
Convex.jl Documentation

Release 0.1

Karanveer Mohan, Madeleine Udell, David Zeng, Jenny Hong, Steven

August 26, 2016

| | | |
|----------|--------------------------------|----------|
| 1 | In Depth Documentation: | 3 |
| 1.1 | Installation | 3 |
| 1.2 | Quick Tutorial | 3 |
| 1.3 | Basic Types | 4 |
| 1.4 | Operations | 6 |
| 1.5 | Examples | 9 |
| 1.6 | Solvers | 9 |
| 1.7 | FAQ | 10 |
| 1.8 | Optimizing in a Loop | 12 |
| 1.9 | Advanced Features | 12 |
| 1.10 | Contributing | 14 |
| 1.11 | Credits | 15 |

Convex.jl is a Julia package for [Disciplined Convex Programming](#) (DCP). Convex.jl makes it easy to describe optimization problems in a natural, mathematical syntax, and to solve those problems using a variety of different (commercial and open-source) solvers. Convex.jl can solve

- linear programs
- mixed-integer linear programs and mixed-integer second-order cone programs
- dcp-compliant convex programs including
 - second-order cone programs (SOCP)
 - exponential cone programs
 - semidefinite programs (SDP)

Convex.jl supports many solvers, including [Mosek](#), [Gurobi](#), [ECOS](#), [SCS](#) and [GLPK](#), through the [MathProgBase](#) interface.

Note that Convex.jl was previously called CVX.jl. This package is under active development; we welcome bug reports and feature requests. For usage questions, please contact us via the [JuliaOpt mailing list](#).

In Depth Documentation:

1.1 Installation

Installing Convex.jl is a one step process. Open up Julia and type

```
Pkg.update()
Pkg.add("Convex")
```

This does not install any solvers. If you don't have a solver installed already, you will want to install a solver such as [SCS](#) by running

```
Pkg.add("SCS")
```

To solve certain problems such as mixed integer programming problems you will need to install another solver as well, such as [GLPK](#). If you wish to use other solvers, please read the section on solvers.

1.2 Quick Tutorial

Consider a constrained least squares problem

$$\begin{array}{ll} \text{minimize} & \|Ax - b\|_2^2 \\ \text{subject to} & x \geq 0 \end{array}$$

with variable $x \in \mathbf{R}^n$, and problem data $A \in \mathbf{R}^{m \times n}$, $b \in \mathbf{R}^m$.

This problem can be solved in Convex.jl as follows:

```
# Make the Convex.jl module available
using Convex

# Generate random problem data
m = 4; n = 5
A = randn(m, n); b = randn(m, 1)

# Create a (column vector) variable of size n x 1.
x = Variable(n)

# The problem is to minimize ||Ax - b||^2 subject to x >= 0
# This can be done by: minimize(objective, constraints)
problem = minimize(sumsquares(A * x - b), [x >= 0])
```

```
# Solve the problem by calling solve!
solve!(problem)

# Check the status of the problem
problem.status # :Optimal, :Infeasible, :Unbounded etc.

# Get the optimum value
problem.optval
```

1.3 Basic Types

The basic building block of Convex.jl 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.

1.3.1 Variables

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

```
# Scalar variable
x = Variable()

# Column vector variable
x = Variable(5)

# Matrix variable
x = Variable(4, 6)
```

Variables may also be declared as having special properties, such as being

- (entrywise) positive: `x = Variable(4, Positive())`
- (entrywise) negative: `x = Variable(4, Negative())`
- integral: `x = Variable(4, :Int)`
- binary: `x = Variable(4, :Bin)`
- (for a matrix) being symmetric, with nonnegative eigenvalues (ie, positive semidefinite): `z = Semidefinite(4)`

1.3.2 Constants

Numbers, vectors, and matrices present in the Julia environment are wrapped automatically into a *Constant* expression when used in a Convex.jl expression.

