

Convex Optimization in Julia

Madeleine Udell
udell@stanford.edu

Karanveer Mohan
kvmohan@stanford.edu

David Zeng
dzeng0@stanford.edu

Jenny Hong
jyunhong@stanford.edu

Steven Diamond
stevend2@stanford.edu

Stephen Boyd
boyd@stanford.edu

ABSTRACT

This paper describes **Convex**¹, a convex optimization modeling framework in Julia. **Convex** translates problems from a user-friendly functional language into an abstract syntax tree describing the problem. This concise representation of the global structure of the problem allows **Convex** to infer whether the problem complies with the rules of disciplined convex programming (DCP), and to pass the problem to a suitable solver. These operations are carried out in Julia using multiple dispatch, which dramatically reduces the time required to verify DCP compliance and to parse a problem into conic form. **Convex** then automatically chooses an appropriate backend solver to solve the conic form problem.

Categories and Subject Descriptors

G.1.6 [Numerical Analysis]: Optimization—*Convex Programming*

; D.2.4 [Software Engineering]: Software/Program Verification—*Model checking*

General Terms

Languages

Keywords

Convex programming, automatic verification, symbolic computation, multiple dispatch

1. INTRODUCTION

The purpose of an optimization modeling language is to translate an optimization problem from a user-friendly language into a solver-friendly language. In this paper, we present an approach to bridging the gap. We show that mathematical optimization problems can be parsed from a simple human-readable form, using the trope of function

¹available online at <http://github.com/cvxgrp/Convex.jl>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

HPTCDL November 16-21, 2014, New Orleans, Louisiana, USA
Copyright 2014 ACM 978-1-4799-5500-8/14 ...\$15.00.

composition, into an abstract syntax tree (AST) representing the problem.

Representing the problem as an AST facilitates subsequent computations. For example, **Convex** can efficiently check if a problem is convex by applying the rules of *disciplined convex programming* (DCP), pioneered by Michael Grant and Stephen Boyd in [23, 22]. **Convex** can also use the AST to convert the problem into a conic form optimization problem, allowing a solver access to a complete and computationally concise global description of the problem [41].

Julia.

The Julia language [6] is a high-level, high-performance dynamic programming language for technical computing. With a syntax familiar to users of other technical computing languages such as Matlab, it takes advantage of LLVM-based just-in-time (JIT) compilation [29] to approach and often match the performance of C [5]. In **Convex**, we make particular use of *multiple dispatch* in Julia, an object-oriented paradigm in which different methods may be called to implement a function depending on the data types (classes) of the arguments to the function, rather than having functions encapsulated inside classes [4].

The **Convex** project supports the assertion that “technical computing is [the] killer application [of multiple dispatch]” [4]. We show in this paper that Julia’s multiple dispatch allows the authors of a technical computing package to write extremely performant code using a high level of abstraction. Indeed, the abstraction and generality of the code has pushed the authors of this paper toward more abstract and general formulations of the mathematics of disciplined convex programming, while producing code whose performance rivals and often surpasses codes with similar functionality in other languages.

Moreover, multiple dispatch enforces a separation of the fundamental objects in **Convex**— functions — from the methods for operating on them. This separation has other benefits: in **Convex**, it is easy to implement new methods for operating on the AST of an optimization problem, opening the door to new structural paradigms for parsing and solving optimization problems.

Mathematical optimization.

A traditional and familiar form for specifying a mathematical optimization problem (MOP) is to write

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

with variable $x \in \mathbf{R}^n$. Here, a problem instance is specified as list of *functions* $f_0, \dots, f_m, h_1, \dots, h_p : \mathbf{R}^n \rightarrow \mathbf{R}$. The function f_0 is called the objective function; f_1, \dots, f_m are called the inequality constraint functions; and h_1, \dots, h_p are called the equality constraint functions. The field of mathematical optimization is concerned with finding methods for solving MOP.

Some special structural forms of MOP are well known. For example, if all the functions f_0, \dots, f_m and h_1, \dots, h_p are affine, the MOP is called a linear program (LP). When each of the functions f_0, \dots, f_m is convex, and h_1, \dots, h_p are all affine, it is called a convex program (CP). This paper focusses on solving CPs, which can be solved much more quickly than general MOPs [26].

Perhaps the most consistent theme of the convex optimization literature concerns the importance of using the *structural form* of the CP (that is, the properties of the functions $f_0, \dots, f_m, h_1, \dots, h_p$) in order to devise faster solution methods. The advantage of this structural approach is that it allows for a division of labor between *users* of optimization and *designers* of optimization algorithms. Designers of optimization algorithms write solvers specialized to a particular structural form of MOP. Users can take advantage of innovations in optimization algorithms and theory so long as they can *identify* the structural form of their problem, and apply the appropriate solver.

A number of general purpose solvers have been developed, including ACADO [25, 24, 2], Ipopt [7], NLOpt [26]), and the solvers in AMPL [17], GAMS [10, 11], and Excel [19], along with associated modeling languages that make it easy to formulate problems for these solvers, including AMPL [17], GAMS [10, 11], and JuMP [33]. However, these solvers can be significantly slower than ones designed for special problem classes [28, 38, 40]. A modeling language that retains information about the structure of the problem can often choose a solver that will significantly outperform a general purpose solver.

Cone programs.

Convex targets a special class of convex programs called *cone programs*, which includes LPs, QPs, QCQPs, SOCPs, SDPs, and exponential cone programs as special cases (see §4 for more details). Cone programs occur frequently in a variety of fields [31, 51, 52], can be solved quickly, both in theory and in practice [42].

A number of modeling systems have been developed to automatically perform translations of a problem into standard conic form. The first parser-solver, SDPSOL, was written using Bison and Flex in 1996, and was able to automatically parse and solve SDPs with a few hundred variables [54]. YALMIP [32] was its first modern successor, followed shortly thereafter by CVX [23], both of which embedded a convex optimization modeling language into MATLAB, a proprietary language. Diamond *et. al.* followed with CVXPY [15], a convex optimization modeling language in python which uses an object-oriented approach. The present paper concerns **Convex**, which borrows many ideas from CVXPY, but takes advantage of language features in Julia (notably, multiple dispatch) to make the modeling layer simpler and faster.

