmip is called *disciplined convex programming* (DCP) (Grant, 2004). These DSLs require users to express their problems by means of DCP, and while not all convex problems adhere to DCP, non-DCP convex problems can in practice be made DCP-compliant with at most a moderate amount of human effort and expertise. DCP-compliant problems can be verified as convex using a simple set of rules that are readily automated.

Our CVXPY example does not procedurally describe the method of procuring a solution. optimization DSLs are in this sense *declarative languages*. An optimization problem is simply a precise articulation of preferences and constraints among the values of the variables; how the problem is to be solved is another story. This story is the topic of our next section. But an abridged version is as follows: A subroutine inspects the problem and invokes a numerical solver capable of solving it, and most DSLs also allow users to mandate that a specific solver be used to solve any given problem – in CVXPY, for example, users may select a solver via the `solve` method’s keyword argument `solver`.

1.2. Numerical solvers

A numerical solver is a low-level tool that takes as input an optimization problem encoded in a rigid format and returns a solution for it. Every solver is tied to one or more *classes* of problems, insofar as problems supplied to a solver must be instances of one of its supported classes. One of the oldest problem classes is least squares, dating back to works authored by Legendre and Gauss in the late eighteenth and early nineteenth centuries (see Stigler, 1981, for a discussion of Legendre and Gauss’ contributions to the methodology, and Gauss, 1995, for a translation of Gauss’ manuscript on the same). Other well-studied convex optimization classes include linear programs, popularised by Dantzig following the Second World War (Dantzig, 1963) and cone programs, introduced in the late twentieth century by Nesterov and Nemirovski (1992, 1994, §4.1).

Various classes of convex problems fit into a hierarchy, as depicted in Figure 1. Every linear program reduces to a quadratic program (Boyd & Vandenberghe, 2004, §4.4); every quadratic program reduces to a second-order cone program (Nesterov & Nemirovski, 1994, §6.2.3); every second-order cone program reduces to a semidefinite program (Vandenberghe & Boyd, 1996); every semidefinite program reduces to a cone program; and every cone program reduces to a graph form program (Parikh & Boyd, 2014). Problem




Figure 1. Hierarchy of convex optimization problems. Rewriting systems should reduce instances of classes higher in the hierarchy to lower ones, when possible. (LP: linear program, QP: quadratic program, SOCP: second-order cone program, SDP: semidefinite program, CP: cone program, GFP: graph form program.)




Figure 2. Generic representation of a rewriting, capped with an invocation to a solver. The user-posed problem p_0 is *canonicalized* to a form p_n compatible with the targeted solver via a sequence of n reductions, yielding the intermediate problems p_1 through p_{n-1} . The solver solves p_n and returns a solution s_n for p_n ; we then *retrieve* solutions for the intermediate problems in reverse order, terminating with a solution s_0 for p_0 .

classes increase in both generality and difficulty as one goes up the hierarchy. For example, whereas every second-order cone program is also a semidefinite program, it is not advisable to solve the second-order cone programs using semidefinite program solvers (Alizadeh & Goldfarb, 2001). This fact motivates the existence of numerical solvers for different classes of problems – all other things equal, it is better to use as specific a solver as possible for the problem at hand.

In 1953, Hays and Dantzig developed for the RAND Corporation one of the earliest linear program solvers (Dantzig, 1963, §2.1). The universe of solvers has since grown in lockstep with the universe of optimization algorithms. There exist today solvers for many types of convex programs. Examples of solvers specialised to each class include GLPK (Makhorin, 2016) and CBC (Forrest & Lougee-Heimer, 2005) for linear programs; qpOASES (Ferreau, Kirches, Potschka, Bock, & Diehl, 2014) and OSQP (Banjac, Goulart, Stellato, & Boyd, 2017) for quadratic programs; ECOS (Domahidi, Chu, & Boyd, 2013) and Gurobi (Gurobi optimizer reference manual, 2017) for second-order cone programs; MOSEK (MOSEK optimization suite, 2017) and SeDuMi (Sturm, 1999) for semidefinite programs; SCS (O’Donoghue, Chu, Parikh, & Boyd, 2016) for more general cone programs;

and POGS for graph form programs (Fougner & Boyd, 2015). Many of the solvers listed also support non-convex problems such as mixed-integer and non-linear programs.

When we say that numerical solvers are low-level tools, we mean that it is onerous to translate problems to forms acceptable to solvers – even deciding which solver to use for a particular problem is a skill that requires training. These observations provide two of the *raison d'être* for optimization DSLs.

1.3. Canonicalization

The process of converting an optimization problem encoded in a DSL to a solver-compatible form – for example, the process by which CVXPY transformed our toy problem to the above representation – is called *canonicalization* (Grant, 2004, §4). Consider once again the toy problem from Section 1.1, the one transcribed in CVXPY. This problem can be transformed to an equivalent linear program, i.e. a problem of the form:

$$\begin{aligned} & \text{minimize} && c^T x \\ & \text{subject to} && Gx \leq h \\ & && Ax = b, \end{aligned}$$

where x is the (vector) variable, the matrices G and A , and vectors c , h and b , are constants, and the inequality is component-wise. Upon invoking the solve method, CVXPY canonicalizes the toy problem to the above standard form to make it compatible with a numerical (linear program) solver. The transformed problem has a variable $x \in \mathbb{R}^3$ whose first and second components represent `alice` and `bob`, respectively, and whose third component is an auxiliary variable introduced in the canonicalization. The problem data are:

$$G = \begin{bmatrix} 1 & 1 & -1 \\ -1 & -1 & -1 \\ 1 & 0 & 0 \end{bmatrix}, \quad A = [0 \ 1 \ 0], \quad c = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \quad h = \begin{bmatrix} -2 \\ 0 \\ 0 \end{bmatrix}, \quad b = -0.5.$$

This canonical form of our toy problem is the result of applying the transformation described in §6.5 of Dantzig and Thapa's (1997) text, and as such it could have easily been produced manually. Yet as problems of interest grow larger and more complex, producing canonical forms by hand quickly becomes a tedious, laborious and error-prone task. Many users of convex optimization still canonicalize problems manually – instead of letting optimization DSLs do the work for them – due to reliance on legacy systems, performance concerns, or in some cases, ignorance of the existence of DSLs that can automate the task.