1.3.3 Expressions

Expressions in Convex.jl are formed by applying any *atom* (mathematical function defined in Convex.jl) to variables, constants, and other expressions. For a list of these functions, see [operations](#). Atoms are applied to expressions using operator overloading. For example, `2+2` calls Julia's built-in addition operator, while `2+x` calls the Convex.jl addition method and returns a Convex.jl expression. Many of the useful language features in Julia, such as arithmetic, array

indexing, and matrix transpose are overloaded in Convex.jl so they may be used with variables and expressions just as they are used with native Julia types.

Expressions that are created must be DCP-compliant. More information on DCP can be found [here](#).

```
x = Variable(5)
# The following are all expressions
y = sum(x)
z = 4 * x + y
z_1 = z[1]
```

Convex.jl allows the values of the expressions to be evaluated directly.

```
x = Variable()
y = Variable()
z = Variable()
expr = x + y + z
problem = minimize(expr, x >= 1, y >= x, 4 * z >= y)
solve!(problem)

# Once the problem is solved, we can call evaluate() on expr:
evaluate(expr)
```

1.3.4 Constraints

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

```
x = Variable(5, 5)
# Equality constraint
constraint = x == 0
# Inequality constraint
constraint = x >= 1
```

Matrices can also be constrained to be positive semidefinite.

```
x = Variable(3, 3)
y = Variable(3, 1)
z = Variable()
# constrain [x y; y' z] to be positive semidefinite
constraint = ([x y; y' z] in :SDP)
# or equivalently,
constraint = ([x y; y' z] 0)
```

1.3.5 Objective

The objective of the problem is a scalar expression to be maximized or minimized by using `maximize` or `minimize` respectively. Feasibility problems can be expressed by either giving a constant as the objective, or using `problem = satisfy(constraints)`.

1.3.6 Problem

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

```

problem = minimize(objective, constraints)
# or
problem = maximize(objective, constraints)
# or
problem = satisfy(constraints)

```

Constraints can be added at any time before the problem is solved.

```

# No constraints given
problem = minimize(objective)
# Add some constraint
problem.constraints += constraint
# Add many more constraints
problem.constraints += [constraint1, constraint2, ...]

```

A problem can be solved by calling `solve!`:

```
solve!(problem)
```

After the problem is solved, `problem.status` records the status returned by the optimization solver, and can be `:Optimal`, `:Infeasible`, `:Unbounded`, `:Indeterminate` or `:Error`. If the status is `:Optimal`, `problem.optval` will record the optimum value of the problem. The optimal value for each variable `x` participating in the problem can be found in `x.value`. The optimal value of an expression can be found by calling the `evaluate()` function on the expression as follows: `evaluate(expr)`.

1.4 Operations

Convex.jl currently supports the following functions. These functions may be composed according to the [DCP](#) composition rules to form new convex, concave, or affine expressions. Convex.jl transforms each problem into an equivalent [cone program](#) in order to pass the problem to a specialized solver. Depending on the types of functions used in the problem, the conic constraints may include linear, second-order, exponential, or semidefinite constraints, as well as any binary or integer constraints placed on the variables. Below, we list each function available in Convex.jl organized by the (most complex) type of cone used to represent that function, and indicate which solvers may be used to solve problems with those cones. Problems mixing many different conic constraints can be solved by any solver that supports every kind of cone present in the problem.

In the notes column in the tables below, we denote implicit constraints imposed on the arguments to the function by IC, and parameter restrictions that the arguments must obey by PR. (Convex.jl will automatically impose ICs; the user must make sure to satisfy PRs.) Elementwise means that the function operates elementwise on vector arguments, returning a vector of the same size.