A few more specialized conic parsing/modeling tools have also appeared targeting small QPs and SOCPs with very tight performance requirements, including CVXGEN [35] and QCML [12].

Organization.

This paper is organized as follows. In §2, we discuss how **Convex** represents optimization problems, and show how this representation suits the requirements of both users and solvers. In §3, we show how to use the AST to verify convexity of the problem using the rules of disciplined convex programming. In §4, we show how to use the AST to transform the problem into conic form, and pass the problem to a conic solver. Finally, in §5 we compare the performance of **Convex** to other codes for disciplined convex programming.

2. FUNCTIONAL REPRESENTATION

In this section, we describe the functional representation used by **Convex**. The basic building block of **Convex** is called an *expression*, which can represent a variable, a constant, or a function of another expression. We discuss each kind of expression in turn.

Variables.

The simplest kind of expression in **Convex** is a variable. Variables in **Convex** are declared using the **Variable** keyword, along with the dimensions of the variable.

```
# Scalar variable
x = Variable()
```

```
# 4x1 vector variable
y = Variable(4)
```

```
# 4x2 matrix variable
z = Variable(4, 2)
```

Variables may also be declared as having special properties, such as being (entrywise) positive or negative, or (for a matrix) being symmetric, with nonnegative eigenvalues (*i.e.*, positive semidefinite). These properties are treated as constraints in any problem constructed using the variables.

```
# Positive scalar variable
x = Variable(Positive())
```

```
# Negative 4x1 vector variable
y = Variable(4, Negative())
```

```
# Symmetric positive semidefinite
# 4x4 matrix variable
z = Semidefinite(4)
```

Constants.

Numbers, vectors, and matrices present in the Julia environment are wrapped automatically into a **Constant** expression when used in a **Convex** expression. The **Constant** expression stores a pointer to the underlying value, so the expression and any derived expression will reflect changes made to the value after the expression is formed. In other words, a **Convex** expression is parametrized by the constants it encapsulates.

Atoms.

Functions of expressions may be composed using elements called *atoms*, which represent simple functions with known properties such as **+**, *****, **abs**, and **norm**. A list of some of the atoms available in **Convex** is given in Table 3.

Atoms are applied to expressions using operator overloading. Hence, `2+2` calls Julia’s built-in addition operator, while `2+x` calls the `Convex` addition method and returns a `Convex` expression with top-level atom `+` and arguments `2` and `x`. Many of the useful language features in Julia, such as arithmetic, array indexing, and matrix transpose are overloaded in `Convex` so they may be used with variables and expressions just as they are used with native Julia types.

For example, all of the following form valid expressions.

```
# indexing, multiplication, addition
e1 = y[1] + 2*x

# expressions can be affine, convex, or concave
e3 = sqrt(x) + log(x)

# more atoms
e2 = 4 * pos(x) + max(abs(y)) + norm(z[:,1],2)
```

For each atom, we define five methods that allow `Convex` to solve problems involving that atom. The methods are called on expressions, and are dispatched based on the type (*i.e.*, the top-level atom) of the expression.

- `sign` returns the *sign* of the function the top-level atom represents (positive, negative, or no sign) over the range of its arguments.
- `evaluate` evaluates the function the top-level atom represents on the values of its arguments.
- `curvature` returns the curvature of the function the top-level atom represents (convex, concave, affine, or neither) over the range of its arguments.
- `monotonicity` returns the curvature of the function the atom represents (nondecreasing, nonincreasing, or not monotonic) over the range of its arguments.
- `conic_form` returns the conic form of an expression; see §4 for details.

These methods are computed using local combination rules, so they are never wrong, but sometimes miss strong classifications that cannot be deduced from local information. For example, `square(x) + x - x` is nonnegative for any value of `x`, but `Convex` will return no sign.

To define a new atom that can be used to solve problems in `Convex`, it suffices to define these five methods. We discuss how these methods are used to verify convexity and solve conic problems in §3) and §4, respectively.

Atoms may also be defined by composing other atoms; we call these *composite* atoms. For example, `norm(x, 1)` is implemented as `sum(abs(x))`.

Expressions.

Formally, expressions are defined recursively to be a variable, a constant, or the result of applying an atom to one or more expressions. In `Convex`, each expression is represented as an AST with a variable or constant at each leaf node, and an atom at each internal node. We refer to the arguments to the atom at a certain node as the *children* of that node. The top-level atom is called the *head* of the AST.

Considered as a tree without leaves, the AST may be seen as representing a function formed by composing the atoms at each node according to the structure of the tree; with the

leaves, it may be interpreted as a function with a closure, a curried function, or simply a function of named arguments.

The AST is a directed acyclic graph (DAG) [18]. The idea of using a DAG to study global properties of optimization problems dates back to work by Kantorovich in the 1950s [27].

Expressions may be evaluated when every leaf node in their AST has already been assigned a value. For example,

```
x = Variable()
e = max(x,0)
x.value = -4
evaluate!(e)
```

evaluates to 0.

One novel aspect of `Convex` is that expressions are identified via a unique identifier computed as a hash on the name of the atom at the head of the AST, together with the unique identifiers of its children in the AST. Thus, two different instances of the same expression will automatically be assigned the same unique identifier. This identification allows `Convex` to identify when the same expression is used in multiple places in the problem, to memoize and reuse the results of computations on a given expression, and to reduce the complexity of the problem that is ultimately passed to the solver.

Constraints.

Constraints in `Convex` are declared using the standard comparison operators `<=`, `>=`, and `==`. They specify relations that must hold between two expressions. `Convex` does not distinguish between strict and non-strict inequality constraints.

```
# affine equality constraint
A = randn(3,4); b = randn(3)
c1 = A*y == b

# convex inequality constraint
c2 = norm(y,2) <= 2
```

Problems.