Problems written in CVX, Convex.jl and CVXPY are automatically canonicalized to conic form, and Yalmip supports other classes of problems as well. For all of these DSLs, however, the solver selection and canonicalization procedures are implemented in ad hoc fashions that cannot easily be modified or extended, for instance to target new problem classes.

1.4. This paper

This paper is about treating canonicalization as a first-class component of the software ecosystem for convex optimization. We propose a *rewriting system* that sits between DSLs and numerical solvers, translating problems expressed in the former to forms compatible with the latter.

The atomic rewriting unit in our rewriting system is the *reduction*. A reduction is a function that converts problems of one form to equivalent problems of another form. Two problems are *equivalent* if a solution for one can be readily converted to a solution for the other (as defined in [Boyd & Vandenberghe, 2004](#), §4.1.3); readers from computer science will recognise that our definitions are in the same spirit as the more formal ones from their field (see e.g. §10.3 of [Papadimitriou, 1994](#), and §2.2 of [Arora & Barak, 2009](#)).

Every canonicalization is a reduction. In many cases, canonicalizations are complex enough to merit decomposing them into compositions of reductions. For a rewriting system to be useful, it must *retrieve* solutions for canonicalized problems to solutions for their provenances; because reductions output equivalent problems, they by definition support retrieval. The dual processes of canonicalization and retrieval are diagrammed in [Figure 2](#). Any rewriting that is cast as a composition of reductions is provably correct: The rewriting will either (a) output an equivalent problem, if the reductions in the canonicalization are mutually compatible and the first reduction is applicable to the source problem, or otherwise (b) audibly fail. There are at least three other benefits enjoyed by placing special emphasis on problem rewritings and reductions: doing so provides a structured way of preferentially targeting some solvers over others, simplifies the interfacing of DSLs with new solvers, and unifies problem transformations and back-end optimizations like presolves within a single conceptual framework.

The remainder of this paper is structured as follows. In [Section 2](#), we elaborate upon the role and structure of rewriting systems; in [Section 3](#), we list many examples of reductions, from simple to nuanced; and in [Section 4](#), we discuss version 1.0 of CVXPY, an open-source implementation of our proposed rewriting system.

2. An architecture for rewriting systems

2.1. Principles

There are four principles to which all optimization rewriting systems should adhere; these principles are informed by ones adopted by the software compiler community ([Aho, Lam, Sethi, & Ullman, 2006](#), §1.4.2).

- (1) Every rewriting must yield an equivalent problem that is target-compatible.
- (2) For each problem, an effort should be made to select a suitable solver for it.
- (3) The rewriting time must be tolerable.
- (4) The engineering effort required to maintain the rewriting system and add solvers to it must be kept manageable.

Adherence to the first three principles is necessary in order for a rewriting system to be useful, and to faithfully solve the specified problem. The fourth principle expedites not only the work of engineers responsible for rewriting systems but also that of those developing new solvers: Researchers who interface their solvers to popular DSLs gain immediate access

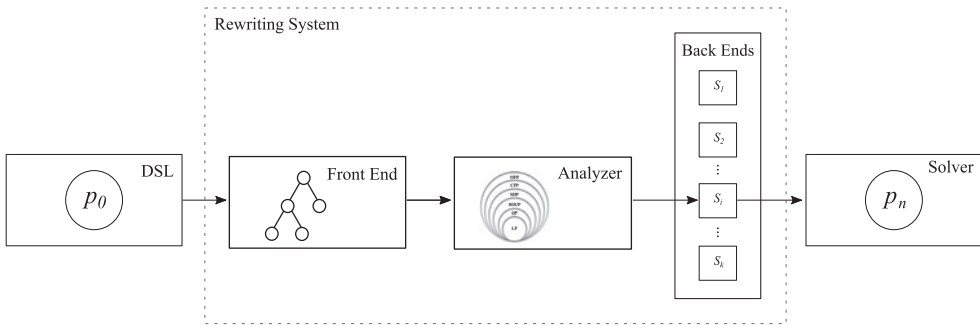


Figure 3. Architecture of a three-phase rewriting system. The system takes as input a problem p_0 written in a DSL. A *front end* specific to that DSL takes p_0 and encodes it in some intermediate representation, e.g., mathematical expression trees. The *analyser* reads the intermediate representation, discovers what it can about the problem’s structure and then selects a target, or solver; the analyser may itself produce further intermediate representations of the problem via reductions. A *back end* takes the intermediate representation produced by the analyser and reduces it to a form p_n that is compatible with the target.

to a rich ecosystem of problems for testing and tuning their algorithms. Indeed, the advent of software compilers fundamentally altered the development cycle of processors – today, compilers are built before processor designs are finalised, for often a processor is only useful if a compiler can exploit it (Aho et al., 2006, §1.5.3). The creation of optimization rewriting systems that satisfy the fourth principle might effect a similar paradigm shift in the design of numerical solvers.

As for how to actually satisfy these principles – the first principle is automatically satisfied by the use of reductions. One way to satisfy the remaining three principles is to decompose the rewriting system architecture into three distinct phases.