1.4.1 Linear Program Representable Functions

An optimization problem using only these functions can be solved by any LP solver.

| operation | description | vexity | slope | notes |
|-----------------|-------------|--------|--|--------------|
| $x+y$ or $x.+y$ | addition | affine | increasing | none |
| $x-y$ or $x.-y$ | subtraction | affine | increasing in x decreasing in y | none none |

Continued on next page

Table 1.1 – continued from previous page

| operation | description | vexity | slope | notes |
|--|---|---------|--|---------------------------------|
| <code>x*y</code> | multiplication | affine | increasing if constant term ≥ 0 decreasing if constant term ≤ 0 not monotonic otherwise | PR: one argument is constant |
| <code>x/y</code> | division | affine | increasing | PR: y is scalar con- stant |
| <code>x .* y</code> | elementwise multipli- cation | affine | increasing | PR: one argument is constant |
| <code>x ./ y</code> | elementwise division | affine | increasing | PR: one argument is constant |
| <code>x[1:4, 2:3]</code> | indexing and slicing | affine | increasing | none |
| <code>diag(x, k)</code> | k -th diagonal of a ma- trix | affine | increasing | none |
| <code>diagm(x)</code> | construct diagonal matrix | affine | increasing | PR: x is a vector |
| <code>x'</code> | transpose | affine | increasing | none |
| <code>vec(x)</code> | vector representation | affine | increasing | none |
| <code>dot(x, y)</code> | $\sum_i x_i y_i$ | affine | increasing | PR: one argument is constant |
| <code>vecdot(x, y)</code> | <code>dot(vec(x), vec(y))</code> | affine | increasing | PR: one argument is constant |
| <code>sum(x)</code> | $\sum_{ij} x_{ij}$ | affine | increasing | none |
| <code>sum(x, k)</code> | sum elements across dimension k | affine | increasing | none |
| <code>sumlargest(x, k)</code> | sum of k largest elements of x | convex | increasing | none |
| <code>sumsmallest(x, k)</code> | sum of k smallest elements of x | concave | increasing | none |
| <code>dotsort(a, b)</code> | <code>dot(sort(a), sort(b))</code> | convex | increasing | PR: one argument is constant |
| <code>reshape(x, m, n)</code> | reshape into $m \times n$ | affine | increasing | none |
| <code>minimum(x)</code> | $\min(x)$ | concave | increasing | none |
| <code>maximum(x)</code> | $\max(x)$ | convex | increasing | none |
| <code>[x y]</code> or <code>[x; y]</code> <code>hcat(x, y)</code> or <code>vcat(x, y)</code> | stacking | affine | increasing | none |
| <code>trace(x)</code> | $\text{tr}(X)$ | affine | increasing | none |
| <code>conv(h, x)</code> | $h \in \mathbb{R}^m$ $x \in \mathbb{R}^m$ $h * x \in \mathbb{R}^{m+n-1}$ entry i is given by $\sum_{j=1}^m h_j x_{i-j}$ | affine | increasing if $h \geq 0$ decreasing if $h \leq 0$ not monotonic otherwise | PR: h is constant |
| <code>min(x, y)</code> | $\min(x, y)$ | concave | increasing | none |
| <code>max(x, y)</code> | $\max(x, y)$ | convex | increasing | none |
| <code>pos(x)</code> | $\max(x, 0)$ | convex | increasing | none |
| <code>neg(x)</code> | $\max(-x, 0)$ | convex | decreasing | none |
| <code>invpos(x)</code> | $1/x$ | convex | decreasing | IC: $x > 0$ |

Continued on next page

Table 1.1 – continued from previous page

| operation | description | vexity | slope | notes |
|-----------|-------------|--------|--|-------|
| abs (x) | $ x $ | convex | increasing on $x \geq 0$ decreasing on $x \leq 0$ | none |

1.4.2 Second-Order Cone Representable Functions

An optimization problem using these functions can be solved by any SOCP solver (including ECOS, SCS, Mosek, Gurobi, and CPLEX). Of course, if an optimization problem has both LP and SOCP representable functions, then any solver that can solve both LPs and SOCPs can solve the problem.