A *problem* in `Convex` consists of a *sense* (minimize, maximize, or satisfy), an objective (an expression to which the sense verb is to be applied), and a list of zero or more constraints which must be satisfied at the solution.

For example, the problem

$$\begin{aligned} & \text{minimize} && \|x\|_\infty \\ & \text{subject to} && x_1 + x_2 = 5 \\ & && x_3 \leq x_2, \end{aligned}$$

where $x \in \mathbf{R}^3$ is the optimization variable, can be expressed in `Convex` as

```
x = Variable(3)
constraints = [x[1]+x[2] == 5, x[3] <= x[2]]
p = minimize(norm_inf(x), constraints)
```

When the user types this code, the input problem description is parsed internally into a `Problem` object with two attributes: an objective and a list of constraints. Each expression appearing in the problem is represented as an AST. For example, the problem `p` constructed above has the following properties:

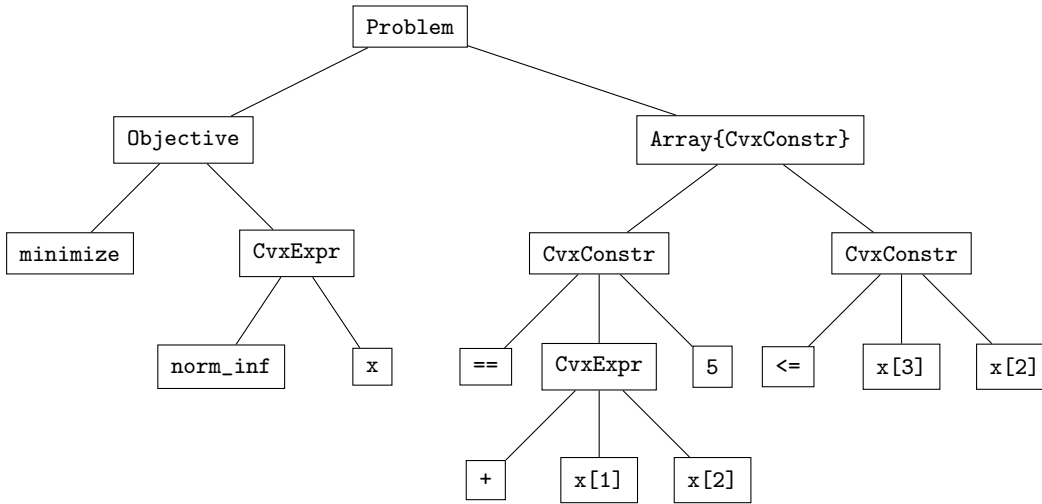


Figure 1: Graphical representation of a problem

```

p.objective = (:norm_inf, x)
p.constraints = [(:==, 5, (:+, (:getindex, x, 1),
                          (:getindex, x, 2))),
                (:<=, getindex(x, 3), getindex(x, 2))]
  
```

Here, we use Polish (prefix) notation to display the ASTs. Figure 1 gives a graphical representation of the structure of problem `p`.

Solving problems.

The `solve!` method in `Convex` checks that the problem is a disciplined convex program, converts it into conic form, and passes the problem to a solver:

```
solve!(p)
```

After `solve!` has been called on problem `p`, the optimal value of `p` can be queried with `p.optval`, and any expression `x` used in `p` is annotated with a value, accessible via `x.value`. If the solver computes dual variables, these are populated in the constraints; for example, the dual variable corresponding to the first constraint can be accessed as `p.constraints[1].dual_value`.

3. DISCIPLINED CONVEX PROGRAMMING

Checking if a function is convex is a difficult problem in general. Many approaches to verifying convexity have been proposed [45, 14], and implemented in modeling systems ranging from AMPL to Microsoft Excel [18, 39, 19].

In `Convex`, we have chosen to use the framework of *disciplined convex programming* (DCP) to verify problem convexity, which has a number of advantages. First, the simplicity of the DCP rules makes the determination of convexity transparent to the user [9]. Second, a problem that complies with the DCP rules can be converted to an equivalent conic form problem. Hence `Convex` is able to take advantage of extremely fast and reliable solvers for any problem whose convexity it can verify.

Disciplined convex expressions.

In `Convex`, we define five different kinds of curvature: constant, affine, convex, concave and not DCP. The simplest

expressions are single variables, which are affine, and constants, which are (unsurprisingly) constant. A constant expression is also affine, and an expression is affine if and only if it is convex and concave.

We determine the curvature of other expressions recursively using the following rule. Suppose that the curvatures of the expressions e_1, \dots, e_n are known. Then the expression $f(e_1, \dots, e_n)$ is convex if the function f is convex on the range of (e_1, \dots, e_n) , and for each $i = 1, \dots, n$,

- f is nondecreasing in its i th argument over the range of (e_1, \dots, e_n) , and e_i is convex;
- f is nonincreasing in its i th argument over the range of (e_1, \dots, e_n) , and e_i is concave; or
- e_i is affine.

The expression $f(e_1, \dots, e_n)$ is concave if $(-f)(e_1, \dots, e_n)$ is convex, and it is affine if it is both convex and concave. Note that the curvature of an expression need not be the same as the curvature of the top-level atom in that expression. For example, the atom `+` has affine curvature, but the expression `2 + square(x)` is convex.

Disciplined convex expressions (also called DCP compliant, or simply DCP, expressions) are defined inductively. We say that an expression is DCP if its curvature (affine, convex, or concave) can be inferred from these rules; otherwise, we say it is not DCP. Hence in order to decide if an expression is DCP, we must be able to calculate the curvature of f and the monotonicity of f in each argument for every atom f appearing in the expression.

The set of DCP expressions depends on the set of atoms. For example, $\log(\sum_{i=1}^n \exp x_i)$ is convex, but its convexity cannot be derived from the DCP rules using only the atoms `log` and `exp` [22]. In `Convex`, it is easy to add new atoms, allowing sophisticated users to expand the set of functions that `Convex` can recognize as convex.