2.2. Phases

Just as it is natural to decompose software compilers into three phases (Aho et al., 2006, §1.2), so it is with optimization rewriting systems (Figure 3). The first and third phases of both systems are analogous. In the first phase, a *front end* takes a human-readable specification of a program and converts it to an intermediate representation; both compilers and rewriting systems often use abstract syntax trees as their intermediate representations (Aho et al., 2006, §2.5.1; Diamond & Boyd, 2016b; Udell et al., 2014). In the third phase, a *back end* takes an intermediate representation and translates it to a target-compatible form; targets for software compilers are processor architectures, while targets for optimization rewriting systems are numerical solvers. Our front end does not make use of reductions, as the problem of constructing a syntax tree is one of parsing, but our back ends most certainly do – in fact, our back ends are nothing more than sequences of reductions. In other words, each of the back ends supported by the rewriting system is a different canonicalization procedure.

Whereas the second phase of a compiler optimizes programs (Aho et al., 2006, §1.2.5), the second phase of our rewriting system *analyses* problems, selecting for each a suitable target; an imperfect analogy might compare our analysis phase to a static analysis tool like Frama-C (Kirchner, Kosmatov, Prevosto, Signoles, & Yakobowski, 2015). Compilers like the GNU Compiler Collection (Stallman & GCC Developer Community, 2017) can

translate a given source program to any of their targeted architectures, since every architecture is Turing complete (Sipser, 1996). Rewriting systems provide no such guarantees – solvers are specific to particular classes of problem, i.e. a given optimization problem might be compatible with but a subset of the solvers targeted by a rewriting system, hence the need for analysis. A sensible analysis policy is to identify the most specific class to which a problem belongs and choose a target supporting that class (see Figure 1). Solver selection might also reflect desiderata like accuracy, scalability and speed. For example, first-order solvers like SCS scale to larger problems than interior-point methods like ECOS, though the latter typically provide more accurate solutions than do the former. Analysis may of course fail to find a target that can handle a particular problem. When this happens, the rewriting system should abort with a descriptive error code or message.

This three-phase architecture satisfies principles two through four listed in Section 2.1. The existence of an analysis phase satisfies the second principle, a best-effort implementation of the analysis phase will satisfy the third principle, and the separation of back ends from front ends, together with the use of reductions as modular rewriting units, satisfies the fourth principle.

3. Reductions

In this section, we list several examples of reductions. Some of the listed reductions might be used to perform routine operations common among canonicalizations, while other more involved ones might be used to reduce problems to instances of simpler classes.

Notation. In all of the following examples, the variable x denotes the optimization variable wherever it appears, which may be a scalar or a vector. If we refer to the variables, plural, of an optimization problem, we mean to refer to the individual components of x . An *equality constraint* is one of the form $f(x) = g(x)$. An *inequality constraint* is one of the form $f(x) \leq g(x)$, where f and g are (real) vector-valued functions and the inequality is component-wise; the *constraint function* of a constraint is $f - g$. Equality constraints are convex if the constraint function is affine, and inequality constraints are convex if the function on the left (f) is convex and the function on the right (g) is concave.

3.1. Simple reductions

In this section, we present some simple but useful reductions.

Flipping objectives. The action of inverting a maximization objective to a minimization is a reduction: maximizing a function f over its domain is equivalent to minimizing $-f$, as solutions to the two problems are related by an inversion of sign. No work is required to retrieve a solution from the reduced-to problem, as the optimal sets of both problems are the same. (The optimal value of the problem must be negated.)

Moving expressions to the left of a relation. Subtracting the right-hand side of each constraint of a problem from both sides yields an equivalent problem in which all constraints have zero-valued right-hand sides. As in the case of the previous reduction, solution retrieval requires no work.

Eliminating general linear inequality constraints. Linear inequality constraints (i.e. those with f and g both affine) can be replaced with equality constraints by introducing non-negative variables. To wit, the constraint $f(x) \leq g(x)$ holds if and only if $f(x) + s = g(x)$, for some $s \geq 0$; the auxiliary variable s is called a *slack variable* (see [Dantzig, 1963](#), §4.3). Applying this reduction gives an equivalent problem in which every linear constraint is either an equality constraint or a non-negativity constraint on a slack variable. A solution for the original problem may be retrieved from a solution for the transformed problem by simply fetching the value of x and discarding the slack variables.

Monotone transformations of objective and constraints. Composing any monotone increasing function with the objective function of a problem yields an equivalent problem; so does transforming any number of constraints by applying any monotone increasing function to both sides. The retrieval method for this reduction is essentially a no-op, as the feasible and optimal sets for the two problems are identical. This reduction has been employed for centuries – squaring the Euclidean norm when it appears as an objective function to render it differentiable is, at least historically, standard mathematical practice.

Stated in the opposite direction, if the objective function of a convex problem is monotone increasing, then the problem given by replacing the objective function with a composition of its inverse and itself (and similarly for the constraints) is an equivalent convex problem. This reduction might be used to eliminate exponentials and logarithms from a problem, which in turn might make the problem amenable to more solvers.

Changing variables. Let ϕ be any one-to-one function whose image covers the problem domain. Replacing the optimization variable x wherever it appears with $\phi(z)$ is a reduction, yielding an equivalent problem with optimization variable z . To retrieve a solution for the original problem from one for the transformed problem, simply take $x^* = \phi(z^*)$, where x^* denotes an optimal point for the original problem and z^* an optimal point for the transformed problem.