| operation | description | vexity | slope | notes |
|-------------------------|---|--------------------------------|---|--|
| norm(x, p) | $(\sum x_i^p)^{1/p}$ | convex | increasing on $x \geq 0$ decreasing on $x \leq 0$ | PR: $p \geq 1$ |
| vecnorm(x, p) | $(\sum x_{ij}^p)^{1/p}$ | convex | increasing on $x \geq 0$ decreasing on $x \leq 0$ | PR: $p \geq 1$ |
| quadform(x, P) | $x^T P x$ | convex in x affine in P | increasing on $x \geq 0$ decreasing on $x \leq 0$ increasing in P | PR: either x or P must be constant; if x is not constant, then P must be symmetric and positive semidefinite |
| quadoverlin(x, y) | $x^T x / y$ | convex | increasing on $x \geq 0$ decreasing on $x \leq 0$ decreasing in y | IC: $y > 0$ |
| sumsquares(x) | $\sum x_i^2$ | convex | increasing on $x \geq 0$ decreasing on $x \leq 0$ | none |
| sqrt(x) | \sqrt{x} | concave | decreasing | IC: $x > 0$ |
| square(x), x.^2, x^2 | x^2 | convex | increasing on $x \geq 0$ decreasing on $x \leq 0$ | elementwise |
| geomean(x, y) | \sqrt{xy} | concave | increasing | IC: $x \geq 0, y \geq 0$ |
| huber(x, M=1) | $\begin{cases} x^2 & x \leq M \\ 2M x - M^2 & x > M \end{cases}$ | convex | increasing on $x \geq 0$ decreasing on $x \leq 0$ | PR: $M \geq 1$ |

1.4.3 Exponential Cone Representable Functions

An optimization problem using these functions can be solved by any exponential cone solver (SCS).

| operation | description | vexity | slope | notes |
|-----------------|----------------------------------|---------|---------------|-------------|
| logsumexp(x) | $\log(\sum_i \exp(x_i))$ | convex | increasing | none |
| exp(x) | $\exp(x)$ | convex | increasing | none |
| log(x) | $\log(x)$ | concave | increasing | IC: $x > 0$ |
| entropy(x) | $\sum_{ij} -x_{ij} \log(x_{ij})$ | concave | not monotonic | IC: $x > 0$ |
| logisticloss(x) | $\log(1 + \exp(x_i))$ | convex | increasing | none |

1.4.4 Semidefinite Program Representable Functions

An optimization problem using these functions can be solved by any SDP solver (including SCS and Mosek).

| operation | description | vexity | slope | notes |
|-------------------------------|-------------------------------|----------|---------------|--------------------------------|
| <code>nuclearnorm(x)</code> | sum of singular values of x | convex | not monotonic | none |
| <code>operatornorm(x)</code> | max of singular values of x | convex | not monotonic | none |
| <code>lambdamax(x)</code> | max eigenvalue of x | convex | not monotonic | none |
| <code>lambdamin(x)</code> | min eigenvalue of x | con-cave | not monotonic | none |
| <code>matrixfrac(x, P)</code> | $x^T P^{-1} x$ | convex | not monotonic | IC: P is positive semidefinite |

1.4.5 Exponential + SDP representable Functions

An optimization problem using these functions can be solved by any solver that supports exponential constraints *and* semidefinite constraints simultaneously (SCS).

| operation | description | vexity | slope | notes |
|------------------------|---------------------------|---------|------------|----------------------------------|
| <code>logdet(x)</code> | log of determinant of x | concave | increasing | IC: x is positive semidefinite |

1.4.6 Promotions

When an atom or constraint is applied to a scalar and a higher dimensional variable, the scalars are promoted. For example, we can do `max(x, 0)` gives an expression with the shape of x whose elements are the maximum of the corresponding element of x and 0.

1.5 Examples

- Basic usage
- Control
- Tomography
- Time series
- Section allocation
- Binary knapsack
- Entropy maximization
- Logistic regression

1.6 Solvers

Convex.jl transforms each problem into an equivalent [cone program](#) in order to pass the problem to a specialized solver. Depending on the types of functions used in the problem, the conic constraints may include linear, second-order, exponential, or semidefinite constraints, as well as any binary or integer constraints placed on the variables.