The set of DCP expressions also depends on what is known about the monotonicity of atoms as a function of their arguments. For example, the quadratic-over-linear function $f(x, y) = x^T x / y$ is not convex on \mathbf{R}^2 , but it is convex on

$\mathbf{R} \times (0, \infty)$. This observation is known as a *signed* convexity rule [15]. Our formulation of the DCP rule above in terms of the *ranges* of the expressions (e_1, \dots, e_n) generalizes this observation. `Convex` implements a signed monotonicity rule: the monotonicity of an atom depends on the sign of the atom’s children.

Verifying DCP using multiple dispatch.

`Convex` checks that expressions are DCP using multiple dispatch.

Arithmetic is defined directly on `Vexity` and `Monotonicity` types so that addition of curvatures and multiplication of curvatures by monotonicities enforce the DCP rule. For example, adding `ConvexVexity` to `AffineVexity` results in `ConvexVexity`; multiplying `ConvexVexity` by `Nonincreasing` monotonicity results in `ConcaveVexity`. The curvature of an expression can then be calculated by applying the following function:

```
function vexity(x::AbstractExpr)
    monotonicities = monotonicity(x)
    vexity = curvature(x)
    for i = 1:length(x.children)
        vexity += monotonicities[i]
            * vexity(x.children[i])
    end
    return vexity
end
```

This approach to checking convexity has a number of advantages over using if-else statements to enforce the DCP rules. First, the implementation is aesthetically appealing: the code follows the mathematics. Second, this structure makes it easy to write new atoms, since one only needs to implement the `curvature` and `monotonicity` methods to allow `Convex` to determine the curvature of expressions involving the new atom. Third, it reduces the time needed to verify convexity, since multiple dispatch is implemented as a lookup table rather than via a (slower) `if` statement.

Disciplined convex constraints.

A constraint is called a *disciplined convex constraint* (or DCP constraint) if it has the form

- $e_1 \leq e_2$, where e_1 is a convex DCP expression and e_2 is a concave DCP expression;
- $e_1 \geq e_2$, where e_1 is a concave DCP expression and e_2 is a convex DCP expression; or
- $e_1 = e_2$, where e_1 and e_2 are both affine DCP expressions.

(For the purposes of this definition, remember that constant expressions are affine, and affine expressions are both convex and concave.)

Disciplined convex programs.

A problem is called a *disciplined convex program* if

1. the (sense, objective) are
 - (minimize, convex DCP expression),
 - (maximize, concave DCP expression), or
 - (satisfy); and

2. every constraint in the problem is DCP compliant.

(Note that the sense satisfy does not take an objective.) When an optimization problem is formed, `Convex` applies the DCP rules recursively to determine whether the problem is DCP.

4. CONIC FORM OPTIMIZATION

A conic form optimization problem is written as

$$\begin{aligned} & \text{minimize} && c^T x \\ & \text{subject to} && Ax = b \\ & && x \in \mathcal{K}, \end{aligned} \tag{1}$$

with variable x . Here \mathcal{K} is a *cone*: a set of points such that $x \in \mathcal{K}$ iff $rx \in \mathcal{K}$ for every $r \geq 0$. If in addition \mathcal{K} is convex, it is called a convex cone.

Examples of convex cones include the following.

- *The zero cone.* $\mathcal{K}_0^1 = \{0\}$
- *The free cone.* $\mathcal{K}_{\text{free}}^n = \mathbf{R}^n$
- *The positive orthant.* $\mathcal{K}_+^n = \{x \in \mathbf{R}^n : x \geq 0\}$
- *The second order cone.*
 $\mathcal{K}_{\text{SOC}}^{n+1} = \{(x, t) \in \mathbf{R}^n : \|x\| \leq t\}$
- *The semidefinite cone.*
 $\mathcal{K}_{\text{SDP}}^{n^2} = \{X \in \mathbf{R}^{n \times n} : \lambda_{\min}(X) \geq 0, X = X^T\}$
- *The exponential cone.*
 $\mathcal{K}_{\text{exp}}^3 = \{(x, y, z) \in \mathbf{R}^3 : ye^{x/y} \leq z, y > 0\}$

In general, a cone \mathcal{K} may also be given as a product of simpler cones,

$$\mathcal{K} = \mathcal{K}_1 \times \dots \times \mathcal{K}_p.$$

A conic form optimization problem is specified by the problem data A , b , c , and a cone K . In this paper, we extend this definition: we say a problem is in conic form if every expression in the problem is affine, and the constraints are all conic constraints, since it is trivial to rewrite such a problem in the form 1.

A number of fast solvers exist for conic form problems in a variety of languages. We list some of these solvers in Table 4.

`Convex` is able to automatically rewrite a problem specified as a DCP in standard cone format, and to pick a cone solver automatically depending on the cones required to represent the problem and the solvers available on the user’s computer, using the `loadconicproblem!` interface in MathProgBase [33]. This allows `Convex` to, for example, choose a SOCP solver if no exponential or SDP cones are present in the problem, which can result in significantly faster solve time [1].

4.1 Example

Here, we describe how `Convex` transforms a DCP problem into conic form.

Conic form expression.

The value of the function $f(x)$ (for fixed x) is the optimal value of the trivial optimization problem

$$\text{minimize } f(x),$$

with no variable. We say a function is *cone-representable* if the value of the function is the optimal value of a (nontrivial) conic form optimization problem. This conic form optimization problem is usually constructed by introducing auxiliary variables. For example, $f(x) = |x|$ is cone-representable, since $|x|$ is the optimal value of the conic problem

$$\begin{aligned} & \text{minimize} && t \\ & \text{subject to} && t - x \in \mathcal{K}_+ \\ & && t + x \in \mathcal{K}_+, \end{aligned}$$

with variable t .