Changing variables can convert non-convex problems to equivalent convex problems. A classic example is the technique for convexifying geometric programs; this technique both changes variables and transforms the objective and constraints. A geometric program is an optimization problem of the form:

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

where the functions f_i are *posynomials* and the functions h_i are *monomials*. A monomial is a function $f : \mathbf{R}^n \rightarrow \mathbf{R}$ over the Non-negative orthant defined as:

$$f(x) = cx_1^{a_1} x_2^{a_2} \cdots x_n^{a_n},$$

where $c > 0$ and $a_i \in \mathbf{R}$, and a polynomial is a sum of monomials. Performing a change of variables with $x_i = \exp(z_i)$ and taking logarithms of both the objective and constraints results in a convex problem. In fact, if the f_i are all monomials, then the resulting problem is a linear program. For a brief introduction to geometric programming, see ([Boyd & Vandenberghe, 2004](#), §4.5); for a longer survey, see ([Boyd, Kim, Vandenberghe, & Hassibi, 2007](#)).

Eliminating complex numbers. It is possible to reduce an optimization problem over a complex domain to one over a real-domain. (While such a problem has complex variables and expressions, the constraint and objective functions must all be real-valued.)

We provide here a partial specification of the reduction. Absolute values of complex numbers reduce to Euclidean norms of their concatenated real and imaginary parts, sums of complex numbers reduce to sums of their real and imaginary parts, and equality constraints between two complex numbers reduce to equality constraints on the implicated expressions' real and imaginary parts. Perhaps more interesting, positive semidefinite constraints on Hermitian matrices reduce to positive semidefinite constraints on real symmetric matrices. As presented in [Goemans and Williamson \(2004\)](#), a Hermitian matrix X is positive semidefinite if and only if the real symmetric matrix $T(X)$ is positive semidefinite, where the mapping T is defined as:

$$T(X) = \begin{bmatrix} \operatorname{Re} X & -\operatorname{Im} X \\ \operatorname{Im} X & \operatorname{Re} X \end{bmatrix}.$$

As this reduction expands the optimization variable into its real and imaginary parts, retrieving a solution for the complex-domain problem from a solution for the real-domain problem requires but a bit of book-keeping to map variable values from the latter solution to the real and imaginary parts of the original variable.

3.2. Presolves

A presolve is any reduction that is meant to decrease the computational cost incurred when solving a problem. Presolves are typically performed immediately before problems are solved, with some but not all numerical solvers subjecting problems to a battery of presolves prior to solving them. Many presolves are applicable across solvers, i.e. a presolve that helps one solver is likely to help many others. This motivates folding presolves into rewriting systems and excising them from numerical solvers whenever possible. As recommended by [Bradley, Brown, and Graves \(1983\)](#), one might even treat the application of presolves as a fixed point iteration, cyclically applying presolves until the problem cannot be further simplified. This approach resembles the multiple passes an optimizing software compiler may make over an intermediate code representation ([Aho et al., 2006](#), §8 and §9).

There is substantial literature on presolves. [Andersen and Andersen \(1995\)](#) cast presolves as reductions, listing many examples for linear programs together with methods for retrieving their solutions. Earlier surveys of linear programming presolves include [Brearley, Mitra, and Williams \(1975\)](#), [Bradley et al. \(1983\)](#), and [Tomlin and Welch \(1986\)](#). [Tomlin \(1975\)](#) discusses the problem of scaling data matrices to coax faster performance out of the simplex algorithm, while [Bradley \(2010\)](#) provides a more modern perspective on scaling for a wider class of algorithms. Here, we present a sample of some of the presolves covered by these and other references.

Eliminating fixed variables. Any variable that is constrained to be a constant is called a fixed variable; replacing every occurrence of it with the value of the constant yields an equivalent problem. In the software compilers literature, this technique is called *constant propagation* ([Aho et al., 2006](#), §9.4). Solution retrieval simply requires setting the values of the fixed variables to their respective constant values and copying all other variable values.

Eliminating free variables. Any variable that does not have upper and lower bounds is called a free variable; replacing every occurrence of it with the difference of two auxiliary non-negative variables yields an equivalent problem for which a solution can be retrieved in the obvious way (Dantzig, 1963, §4.3). In the setting of cone programs, free variables are those that are not restricted to lie in a cone (other than \mathbf{R}^n). There are a number of ways to exploit free variables in cone programs, some of which are outlined by Anjos and Burer (2008).

Eliminating redundant constraints. Any constraint whose removal leaves the feasible region unchanged is redundant; deleting such constraints yields an equivalent problem. For example, any equation in a linear system that is a linear combination of the others is redundant (Dantzig & Thapa, 1997, §B.2). As another example, if it is required that $x \leq b$ and $x \leq c$, and moreover if it is known that $b \leq c$, then the constraint $x \leq c$ is redundant. Solution retrieval for this reduction is a no-op.

Scaling. Scaling both sides of a constraint by a positive constant is a presolve; this presolve is an instance of monotonically transforming constraints (see Section 3.1). It has long been known that scaling matrices (i.e. scaling constraints of the form $Ax \leq b$ or $Ax = b$) can by lowering the condition numbers of these matrices dramatically affect the performance of first-order methods for convex optimization (see, e.g. Tomlin, 1975). One scaling technique, called diagonal preconditioning, premultiplies such constraints by diagonal matrices and also changes variables by premultiplying the optimization variable by another diagonal matrix (Kelley, 1995, §2.5; Pock & Chambolle, 2011; Takapoui & Javadi, 2016). A popular instantiation of this technique is matrix equilibration, which chooses the diagonal matrices so that all rows of the scaled data matrix have one p -norm and all columns have another, with the two equal for square matrices. The literature on equilibration spans decades (see e.g. Bradley, 2010; Diamond & Boyd, 2016c; Sluis, 1969, and the references therein).