By default, Convex.jl does not install any solvers. Many users use the solver [SCS](#), which is able to solve problems with linear, second-order cone constraints (SOCPs), exponential constraints and semidefinite constraints (SDPs). Any other solver in [JuliaOpt](#) may also be used, so long as it supports the conic constraints used to represent the problem.

Most other solvers in the JuliaOpt ecosystem can be used to solve (mixed integer) linear programs (LPs and MILPs). Mosek and Gurobi can be used to solve SOCPs (even with binary or integer constraints), and Mosek can also solve SDPs. For up-to-date information about solver capabilities, please see the table [here](#) describing which solvers can solve which kind of problems.

Installing these solvers is very simple. Just follow the instructions in the documentation for that solver.

To use a specific solver, you can use the following syntax

```
solve!(p, GurobiSolver())
solve!(p, MosekSolver())
solve!(p, GLPKSolverMIP())
solve!(p, GLPKSolverLP())
solve!(p, ECOSolver())
solve!(p, SCSSolver())
```

(Of course, the solver must be installed first.) For example, we can use GLPK to solve a MILP

```
using GLPKMathProgInterface
solve!(p, GLPKSolverMIP())
```

You can set or see the current default solver by

```
get_default_solver()
using Gurobi
set_default_solver(GurobiSolver()) # or set_default_solver(SCSSolver(verbose=0))
# Now Gurobi will be used by default as a solver
```

Many of the solvers also allow options to be passed in. More details can be found in each solver's documentation.

For example, if we wish to turn off printing for the SCS solver (ie, run in quiet mode), we can do so by

```
using SCS
solve!(p, SCSSolver(verbose=false))
```

If we wish to increase the maximum number of iterations for ECOS or SCS, we can do so by

```
using ECOS
solve!(p, ECOSolver(maxit=10000))
using SCS
solve!(p, SCSSolver(max_iters=10000))
```

To turn off the problem status warning issued by Convex when a solver is not able to solve a problem to optimality, use the keyword argument `verbose=false` of the solve method, along with any desired solver parameters:

```
solve!(p, SCSSolver(verbose=false), verbose=false)
```

1.7 FAQ

- *Where can I get help?*
- *How does Convex.jl differ from JuMP?*
- *Where can I learn more about Convex Optimization?*
- *Are there similar packages available in other languages?*
- *How does Convex.jl work?*
- *How do I cite this package?*

1.7.1 Where can I get help?

For usage questions, please contact us via the [JuliaOpt mailing list](#). If you're running into bugs or have feature requests, please use the [Github Issue Tracker](#).

1.7.2 How does Convex.jl differ from JuMP?

Convex.jl and JuMP are both modelling languages for mathematical programming embedded in Julia, and both interface with solvers via the MathProgBase interface, so many of the same solvers are available in both. Convex.jl converts problems to a standard conic form. This approach requires (and certifies) that the problem is convex and DCP compliant, and guarantees global optimality of the resulting solution. JuMP allows nonlinear programming through an interface that learns about functions via their derivatives. This approach is more flexible (for example, you can optimize non-convex functions), but can't guarantee global optimality if your function is not convex, or warn you if you've entered a non-convex formulation.

For linear programming, the difference is more stylistic. JuMP's syntax is scalar-based and similar to AMPL and GAMS making it easy and fast to create constraints by indexing and summation (like `sum{x[i], i=1:numLocation}`). Convex.jl allows (and prioritizes) linear algebraic and functional constructions (like `max(x, y) < A*z`); indexing and summation are also supported in Convex.jl, but are somewhat slower than in JuMP. JuMP also lets you efficiently solve a sequence of problems when new constraints are added or when coefficients are modified, whereas Convex.jl parses the problem again whenever the *solve!* method is called.

1.7.3 Where can I learn more about Convex Optimization?

See the freely available book [Convex Optimization](#) by Boyd and Vandenberghe for general background on convex optimization. For help understanding the rules of Disciplined Convex Programming, we recommend this [DCP tutorial website](#).

1.7.4 Are there similar packages available in other languages?

Indeed! You might use [CVXPY](#) in Python, or [CVX](#) in Matlab.