In general, define the conic form template of a function f to be a constructor for a conic form optimization problem. The conic form template of f takes a list of arguments x_1, \dots, x_n , and returns a conic form optimization problem whose optimal value is $f(x_1, \dots, x_n)$. Concretely, it returns a sense, an objective, and a (possibly empty) set of constraints (s, o, \mathcal{C}) such that every expression appearing in the objective or constraints is affine in x_1, \dots, x_n . Note that the objective of the conic form problem may be vector or matrix valued. For example, the conic form template for $f(x) = Ax$ is simply

$$\text{minimize } Ax$$

(as a function of x).

The conic form template of a function is also known as the *graph form* of the function [21]. For convex functions f , any feasible point (x_1, \dots, x_n, t) for the conic form problem is a point in the epigraph $\{(x, t) : f(x) \leq t\}$ of f ; for concave functions, any feasible point of the conic form problem lies in the hypograph of f .

`Convex` uses the conic form templates of the atoms in an expression recursively in order to construct the conic form of an expression. Recall that every expression is represented as an AST with a function (atom) as its top-level node. We refer to the arguments to the function at a certain node as the children of that node. The conic form of an expression e is constructed recursively as follows:

- If e is affine, return (e, \emptyset) .
- Otherwise:
 1. Compute the conic forms $(s_i, o_i, \mathcal{C}_i)$ for each child $i = 1, \dots, n$ of the top-level node of e .
 2. Apply the conic form template for the top-level node to the list of objectives (o_1, \dots, o_n) , producing a conic form problem (s, o, \mathcal{C}) .
 3. Return $(s, o, \mathcal{C} \cup (\cup_i \mathcal{C}_i))$.

Conic form problem.

Now that we have defined the conic form for an expression, it is simple to define the conic form for an optimization problem. The conic form of an optimization problem, given as an objective and constraint set, is computed as follows:

1. Let (s, o, \mathcal{C}) be the conic form of the objective.
2. For each constraint in the constraint set,
 - (a) Compute the conic forms $(s_l, o_l, \mathcal{C}_l)$ and $(s_r, o_r, \mathcal{C}_r)$ for the left and right hand sides of the constraint.
 - (b) Add \mathcal{C}_l and \mathcal{C}_r to \mathcal{C} .

(c) If the sense of the constraint is

- \leq , add $o_r - o_l \in \mathcal{K}_+$ to \mathcal{C} ;
- \geq , add $o_l - o_r \in \mathcal{K}_+$ to \mathcal{C} ;
- $=$, add $o_r - o_l \in \mathcal{K}_0$ to \mathcal{C} .

(d) Return (s, o, \mathcal{C}) .

Notice that we have not used the sense of the conic forms of the arguments to an expression in constructing its conic form. This would be worrisome, were it not for the following theorem:

THEOREM 1 ([22]). *Let p be a problem with variable x and dual variable λ , and let $\Phi(p)$ be the conic form of the problem with variables x and t and dual variables λ and μ . Here we suppose t and μ are the primal and dual variables, respectively, that have been introduced in the transformation to conic form. If p is DCP, then any primal-dual solution (x, t, λ, μ) to $\Phi(p)$ provides a solution (x, λ) to p .*

To build intuition for this theorem, note that in a DCP expression, a convex function f will have objectives with the sense minimize spliced into argument slots in which f is increasing, and objectives with the sense maximize spliced into argument slots in which f is decreasing. At the solution, the coincidence of these senses, monotonicity, and curvature will force the variables participating in the conic form of each atom to lie on the graph of the atom, rather than simply in the *epigraph*. Using this reasoning, the theorem can be proved by induction on the depth of the AST.

4.2 Solvers

Julia has a rich ecosystem of optimization routines. The `JuliaOpt` project collects mathematical optimization solvers, and interfaces to solvers written in other languages, in a single GitHub repository, while `MathProgBase` enforces consistent interfaces to these solvers. Through integration with `MathProgBase`, `Convex` is able to use all of the solvers that accept problems in conic form, which includes all the linear programming solvers in `JuliaOpt` (including GLPK [34] and the commercial solvers CPLEX [13], Gurobi [13], and Mosek [36]), as well as the open source interior-point SOCP solver ECOS [16], and the open source first-order primal-dual conic solver SCS [43]. A list of solvers which can be used with `Convex` is presented in Table 4.

5. SPEED

Here we present a comparison of `Convex` with CVXPY² [15] and CVX³ [23] on a number of representative problems. The problems are chosen to be easy to *solve* (indeed, they are all easily solved by inspection) but difficult to *parse*. We concentrate on problems involving a large number of affine expressions; other atoms are treated very similarly in the different frameworks. The code for the tests in the three languages may be found in Appendix A.

We compare both the time required to convert the problem to conic form, and the time required to solve the resulting problem, since in general the conic form problems

²available at <https://pypi.python.org/pypi/cvxpy> as of October 15, 2014

³CVX v3.0 beta, available at www.cvxr.com as of October 15, 2014

produced in different modeling frameworks may be different. We use all three modeling languages with the solver ECOS [16], and call the solver with the same parameters (`abstol=1e-7`, `reltol=1e-7`, `feastol=1e-7`, `maxit=100`) so that solve times are comparable between the modeling frameworks.

Parse times are presented in Table 1 and solve times in Table 2. Parse times are computed by subtracting solve time from the total time to form and solve the problem.

We present two different timings for `Convex`. Julia uses just in time (JIT) compilation; the code is compiled at the first evaluation, and the faster compiled version is used on subsequent calls, so the first time a function is called is in general slower than the second time. The third column shows the runtime of the first (uncompiled) evaluation, and the fourth column the second (compiled) time. While the performance of `Convex` is comparable to CVX and CVXPY at first evaluation, it substantially outperforms upon the second evaluation. For applications in which convex optimization routines are called in an *inner* loop of a larger optimization algorithm (for example, in sequential convex programming), users will see the fast performance on every iteration of the loop after the first.