3.3. Conic canonicalization of DCP programs

The embedded languages CVX, Convex.jl, CVXPY and YALMIP canonicalize problems to a form compatible with cone program solvers; in particular, the canonicalized objective function is affine and all constraints are conic, imposed only on affine expressions of the optimization variable. These tools canonicalize problems in the same fashion, and the methodology shared among them is the subject of this section. The methodology – which is a reduction if the problem operated upon is DCP-compliant – proceeds in three steps: first, the problem is lifted into a higher dimension via its *Smith form*, making affine the arguments of each atom; second, the lifted problem is *relaxed* to a convex problem; and third, every non-linear atom is replaced with conic constraints that encode its *graph implementation*. Our exposition in this section draws from work by Smith (1996), who introduced Smith form, Grant and Boyd (2008), who introduced graph implementations and Chu, Parikh, Domahidi, and Boyd (2013), who illustrated these three steps with a clear example.

Smith form. It is natural to view an optimization problem as composed of a forest of mathematical expression trees, with one tree for the objective function and two trees for every constraint, one for each side of the constraint. The inner nodes of an expression

tree represent mathematical functions, or atoms, and the leaves represent variables and constants. Every inner node is evaluated at its children, i.e. the children of an atom are its arguments. For example, the expression $f(x) + c$, in which f is an atom and c is a constant, parses into a tree where the summation atom is the root, f and c are the children of the root, and x is the child of f .

Converting an optimization problem to Smith form requires making a single pass over every expression tree present in the problem. Recursively, beginning with the root, each subexpression is replaced with an auxiliary variable, and equality constraints are introduced between the auxiliary variables and the subexpressions they replaced. The resulting problem is said to be in Smith form, a key property of which is that the arguments of each atom within the problem are affine (indeed, they are unadorned variables). Transforming a problem to Smith form is always a reduction. This reduction does not however preserve convexity, as any convex function present in the original problem will appear as the constraint function of an equality constraint in the transformed problem.

Relaxed Smith form. If the original problem is DCP-compliant, then its Smith form can be relaxed to an equivalent convex problem. In particular, relaxing in the correct direction the non-convex equality constraints between the auxiliary variables and their associated atoms is in this case a reduction, in the following sense: if (x^*, t^*) is optimal for the relaxed problem, t a vector of the auxiliary variables and x the original variable, then x^* is optimal for the original problem.

Graph implementations. The final step in this reduction replaces every constraint in which a Non-linear convex atom appears with conic constraints that encode the atom's epigraph. The *epigraph* of a function is defined as the set of points that lie above its graph: For a function f , its epigraph is defined as the set of points $\{(x, y) \mid f(x) \leq y\}$. As a simple example of such a replacement, the epigraph of the function $|x|$ is the set $\{(x, y) \mid x \leq y, -x \leq y\}$; accordingly, the constraint $|t_1| \leq t_2$ would be replaced with the constraints $t_1 \leq t_2$ and $-t_1 \leq t_2$. Constraints encoding the epigraph of an atom are called the *graph implementation* of the atom, coined in Grant (2004), though such constraints might more aptly be referred to as an epigraph implementation; the action of replacing a non-linear atom with its graph implementation is called a *graph expansion*.

Graph implementations are useful outside of conic canonicalization as well. For example, one might choose to only perform graph expansions for piecewise-linear atoms such as `abs`, `max` and `sum_k_largest`, which sums the k largest entries of a vector or matrix. This process of *eliminating piecewise-linear* atoms is itself a reduction if the problem to which it is applied is DCP-compliant.

3.4. Other reductions

The reductions presented are somewhat subtle and problem specific, examples of the kinds of experimental reductions one might include in a rewriting system.

Decomposing second-order cone constraints. A *second-order cone constraint* on a block vector (x, t) is a constraint of the form $\|x\|_2 \leq t$, where $x \in \mathbf{R}^n$ is a vector and t is a scalar. The dimension of such a second-order cone constraint is $n + 1$. Any second-order constraint of dimension $n + 1$, $n \geq 2$, can be reduced to $n - 1$ three-dimensional second-order cone constraints by the following chain of observations: The constraint $\|x\|_2 \leq t$

holds if and only if $x_1^2 + x_2^2 + \dots + x_n^2 \leq t^2$ and $t \geq 0$, which in turn holds if and only if $x_2^2 + \dots + x_n^2 \leq u^2$, $u^2 \leq t^2 - x_1^2$, and $u, t \geq 0$, where u is a scalar variable, i.e. if and only if (x_2, \dots, x_n, u) and (x_1, u, t) are in the second-order cone. The result follows by recursing on (x_2, \dots, x_n, u) .

Decomposing semidefinite constraints. A *semidefinite program* is a convex optimization problem of the form:

$$\begin{aligned} & \text{minimize} && \text{tr}(CX) \\ & \text{subject to} && \text{tr}(A_i X) = b_i, \quad i = 1, \dots, p \\ & && X \in \mathbf{S}_+^n, \end{aligned}$$

where the constraint $X \in \mathbf{S}_+^n$ requires X to be an $n \times n$ positive semidefinite matrix. Semidefinite programs can become significantly harder to solve as the size of the matrices involved increases, motivating reductions that decompose the matrices into smaller ones.

A semidefinite program that exhibits a chordal aggregate sparsity pattern can often be reduced to a program involving smaller matrices. The sparsity pattern of an $n \times n$ matrix is a set E of pairs (i, j) , $i, j \in \{1, 2, \dots, n\}$, such that $A_{ij} = 0$ if $(i, j) \notin E$ and $i \neq j$; the aggregate sparsity pattern of a semidefinite program is the union of the edge sets of the sparsity patterns of C and A_1, A_2, \dots, A_p . A sparsity pattern is *chordal* if the induced graph (with vertices $V = \{1, 2, \dots, n\}$ and edges E) is chordal, i.e. if every cycle of length greater than three contains an edge between two non-consecutive vertices, called a chord.

