

CVXPY: A Python-Embedded Modeling Language for Convex Optimization

Steven Diamond

DIAMOND@CS.STANFORD.EDU

Stephen Boyd

BOYD@STANFORD.EDU

*Departments of Computer Science and Electrical Engineering
Stanford University
Stanford, CA 94305, USA*

Editor: Antti Honkela

Abstract

CVXPY is a domain-specific language for convex optimization embedded in Python. It allows the user to express convex optimization problems in a natural syntax that follows the math, rather than in the restrictive standard form required by solvers. CVXPY makes it easy to combine convex optimization with high-level features of Python such as parallelism and object-oriented design. CVXPY is available at <http://www.cvxpy.org/> under the GPL license, along with documentation and examples.

Keywords: convex optimization, domain-specific languages, Python, conic programming, convexity verification

1. Introduction

Convex optimization has many applications to fields as diverse as machine learning, control, finance, and signal and image processing (Boyd and Vandenberghe, 2004). Using convex optimization in an application requires either developing a custom solver or converting the problem into a standard form. Both of these tasks require expertise, and are time-consuming and error prone. An alternative is to use a domain-specific language (DSL) for convex optimization, which allows the user to specify the problem in a natural way that follows the math; this specification is then automatically converted into the standard form required by generic solvers. CVX (Grant and Boyd, 2014), YALMIP (Lofberg, 2004), QCML (Chu et al., 2013), PICOS (Sagnol, 2015), and Convex.jl (Udell et al., 2014) are examples of such DSLs for convex optimization.

CVXPY is a new DSL for convex optimization. It is based on CVX (Grant and Boyd, 2014), but introduces new features such as signed disciplined convex programming analysis and parameters. CVXPY is an ordinary Python library, which makes it easy to combine convex optimization with high-level features of Python such as parallelism and object-oriented design.

CVXPY has been downloaded by thousands of users and used to teach multiple courses (Boyd, 2015). Many tools have been built on top of CVXPY, such as an extension for stochastic optimization (Ali et al., 2015).

2. CVXPY Syntax

CVXPY has a simple, readable syntax inspired by CVX (Grant and Boyd, 2014). The following code constructs and solves a least squares problem where the variable's entries are constrained to be between 0 and 1. The problem data $A \in \mathbf{R}^{m \times n}$ and $b \in \mathbf{R}^m$ could be encoded as NumPy ndarrays or one of several other common matrix representations in Python.

```
# Construct the problem.
x = Variable(n)
objective = Minimize(sum_squares(A*x - b))
constraints = [0 <= x, x <= 1]
prob = Problem(objective, constraints)

# The optimal objective is returned by prob.solve().
result = prob.solve()
# The optimal value for x is stored in x.value.
print x.value
```

The variable, objective, and constraints are each constructed separately and combined in the final problem. In CVX, by contrast, these objects are created within the scope of a particular problem. Allowing variables and other objects to be created in isolation makes it easier to write high-level code that constructs problems (see §6).

3. Solvers

CVXPY converts problems into a standard form known as conic form (Nesterov and Nemirovsky, 1992), a generalization of a linear program. The conversion is done using graph implementations of convex functions (Grant and Boyd, 2008). The resulting cone program is equivalent to the original problem, so by solving it we obtain a solution of the original problem.

Solvers that handle conic form are known as cone solvers; each one can handle combinations of several types of cones. CVXPY interfaces with the open-source cone solvers CVXOPT (Andersen et al., 2015), ECOS (Domahidi et al., 2013), and SCS (O'Donoghue et al., 2016), which are implemented in combinations of Python and C. These solvers have different characteristics, such as the types of cones they can handle and the type of algorithms employed. CVXOPT and ECOS are interior-point solvers, which reliably attain high accuracy for small and medium scale problems; SCS is a first-order solver, which uses OpenMP to target multiple cores and scales to large problems with modest accuracy.

4. Signed DCP

Like CVX, CVXPY uses disciplined convex programming (DCP) to verify problem convexity (Grant et al., 2006). In DCP, problems are constructed from a fixed library of functions with known curvature and monotonicity properties. Functions must be composed according to a simple set of rules such that the composition's curvature is known. For a visualization of the DCP rules, visit `dcp.stanford.edu`.

CVXPY extends the DCP rules used in CVX by keeping track of the signs of expressions. The monotonicity of many functions depends on the sign of their argument, so keeping track of signs allows more compositions to be verified as convex. For example, the composition `square(square(x))` would not be verified as convex under standard DCP because the `square` function is nonmonotonic. But the composition is verified as convex under signed DCP because `square` is increasing for nonnegative arguments and `square(x)` is nonnegative.

5. Parameters

Another improvement in CVXPY is the introduction of parameters. Parameters are constants whose symbolic properties (*e.g.*, dimensions and sign) are fixed but whose numeric value can change. A problem involving parameters can be solved repeatedly for different values of the parameters without redoing computations that do not depend on the parameter values. Parameters are an old idea in DSLs for optimization, appearing in AMPL (Fourer et al., 2002).