1.7.5 How does Convex.jl work?

For a detailed discussion of how Convex.jl works, see [our paper](#).

1.7.6 How do I cite this package?

If you use Convex.jl for published work, we encourage you to cite the software using the following BibTeX citation:

```
@article{convexjl,
  title = {Convex Optimization in {J}ulia},
  author = {Udell, Madeleine and Mohan, Karanveer and Zeng, David and Hong, Jenny and Diamond, Steven and ...},
  year = {2014},
  journal = {SC14 Workshop on High Performance Technical Computing in Dynamic Languages},
  archivePrefix = "arXiv",
  eprint = {1410.4821},
  primaryClass = "math-oc",
}
```

1.8 Optimizing in a Loop

1.8.1 Memory Management

Convex uses a module-level dictionary to store the conic forms of every variable and expression created in the same Julia session. These variables and expressions persist even after they are out of scope. If you create large numbers of variables inside a loop, this dictionary can eat a considerable amount of memory.

To flush the memory, you can call

```
Convex.clearmemory()
```

This will remove every variable and expression you've formed before from the memory cache, so that you're starting as fresh as if you'd just reimported Convex.

1.8.2 Caching Expressions

Better yet, take advantage of this cache of variables and expressions! Create variables and expressions outside the loop, and reuse them inside the loop as you tweak parameters. Doing this will allow Convex to reuse the conic forms it has already calculated for previously used expressions.

For example, the following **bad** code will create a new instance of a variable and of the expression `square(x)` for each value of `i`. Don't do this:

```
for i=1:10
    x = Variable()
    p = minimize(square(x), x >= i)
    solve!(p)
end
```

Contrast this with the following **good** code, which will reuse the cached conic form for `square(x)` for each `i`, reducing the memory footprint and speeding up the computation. Do this instead:

```
x = Variable()
obj = square(x)
for i=1:10
    p = minimize(obj, x >= i)
    solve!(p)
end
```

1.8.3 Warmstarts, Parameters, Fixing and Freeing Variables

If you're solving many problems of the same form, or many similar problems, you may also want to use warmstarts, or to dynamically fix and free variables. The former is particularly good for a family of problems related by a parameter; the latter allows for easy implementation of alternating minimization for nonconvex problems. See the Advanced Features section of the documentation for more information.

1.9 Advanced Features

1.9.1 Dual Variables

Convex.jl also returns the optimal dual variables for a problem. These are stored in the `dual` field associated with each constraint.


```
using Convex

x = Variable()
constraint = x >= 0
p = minimize(x, constraint)
solve!(p)

# Get the dual value for the constraint
p.constraints[1].dual
# or
constraint.dual
```

1.9.2 Warmstarting

If you're solving the same problem many times with different values of a parameter, Convex.jl can initialize many solvers with the solution to the previous problem, which sometimes speeds up the solution time. This is called a **warm start**.

To use this feature, pass the optional argument *warmstart=true* to the *solve!* method.

```
# initialize data
n = 1000
y = rand(n)
x = Variable(n)

# first solve
lambda = 100
problem = minimize(sumsquares(y - x) + lambda * sumsquares(x - 10))
@time solve!(problem)

# now warmstart
# if the solver takes advantage of warmstarts,
# this run will be faster
lambda = 105
@time solve!(problem, warmstart=true)
```

1.9.3 Fixing and freeing variables

Convex.jl allows you to fix a variable x to a value by calling the *fix!* method. Fixing the variable essentially turns it into a constant. Fixed variables are sometimes also called parameters.

fix(x , v) fixes the variable x to the value v .

fix(x) fixes x to the value $x.value$, which might be the value obtained by solving another problem involving the variable x .

To allow the variable x to vary again, call *free!*(x).