The reduction in question replaces the optimization variable with smaller matrix variables, one for each clique in the graph induced by the aggregate sparsity pattern, rewriting the objective and equality constraints in terms of these new variables and adding equality constraints to preserve the semantics of the original problem. The particulars of this reduction (and related ones) can be found in [Fukuda, Kojima, Murota, and Nakata \(2001\)](#), whose authors were among the first to exploit chordal sparsity patterns to decompose large semidefinite programs, and ([Vandenberghe & Andersen, 2015](#), §14.2), which thoroughly surveys the topic of chordal graphs as they relate to semidefinite optimization.

Relaxing convex equality constraints. Consider an optimization problem with a convex objective, convex inequality constraints and a single convex equality constraint. This problem is not convex; however, it can in certain cases be coerced into a convex form. Letting x be the problem variable, if there is an index r such that the objective is monotonically increasing in x_r , each inequality constraint function is non-decreasing in x_r , and the equality constraint function is monotonically decreasing in x_r , then relaxing the equality constraint to a non-positive inequality constraint produces an equivalent convex problem, i.e. the relaxation is tight ([Boyd & Vandenberghe, 2004](#), exercise 4.6).

3.5. An example

We address in this section the following quite practical question: What types of convex programs reduce to quadratic programs?

A *quadratic program* is an optimization problem in which the objective function is a convex quadratic and the constraint functions are affine ([Boyd & Vandenberghe, 2004](#), §4.4); quadratic programs have been studied since the 1950s. Evidently, a problem in which every inequality constraint function is piecewise-linear and every equality constraint

function is affine can be reduced to a problem in which every constraint function is affine (by eliminating the piecewise-linear atoms via graph expansions, as described in Section 3.3).

Describing acceptable objective functions requires more work; we will specify acceptable objective functions via their expression trees, and we will specify an acceptable expression tree by providing *regular expressions* (Aho et al., 2006, §3.3.3) for paths beginning at the root and terminating at (the ancestor of) a leaf. Letting A denote an affine atom, P a piecewise-linear atom and Q a quadratic atom, it is clear that any objective function whose root-to-leaf paths satisfy the regular expression $A^*QA^*|P^+$ can be canonicalized to a quadratic by eliminating the piecewise-linear atoms – this is evident because the Hessian of such a function is constant almost everywhere.

Barring non-linear transformations (e.g. squaring a norm), one might reasonably assume that the class of problems reducible to quadratic problem cannot be further generalised, for the regular expression does after all capture both linear programs and quadratic programs. But such an assumption would be incorrect. In fact, any DCP-compliant problem (with acceptable constraints) whose objective function's root-to-leaf paths satisfy the regular expression $A^*QP^*|P^+$ can be reduced to a quadratic program (Figure 4 renders the regular expression as a finite-state machine). The corresponding reduction eliminates the piecewise-linear atoms using graph expansions and then massages the objective into a quadratic.

As an example, the following problem can be canonicalized to a quadratic program:

$$\text{minimize } (\max(x, 0) + \max(x - 1, 0))^2.$$

Note that unlike the simpler class of problems we described, this problem does not have a constant Hessian – its second derivative is 0 when $x < 0$, 2 when $0 < x < 1$, and 8 when $x > 1$. It is nonetheless reducible to a quadratic program, i.e. a problem whose objective function has a constant Hessian.

4. Implementation

We have implemented a number of the reductions from Section 3 in version 1.0 of CVXPY, an open-source implementation of our proposed three-phase rewriting system, available at

<http://www.cvxpy.org/>.

All problem rewriting is facilitated by `Reduction` objects, and every reduction implements three methods: `accepts`, `apply` and `retrieve`. The `accepts` method takes as input a problem and returns a boolean indicating whether or not the reduction can be applied to the problem, the `apply` method takes as input a problem and returns a new equivalent problem, and the `retrieve` method takes a solution for the problem returned by an invocation of `apply` and retrieves from it a solution for its problem of provenance. Some of the tasks carried out by reductions in our system include flipping objectives (Section 3.1), eliminating piecewise-linear atoms (Section 3.3), canonicalizing problems to cone programs, (Section 3.3), and canonicalizing problems to quadratic programs (Section




Figure 4. A finite-state machine for the example in Section 3.5. Any DCP-compliant problem in which the root-to-leaf paths of the objective function’s expression tree are accepted by this state machine can be reduced to a quadratic program, provided that the equality constraint functions are affine and the inequality constraint functions piecewise-linear. Above, A represents affine atoms, Q represents quadratic atoms and P represents piecewise-linear atoms.

3.5). Sequences of reductions are represented by `Chain` objects, which are themselves reductions, and solver back ends are implemented with `Chain` objects.

Creating expressions and constraints in CVXPY invokes behind-the-scenes a front end that parses them into expression trees; this functionality is not new (see [Diamond & Boyd, 2016a](#)). What is new is the method by which solvers are chosen for problems and the methods by which problems are canonicalized to their standard forms. In CVXPY 1.0, invoking the `solve` method of a problem triggers an analyser, phase two of our rewriting system. The analyser determines the most specific class to which the problem belongs by checking which back ends accept the problem; back ends are checked in order of decreasing specificity, and the analysis is short-circuited as soon as a suitable back end is found. Analysis may itself apply simple reductions to the problem. Any such reductions are prepended to the back end to create a `Chain` object that encapsulates the entire rewriting process, and the problem is solved by applying the chained reduction, invoking a solver, and using the chained reduction to retrieve a solution. For example, if a user specifies a problem that is recognisably reducible to a quadratic program, and if a quadratic solver is installed on the user’s device, CVXPY 1.0 will automatically target it.

Acknowledgements

We extend many thanks to Enzo Busseti, for significant contributions to the quadratic program canonicalization procedure, and Yifan Lu, for fruitful conversations about the analogy between our rewriting system and software compilers.