6. DISCUSSION

The parsing and modeling methodology applied here need not be restricted to conic form problems. Indeed, there are many other interesting and useful problem structures. In general, any problem structure that has been identified and for which there are specialized, fast solvers may be a good candidate for inclusion into a modeling framework like `Convex`. Such problem structures include integer linear programs (which have been a focus of the operations research community for decades, yielding fast solvers in practice [37, 53]); biconvex problems (which may be heuristically solved by alternating minimization [49]); sums of prox-capable functions (which can be efficiently solved by splitting methods like ADMM, even when the problem includes spectral functions of large matrices [8, 3]); difference-of-convex programs (for which sequential convex programming often produces successful solutions in practice [30, 50]); and sigmoidal programming problems (which often admit fast solutions via a specialized branch and bound method [47, 48]); to name only a few. The number of structural forms that might be useful far exceeds the number that a potential consumer of optimization might care to remember.

The framework developed in `Convex` makes it easy to write new transformations on the AST to detect and transform these problem types into a standard format. Automatic detection of some of the problem types mentioned above is the subject of ongoing research.

Significant work also remains to allow `Convex` to handle extremely large scale problems. For example, `Convex` currently performs all computations serially; developments in Julia’s parallelism ecosystem, and particularly its nascent threading capabilities, will enable `Convex` to leverage parallelism.

7. CONCLUSIONS

This paper shows that a convex optimization modeling library may be efficiently and simply implemented using multiple dispatch. The software package `Convex` uses language

features of Julia (notably, multiple dispatch and JIT compilation) to parse DCP problems into conic form as fast or faster than other (significantly more mature) codes for modeling convex optimization.

Acknowledgments

This work was developed with support from the National Science Foundation Graduate Research Fellowship program (under Grant No. DGE-1147470), the Gabilan Stanford Graduate Fellowship, the Gerald J. Lieberman Fellowship, and the DARPA X-DATA program.

8. REFERENCES

- [1] F. Alizadeh and D. Goldfarb. Second-order cone programming. *Mathematical programming*, 95(1):3–51, 2003.
- [2] D. Ariens, B. Houska, and H. Ferreau. Acado for Matlab user’s manual. www.acadotoolkit.org, 2010–2011.
- [3] S. Becker, E. Candes, and M. Grant. TFOCS: Flexible first-order methods for rank minimization. In *Low-rank Matrix Optimization Symposium, SIAM Conference on Optimization*, 2011.
- [4] J. Bezanson, J. Chen, S. Karpinski, V. Shah, and A. Edelman. Array operators using multiple dispatch: a design methodology for array implementations in dynamic languages. *arXiv preprint arXiv:1407.3845*, 2014.
- [5] J. Bezanson, S. Karpinski, V. B. Shah, and A. Edelman. The Julia language. <http://julialang.org/>.
- [6] J. Bezanson, S. Karpinski, V. B. Shah, and A. Edelman. Julia: A fast dynamic language for technical computing. *arXiv preprint arXiv:1209.5145*, 2012.
- [7] L. T. Biegler and V. M. Zavala. Large-scale nonlinear programming using ipopt: An integrating framework for enterprise-wide dynamic optimization. *Computers & Chemical Engineering*, 33(3):575–582, 2009.
- [8] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- [9] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [10] A. Brooke, D. Kendrick, A. Meeraus, and R. Raman. The general algebraic modeling system. *GAMS Development Corporation*, 1998.
- [11] M. R. Bussieck and A. Meeraus. General algebraic modeling system (GAMS). In *Modeling languages in mathematical optimization*, pages 137–157. Springer, 2004.
- [12] E. Chu, N. Parikh, A. Domahidi, and S. Boyd. Code generation for embedded second-order cone programming. In *European Control Conference (ECC)*, pages 1547–1552. IEEE, 2013.
- [13] CPLEX, IBM ILOG. V12.1: User’s manual for CPLEX. *International Business Machines Corporation*, 46(53):157, 2009.

Table 1: Speed comparisons: parse time (s)

Test	CVX	CVXPY	Convex	Convex compiled
sum	2.29	4.45	5.46	1.94
index	3.62	9.69	8.40	5.78
transpose	1.24	0.55	3.08	0.40
matrix constraint	1.38	0.54	2.49	0.34

Table 2: Speed comparisons: solve time (s)

Test	CVX	CVXPY	Convex	Convex compiled
sum	0.01	5e-4	2e-4	1e-4
index	0.01	0.08	0.076	0.061
transpose	0.66	1.54	14.23	11.99
matrix constraint	0.66	2.20	3.44	3.29