Fixing and freeing variables can be particularly useful as a tool for performing alternating minimization on nonconvex problems. For example, we can find an approximate solution to a nonnegative matrix factorization problem with alternating minimization as follows. We use warmstarts to speed up the solution.

```
# initialize nonconvex problem
n, k = 10, 1
A = rand(n, k) * rand(k, n)
x = Variable(n, k)
y = Variable(k, n)
```

```
problem = minimize(sum_squares(A - x*y), x>=0, y>=0)

# initialize value of y
y.value = rand(k, n)
# we'll do 10 iterations of alternating minimization
for i=1:10
    # first solve for x
    # with y fixed, the problem is convex
    fix!(y)
    solve!(problem, warmstart = i > 1 ? true : false)
    free!(y)

    # now solve for y with x fixed at the previous solution
    fix!(x)
    solve!(problem, warmstart = true)
    free!(x)
end
```

1.10 Contributing

We'd welcome contributions to the Convex.jl package. Here are some short instructions on how to get started. If you don't know what you'd like to contribute, you could

- take a look at the current [issues](#) and pick one. (Feature requests are probably the easiest to tackle.)
- add a [usage example](#).

Then submit a pull request (PR). (Let us know if it's a work in progress by putting [WIP] in the name of the PR.)

1.10.1 Adding examples

- Take a look at our existing [usage examples](#) and add another in similar style.
- Submit a PR. (Let us know if it's a work in progress by putting [WIP] in the name of the PR.)
- We'll look it over, fix up anything that doesn't work, and merge it!

1.10.2 Adding atoms

Here are the steps to add a new function or operation (atom) to Convex.jl. Let's say you're adding the new function f .

- Take a look at the [nuclear norm atom](#) for an example of how to construct atoms, and see the [norm atom](#) for an example of an atom that depends on a parameter.
- Copy paste (eg) the nuclear norm file, replace anything saying nuclear norm with the name of the atom f , fill in monotonicity, curvature, etc. Save it in the appropriate subfolder of `src/atoms/`.
- Add as a comment a description of what the atom does and its parameters.
- The most mathematically interesting part is the `conic_form!` function. Following the example in the nuclear norm atom, you'll see that you can just construct the problem whose optimal value is $f(x)$, introducing any auxiliary variables you need, exactly as you would normally in Convex.jl, and then call `cache_conic_form!` on that problem.
- Add a test for the atom so we can verify it works in `test/test_<cone>`, where `<cone>` matches the subfolder of `src/atoms`.

- Submit a PR, including a description of what the atom does and its parameters. (Let us know if it's a work in progress by putting [WIP] in the name of the PR.)
- We'll look it over, fix up anything that doesn't work, and merge it!

1.10.3 Fixing the guts

If you want to do a more major bug fix, you may need to understand how Convex.jl thinks about conic form. To do this, start by reading [the Convex.jl paper](#). Then read the conic form code:

- We define data structures for conic objectives and conic constraints, and simple ways of combining them, in [conic_form.jl](#)
- We convert the internal conic form representation into the [standard form for conic solvers](#) in the function [conic_problem](#).
- We solve problems (that is, pass the standard form of the problem to a solver, and put the solution back into the values of the appropriate variables) in [solution.jl](#).

You're now armed and dangerous. Go ahead and open an issue (or comment on a previous one) if you can't figure something out, or submit a PR if you can figure it out. (Let us know if it's a work in progress by putting [WIP] in the name of the PR.)

PRs that comment the code more thoroughly will also be welcomed.

1.11 Credits

Currently, Convex.jl is developed and maintained by:

- [Jenny Hong](#)
- [Karanveer Mohan](#)
- [Madeleine Udell](#)
- [David Zeng](#)

The Convex.jl developers also thank:

- the [JuliaOpt](#) team: [Iain Dunning](#), [Joey Huchette](#) and [Miles Lubin](#)
- [Stephen Boyd](#), co-author of the book [Convex Optimization](#)
- [Steven Diamond](#), developer of [CVXPY](#) and of a [DCP tutorial website](#) to teach disciplined convex programming.
- [Michael Grant](#), developer of [CVX](#).
- [John Duchi](#) and [Hongseok Namkoong](#) for developing the [representation of power cones in terms of SOCP constraints](#) used in this package.