Disclosure statement

No potential conflict of interest was reported by the authors.

Notes on contributors



Akshay Agrawal holds MS and BS degrees in computer science from Stanford University, and he is currently a member of the Google Brain team, where he works on TensorFlow. His research interests include mathematical optimization and its applications, machine learning, and distributed systems.



Robin Verschueren is currently a PhD student with Moritz Diehl at the University of Freiburg, Germany with a Marie Curie grant from the European Union's seventh framework program. He obtained a master in Mathematical Engineering from the KU Leuven, Belgium, in 2014. His research interests are embedded optimization and non-linear model predictive control.



Steven Diamond is a PhD candidate in Computer Science at Stanford University, advised by Stephen Boyd. He received his BS in Computer Science at Stanford University in 2014. His research interests include domain-specific languages for optimization, matrix-free optimization, and applications of optimization to finance and imaging.



Stephen Boyd is the Samsung professor of Engineering and professor of Electrical Engineering in the Information Systems Laboratory at Stanford University, with courtesy appointments in Computer Science and Management Science and Engineering. He received the AB degree in Mathematics from Harvard University in 1980, and the PhD in Electrical Engineering and Computer Science from the University of California, Berkeley, in 1985, and then joined the faculty at Stanford. His current research focus is on convex optimization applications in control, signal processing, machine learning and finance.

References

- Aho, A., Lam, M., Sethi, R., & Ullman, J. (2006). *Compilers: Principles, techniques, and tools* (2nd ed.). Boston, MA: Addison-Wesley Longman.
- Alizadeh, F., & Goldfarb, D. (2001). Second-order cone programming. *Mathematical Programming*, 95, 3–51.
- Andersen, E., & Andersen, K. (1995). Presolving in linear programming. *Mathematical Programming*, 71(2), 221–245.

- Anjos, M., & Burer, S. (2008). On handling free variables in interior-point methods for conic linear optimization. *SIAM Journal on Optimization*, 18(4), 1310–1325.
- Arora, S., & Barak, B. (2009). *Computational complexity: A modern approach* (1st ed.). New York, NY: Cambridge University Press.
- Banjac, G., Goulart, P., Stellato, B., & Boyd, S. (2017). Infeasibility detection in the alternating direction method of multipliers for convex optimization. *Optimization Online*. Retrieved from http://www.optimization-online.org/DB_HTML/2017/06/6058.html
- Bertsekas, D. (1991). *Linear network optimization: Algorithms and codes*. Cambridge, MA: MIT Press.
- Boyd, S., El Ghaoui, L., Feron, E., & Balakrishnan, V. (1994). *Linear matrix inequalities in system and control theory*. Philadelphia, PA: Society for Industrial and Applied Mathematics.
- Boyd, S., Kim, S.-J., Vandenberghe, L., & Hassibi, A. (2007). A tutorial on geometric programming. *Optimization and Engineering*, 8(1), 67–127.
- Boyd, S., & Vandenberghe, L. (2004). *Convex optimization*. New York, NY: Cambridge University Press.
- Bradley, A. M. (2010). *Algorithms for the equilibration of matrices and their application to limited-memory quasi-Newton methods* (Unpublished doctoral dissertation). Stanford, CA: Stanford University.
- Bradley, G., Brown, G., & Graves, G. (1983). Structural redundancy in large-scale optimization models. In M. Beckmann & W. Krelle (Eds.), *Redundancy in mathematical programming: A state-of-the-art survey* (pp. 145–169). Berlin: Springer.
- Brearley, A., Mitra, G., & Williams, H. (1975). Analysis of mathematical programming problems prior to applying the simplex algorithm. *Mathematical Programming*, 8(1), 54–83.
- Brook, A., Kendrick, D., & Meeraus, A. (1988). GAMS, a user's guide. *SIGNUM Newsletter*, 23(3–4), 10–11.
- Chu, E., Parikh, N., Domahidi, A., & Boyd, S. (2013). Code generation for embedded second-order cone programming. In *Proceedings of the European Control Conference* (pp. 1547–1552), Zurich.
- Dantzig, G. (1963). *Linear programming and extensions* (Technical Report R-366-PR). Santa Monica, CA: RAND Corporation.
- Dantzig, G., & Thapa, M. (1997). *Linear programming 1: Introduction*. Secaucus, NJ: Springer-Verlag New York Inc.
- Diamond, S., & Boyd, S. (2016a). CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83), 1–5.
- Diamond, S., & Boyd, S. (2016b). Matrix-free convex optimization modeling. In B. Goldengorin (Ed.), *Optimization and its applications in control and data sciences* (Vol. 115, pp. 221–264). Springer. Retrieved from <https://link.springer.com/book/10.1007/978-3-319-42056-1#toc>
- Diamond, S., & Boyd, S. (2016c). Stochastic matrix-free equilibration. *Journal of Optimization Theory and Applications*, 172(2), 436–454.
- Domahidi, A., Chu, E., & Boyd, S. (2013). ECOS: An SOCP solver for embedded systems. In *Proceedings of the European Control Conference* (pp. 3071–3076), Zurich.
- Dunning, I., Huchette, J., & Lubin, M. (2017). JuMP: A modeling language for mathematical optimization. *SIAM Review*, 59(2), 295–320.
- Ferreau, H., Kirches, C., Potschka, A., Bock, H., & Diehl, M. (2014). qpOASES: A parametric active-set algorithm for quadratic programming. *Mathematical Programming Computation*, 6(4), 327–363.
- Forrest, J., & Lougee-Heimer, R. (2005). CBC user guide. In *Emerging theory, methods, and applications* (pp. 257–277). Catonsville: INFORMS.
- Fougner, C., & Boyd, S. (2015). Parameter selection and pre-conditioning for a graph form solver. arXiv preprint arXiv:1503.08366.
- Fourer, R., Gay, D., & Kernighan, B. (1990). A modeling language for mathematical programming. *Management Science*, 36(5), 519–554.
- Fukuda, M., Kojima, M., Murota, K., & Nakata, K. (2001). Exploiting sparsity in semidefinite programming via matrix completion i: General framework. *SIAM Journal on Optimization*, 11(3), 647–674.