Table 3: Examples of atoms implemented in Convex

Function	Atom	Monotonicity	Curvature
Slicing and shaping atoms			
$f(x, i) = x_i$	<code>getindex(x, i)</code>	nondecreasing	Affine
$f(x, y) = [x \ y]$	<code>hcat(x, y)</code>	nondecreasing	Affine
$f(x, y) = [x, y]$	<code>vertcat(x, y)</code>	nondecreasing	Affine
$f(X) = (X_{11}, \dots, X_{nn})$	<code>diag(X)</code>	nondecreasing	Affine
$f(X) = X^T$	<code>transpose(X)</code>	nondecreasing	Affine
Positive orthant atoms			
$f(x) = \sum_i x_i$	<code>sum(x)</code>	nondecreasing	Affine
$f(x) = (x_1 , \dots, x_n)$	<code>abs(x)</code>	nondecreasing	Convex
$f(x) = \max_i x_i$	<code>max(x)</code>	nondecreasing	Convex
$f(x) = \min_i x_i$	<code>min(x)</code>	nondecreasing	Concave
$f(x) = \max(x, 0)$	<code>pos(x)</code>	nondecreasing	Convex
$f(x) = \max(-x, 0)$	<code>neg(x)</code>	nonincreasing	Convex
$f(x) = \ x\ _1$	<code>norm_1(x)</code>	nondecreasing $x \geq 0$ nonincreasing $x \leq 0$	Convex
$f(x) = \ x\ _\infty$	<code>norm_inf(x)</code>	nondecreasing	Convex
Second-order cone atoms			
$f(x) = \ x\ _2$	<code>norm_2(x)</code>	nondecreasing $x \geq 0$ nonincreasing $x \leq 0$	Convex
$f(X) = \ X\ _F$	<code>norm_fro(X)</code>	nondecreasing $X \geq 0$ nonincreasing $X \leq 0$	Convex
$f(x) = x^2$	<code>square(x)</code>	nondecreasing $x \geq 0$ nonincreasing $x \leq 0$	Convex
$f(x) = \sqrt{x}$	<code>sqrt(x)</code>	nondecreasing $x \geq 0$	Concave
$f(x, y) = \sqrt{xy}$	<code>geo_mean(x, y)</code>	nondecreasing	Concave
$f(x, y) = x^T x / y$	<code>quad_over_lin(x, y)</code>	nonincreasing in y for $y > 0$ nondecreasing in x for $x \geq 0$ nonincreasing in x for $x \leq 0$	Convex
$f(x) = 1/x$	<code>inv_pos(x)</code>	nonincreasing	Convex
$f(x) = \sum_i x_i^2$	<code>sum_squares</code>	nondecreasing in x for $x \geq 0$ nonincreasing in x for $x \leq 0$	Convex
Exponential cone atoms			
$f(x) = \exp(x)$	<code>exp(x)</code>	nondecreasing	Convex
$f(X) = \log(x)$	<code>log(x)</code>	nondecreasing $x > 0$	Concave
$f(x) = \log(\sum_i \exp(x_i))$	<code>logsumexp(x)</code>	nondecreasing	Convex
SDP atoms			
$f(X) = \ X\ _2$	<code>operatornorm(x)</code>		Convex
$f(X) = \ X\ _*$	<code>nuclearnorm(x)</code>		Convex

Table 4: Conic Form Problem Solvers

Name	Language	LP	SOCP	SDP	Exp	Method	Open source
CLP [46]	C++	X				Simplex	X
Gurobi [44]		X	X			Simplex, Interior Point, ...	
GLPK [34]	C	X				Simplex, Interior Point, ...	X
MOSEK [36]		X	X	X		Simplex, Interior Point, ...	
ECOS [16]	C	X	X			Interior Point	X
SDPA [20]	C++	X	X	X		Interior Point	X
SCS [43]	C	X	X	X	X	Primal-Dual Operator Splitting	X

- [14] C. Crusius. *Automated analysis of convexity properties of nonlinear programs*. PhD thesis, Stanford University, 2003.
- [15] S. Diamond, E. Chu, and S. Boyd. CVXPY: A Python-embedded modeling language for convex optimization, version 0.2. <http://cvxpy.org/>, May 2014.
- [16] A. Domahidi, E. Chu, and S. Boyd. ECOS: An SOCP solver for embedded systems. In *European Control Conference*, Zurich, Switzerland, 2013.
- [17] R. Fourer, D. Gay, and B. Kernighan. *AMPL*. Boyd & Fraser, 1993.
- [18] R. Fourer, C. Maheshwari, A. Neumaier, D. Orban, and H. Schicl. Convexity and concavity detection in computation graphs: Tree walks for convexity assessment. *INFORMS Journal on Computing*, 22:26–43, 2010.
- [19] Frontline Solvers. Excel solver, optimization software, Monte Carlo simulation, data mining - Frontline Systems. www.solver.com.
- [20] K. Fujisawa, M. Kojima, and K. Nakata. SDPA (semidefinite programming algorithm) user’s manual, version 4.10. Technical report, Department of Mathematical and Computing Science, Tokyo Institute of Technology, Tokyo, 1998.
- [21] M. Grant and S. Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer, 2008.
- [22] M. Grant, S. Boyd, and Y. Ye. Disciplined convex programming. In L. Liberti and N. Maculan, editors, *Global Optimization: From Theory to Implementation*, Nonconvex Optimization and its Applications, pages 155–210. Springer, 2006.
- [23] M. Grant, S. Boyd, and Y. Ye. CVX: Matlab software for disciplined convex programming, ver. 2.0, build 870. <http://cvxr.com>, Sept. 2012.
- [24] B. Houska and H. Ferreau. ACADO toolkit user’s manual. www.acadotoolkit.org, 2009–2011.
- [25] B. Houska, H. Ferreau, and M. Diehl. ACADO Toolkit – An open source framework for automatic control and dynamic optimization. *Optimal Control Applications and Methods*, 32(3):298–312, 2011.
- [26] S. G. Johnson. The NLOpt nonlinear-optimization package, 2010.
- [27] L. V. Kantorovich. On a mathematical symbolism convenient for performing machine calculations. *Dokl. Akad. Nauk SSSR*, 113(4):738–741, 1957.
- [28] N. Karmarkar. A new polynomial-time algorithm for linear programming. In *Proceedings of the 16th annual ACM Symposium on Theory of Computing*, pages 302–311. ACM, 1984.
- [29] C. Lattner and V. Adve. LLVM: A compilation framework for lifelong program analysis & transformation. In *International Symposium on Code Generation and Optimization*, pages 75–86. IEEE, 2004.
- [30] T. Lipp and S. Boyd. Variations and extensions of the convex-concave procedure. http://stanford.edu/~boyd/papers/cvx_ccv.html, 2014.
- [31] M. S. Lobo, L. Vandenberghe, S. Boyd, and H. Lebret. Applications of second-order cone programming. *Linear Algebra and Its Applications*, 284:193–228, 1998.
- [32] J. Lofberg. YALMIP: A toolbox for modeling and optimization in MATLAB. In *IEEE International Symposium on Computed Aided Control Systems Design*, pages 294–289, Sep 2004.
- [33] M. Lubin and I. Dunning. Computing in operations research using Julia. *arXiv preprint arXiv:1312.1431*, 2013.
- [34] A. Makhorin. GLPK (GNU linear programming kit), 2008.
- [35] J. Mattingley and S. Boyd. CVXGEN: A code generator for embedded convex optimization. *Optimization and Engineering*, 13:1–27, 2012.
- [36] A. Mosek. The MOSEK optimization software. 54, 2010.
- [37] G. Nemhauser and L. Wolsey. *Integer and combinatorial optimization*, volume 18. Wiley New York, 1988.
- [38] A. S. Nemirovsky and D. B. Yudin. *Informational Complexity and Efficient Methods for Solution of Convex Extremal Problems*. Wiley, New York, 1983.
- [39] I. P. Nenov, D. H. Fylstra, and L. V. Kolev. Convexity determination in the Microsoft Excel solver using automatic differentiation techniques. In *Fourth International Workshop on Automatic Differentiation*, 2004.
- [40] Y. Nesterov. Interior-point methods: An old and new approach to nonlinear programming. *Mathematical Programming*, 79(1):285–297, 1997.
- [41] Y. Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer, 2004.
- [42] Y. Nesterov and A. Nemirovskii. *Interior-point*