A common use case for parameters is computing a trade-off curve. The following code constructs a LASSO problem (Boyd and Vandenberghe, 2004) where the positive parameter γ trades off the sum of squares error and the regularization term. The problem data are $A \in \mathbf{R}^{m \times n}$ and $b \in \mathbf{R}^m$.

```
x = Variable(n)
gamma = Parameter(sign="positive") # Must be positive due to DCP rules.
error = sum_squares(A*x - b)
regularization = norm(x, 1)
prob = Problem(Minimize(error + gamma*regularization))
```

Computing a trade-off curve is trivially parallelizable, since each problem can be solved independently. CVXPY can be combined with Python multiprocessing (or any other parallelism library) to distribute the trade-off curve computation across many processes.

```
# Assign a value to gamma and find the optimal x.
def get_x(gamma_value):
    gamma.value = gamma_value
    result = prob.solve()
    return x.value

# Get a range of gamma values with NumPy.
gamma_vals = numpy.logspace(-4, 6)
# Do parallel computation with multiprocessing.
pool = multiprocessing.Pool(processes = N)
x_values = pool.map(get_x, gamma_vals)
```

6. Object-Oriented Convex Optimization

CVXPY enables an object-oriented approach to constructing optimization problems. As an example, consider an optimal flow problem on a directed graph $G = (V, E)$ with vertex

set V and (directed) edge set E . Each edge $e \in E$ carries a flow $f_e \in \mathbf{R}$, and each vertex $v \in V$ has an internal source that generates $s_v \in \mathbf{R}$ flow. (Negative values correspond to flow in the opposite direction, or a sink at a vertex.) The (single commodity) flow problem is (with variables f_e and s_v)

$$\begin{aligned} & \text{minimize} && \sum_{e \in E} \phi_e(f_e) + \sum_{v \in V} \psi_v(s_v), \\ & \text{subject to} && s_v + \sum_{e \in I(v)} f_e = \sum_{e \in O(v)} f_e, \quad \text{for all } v \in V, \end{aligned}$$

where the ϕ_e and ψ_v are convex cost functions and $I(v)$ and $O(v)$ give vertex v 's incoming and outgoing edges, respectively.

To express the problem in CVXPY, we construct vertex and edge objects, which store local information such as optimization variables, constraints, and an associated objective term. These are exported as a CVXPY problem for each vertex and each edge.

```
class Vertex(object):
    def __init__(self, cost):
        self.source = Variable()
        self.cost = cost(self.source)
        self.edge_flows = []

    def prob(self):
        net_flow = sum(self.edge_flows) + self.source
        return Problem(Minimize(self.cost), [net_flow == 0])

class Edge(object):
    def __init__(self, cost):
        self.flow = Variable()
        self.cost = cost(self.flow)

    def connect(self, in_vertex, out_vertex):
        in_vertex.edge_flows.append(-self.flow)
        out_vertex.edge_flows.append(self.flow)

    def prob(self):
        return Problem(Minimize(self.cost))
```

The vertex and edge objects are composed into a graph using the edges' `connect` method. To construct the single commodity flow problem, we sum the vertices and edges' local problems. (Addition of problems is overloaded in CVXPY to add the objectives together and concatenate the constraints.)

```
prob = sum([object.prob() for object in vertices + edges])
prob.solve() # Solve the single commodity flow problem.
```

Acknowledgments

We thank the many contributors to CVXPY. This work was supported by DARPA XDATA.

References

- A. Ali, Z. Kolter, S. Diamond, and S. Boyd. Disciplined convex stochastic programming: A new framework for stochastic optimization. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pages 62–71, 2015.
- M. Andersen, J. Dahl, and L. Vandenberghe. CVXOPT: Python software for convex optimization, version 1.1. <http://cvxopt.org/>, May 2015.
- S. Boyd. EE364a: Convex optimization I. <http://stanford.edu/class/ee364a/>, December 2015.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- E. Chu, N. Parikh, A. Domahidi, and S. Boyd. Code generation for embedded second-order cone programming. In *Proceedings of the European Control Conference*, pages 1547–1552, 2013.
- A. Domahidi, E. Chu, and S. Boyd. ECOS: An SOCP solver for embedded systems. In *Proceedings of the European Control Conference*, pages 3071–3076, 2013.
- R. Fourer, D. Gay, and B. Kernighan. *AMPL: A Modeling Language for Mathematical Programming*. Cengage Learning, 2002.
- 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.
- M. Grant and S. Boyd. CVX: MATLAB software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, March 2014.
- 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.
- J. Lofberg. YALMIP: A toolbox for modeling and optimization in MATLAB. In *Proceedings of the IEEE International Symposium on Computed Aided Control Systems Design*, pages 294–289, 2004.
- Y. Nesterov and A. Nemirovsky. Conic formulation of a convex programming problem and duality. *Optimization Methods and Software*, 1(2):95–115, 1992.
- B. O’Donoghue, E. Chu, N. Parikh, and S. Boyd. Conic optimization via operator splitting and homogeneous self-dual embedding. *Journal of Optimization Theory and Applications*, pages 1–27, 2016.
- G. Sagnol. PICOS: A Python interface for conic optimization solvers, version 1.1. <http://picos.zib.de/index.html>, April 2015.
- M. Udell, K. Mohan, D. Zeng, J. Hong, S. Diamond, and S. Boyd. Convex optimization in Julia. In *Proceedings of the Workshop for High Performance Technical Computing in Dynamic Languages*, pages 18–28, 2014.