- Gauss, C. (1995). *Theory of the combination of observations least subject to errors*. Society for Industrial and Applied Mathematics. Translated from original 1820 manuscript by G. W. Stewart).
- Goemans, M., & Williamson, D. (1995). Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6), 1115–1145.
- Goemans, M., & Williamson, D. (2004). Approximation algorithms for MAX-3-CUT and other problems via complex semidefinite programming. *Journal of Computer System Sciences*, 68(2), 442–470.
- Grant, M. (2004). *Disciplined convex programming* (Unpublished doctoral dissertation). Stanford, CA: Stanford University.
- Grant, M., & Boyd, S. (2008). Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, & H. Kimura (Eds.), *Recent advances in learning and control* (pp. 95–110). London: Springer.
- Grant, M., & Boyd, S. (2014). *CVX: Matlab software for disciplined convex programming version 2.1*. Retrieved from <http://cvxr.com/cvx>
- Gurobi optimizer reference manual*. (2017). Retrieved from <https://gurobi.com/documentation/7.5/refman.pdf>
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction*. New York, NY: Springer.
- Hudak, P. (1996). Building domain-specific embedded languages. *ACM Computing Surveys*, 28(4es). Retrieved from <https://dl.acm.org/citation.cfm?id=242477&dl=ACM&coll=DL&CFID=991057259&CFTOKEN=42798159>
- Kelley, C. (1995). *Iterative methods for linear and nonlinear equations*. Philadelphia, PA: Society for Industrial and Applied Mathematics.
- Kirchner, F., Kosmatov, N., Prevosto, V., Signoles, J., & Yakobowski, B. (2015). Framac-c: A software analysis perspective. *Formal Aspects of Computing*, 27(3), 573–609.
- Löfberg, J. (2004). YALMIP: A toolbox for modeling and optimization in MATLAB. In *Proceedings of the CACSD Conference*. Taipei, Taiwan.
- Makhorin, A. (2016). *GNU Linear Programming Kit v4.60*. Retrieved from <http://www.gnu.org/software/glpk/>
- Mernik, M., Heering, J., & Sloane, A. M. (2005). When and how to develop domain-specific languages. *ACM Computing Surveys*, 37(4), 316–344.
- MOSEK optimization suite*. (2017). Retrieved from <http://docs.mosek.com/8.0/intro.pdf>
- Nesterov, Y., & Nemirovski, A. (1992). Conic formulation of a convex programming problem and duality. *Optimization Methods and Software*, 1(2), 95–115.
- Nesterov, Y., & Nemirovski, A. (1994). *Interior-point polynomial algorithms in convex programming*. Philadelphia, PA: Society for Industrial and Applied Mathematics.
- O'Donoghue, B., Chu, E., Parikh, N., & Boyd, S. (2016). Conic optimization via operator splitting and homogeneous self-dual embedding. *Journal of Optimization Theory and Applications*, 169(3), 1042–1068.
- Papadimitriou, C. (1994). *Computational complexity*. Reading, MA: Addison-Wesley.
- Parikh, N., & Boyd, S. (2014). Block splitting for distributed optimization. *Mathematical Programming Computation*, 6(1), 77–102.
- Parrilo, P. (2003). Semidefinite programming relaxations for semialgebraic problems. *Mathematical Programming*, 96(2), 293–320.
- Pock, T., & Chambolle, A. (2011). Diagonal preconditioning for first order primal-dual algorithms in convex optimization. In *Proceedings of the 2011 International Conference on Computer Vision* (pp. 1762–1769). Washington, DC: IEEE Computer Society.
- Sipser, M. (1996). *Introduction to the theory of computation* (1st ed.). Boston: International Thomson Publishing.
- Sluis, A. (1969). Condition numbers and equilibration of matrices. *Numerical Mathematics*, 14(1), 14–23.
- Smith, E. M. d. B. (1996). *On the optimal design of continuous processes* (Unpublished doctoral dissertation). London: Imperial College London (University of London).

- Stallman, R. M., & GCC Developer Community (2017). *Using the GNU compiler collection: A GNU manual for GCC Version 7.2.0*. Boston, MA: GNU Press.
- Stigler, S. (1981). Gauss and the invention of least squares. *The Annals of Statistics*, 9(3), 465–474.
- Sturm, J. (1999). Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optimization Methods and Software*, 11(1–4), 625–653.
- Takapoui, R., & Javadi, H. (2016). *Preconditioning via diagonal scaling*. arXiv preprint arXiv:1610.03871.
- Tomlin, J. (1975). On scaling linear programming problems. In M. Balinski & E. Hellerman (Eds.), *Computational practice in mathematical programming* (pp. 146–166). Berlin: Springer.
- Tomlin, L., & Welch, J. (1986). Finding duplicate rows in a linear programming model. *Operations Research Letters*, 5(1), 7–11.
- Udell, M., Mohan, K., Zeng, D., Hong, J., Diamond, S., & Boyd, S. (2014). Convex optimization in Julia. In *SC14 Workshop on High Performance Technical Computing in Dynamic Languages*.
- Vandenberghe, L., & Andersen, M. (2015). Chordal graphs and semidefinite optimization. *Foundations and Trends in Optimization*, 1(4), 241–433.
- Vandenberghe, L., & Boyd, S. (1996). Semidefinite programming. *SIAM Review*, 38(1), 49–95.