polynomial algorithms in convex programming,
volume 13. SIAM, 1994.

- [43] B. O’Donoghue, E. Chu, N. Parikh, and S. Boyd. Operator splitting for conic optimization via homogeneous self-dual embedding. *arXiv preprint arXiv:1312.3039*, 2013.
- [44] G. Optimization. Gurobi optimizer reference manual. 2012.
- [45] D. R. Stoutmeyer. Automatic categorization of optimization problems: An application of computer symbolic mathematics. *Operations Research*, 26:773–738, 1978.
- [46] The Computational Infrastructure for Operations Research project. COIN-OR linear program solver. 2014.
- [47] M. Udell and S. Boyd. Maximizing a sum of sigmoids. http://stanford.edu/~boyd/papers/max_sum_sigmoids.html, 2013.
- [48] M. Udell and S. Boyd. Bounding duality gap for separable functions. *arXiv preprint arXiv:1410.4158*, 2014.
- [49] M. Udell, C. Horn, R. Zadeh, and S. Boyd. Generalized low rank models. *arXiv preprint arXiv:1410.0342*, 2014.
- [50] T. Van Voorhis and F. A. Al-Khayyal. Difference of convex solution of quadratically constrained optimization problems. *European Journal of Operational Research*, 148(2):349–362, 2003.
- [51] L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM review*, 38(1):49–95, 1996.
- [52] H. Wolkowicz, R. Saigal, and L. Vandenberghe. *Handbook of semidefinite programming: theory, algorithms, and applications*, volume 27. Springer, 2000.
- [53] L. Wolsey. *Integer programming*, volume 42. Wiley New York, 1998.
- [54] S.-P. Wu and S. Boyd. SDPSOL: A parser/solver for SDP and MAXDET problems with matrix structure. Available at www.stanford.edu/~boyd/old_software/SDPSOL.html, Nov 1995.

APPENDIX

A. SPEED TESTS: CODE

A.1 Summation

Convex .

```
n = 10000
@time begin
x = Variable();
e = 0;
for i = 1:n
    e = e + x;
end
p = minimize(norm2(e-1), x >= 0);
solve!(p, ECOS.ECOSMathProgModel())
end
```

CVX.

```
n = 10000;
tic()
cvx_begin
variable x
e = 0
for i=1:n
    e = e + x;
end
minimize norm(e-1, 2)
subject to
    x >= 0;
cvx_end
toc()
```

CVXPY.

```
%%timeit
n = 10000;
x = Variable()
e = 0
for i in range(n):
    e = e + x
p = Problem(Minimize(norm(e-1,2)), [x>=0])
p.solve("ECOS", verbose=True)
```

A.2 Indexing

Convex .

```
n = 10000
@time begin
x = Variable(n);
e = 0;
for i = 1:n
    e = e + x[i];
end
p = minimize(norm2(e-1), x >= 0);
solve!(p, ECOS.ECOSMathProgModel())
end
```

CVX.

```
n = 10000;
tic()
cvx_begin
variable x(n)
e = 0;
for i=1:n
    e = e + x(i);
end
minimize norm(e - 1, 2)
subject to
    x >= 0;
cvx_end
toc()
```

CVXPY.

```
%%timeit
n = 10000
x = Variable(n)
```

```

e = 0
for i in range(n):
    e += x[i];
p = Problem(Minimize(norm(e-1,2)), [x>=0])
p.solve("ECOS", verbose=True)

```

A.3 Transpose

Convex .

```

n = 500
A = randn(n, n);
@time begin
X = Variable(n, n);
p = minimize(vecnorm(X' - A), X[1,1] == 1);
solve!(p, ECOS.ECOSMathProgModel())
end

```

CVX.

```

n = 500;
A = randn(n, n);
tic()
cvx_begin
variable X(n, n);
minimize(norm(transpose(X) - A, 'fro'))
subject to
    X(1,1) == 1;
cvx_end
toc()

```

CVXPY.

```

n = 500
A = numpy.random.randn(n,n)
%%timeit
X = Variable(n,n)
p = Problem(Minimize(norm(X.T-A,'fro')), [X[1,1] == 1])
p.solve("ECOS", verbose=True)

```

A.4 Matrix constraint

Convex .

```

n = 500;
A = randn(n, n);
B = randn(n, n);
@time begin
X = Variable(n, n);
p = minimize(vecnorm(X - A), X == B);
solve!(p, ECOS.ECOSMathProgModel())
end

```

CVX.

```

n = 500;
A = randn(n, n);
B = randn(n, n);
tic()
cvx_begin
variable X(n, n);

```

```

minimize(norm(X - A, 'fro'))
subject to
    X == B;
cvx_end
toc()

```

CVXPY.

```

n = 500
A = numpy.random.randn(n,n)
B = numpy.random.randn(n,n)
%%timeit
X = Variable(n,n)
p = Problem(Minimize(norm(X-A,'fro')), [X == B])
p.solve(verbose=True)

```