




Convex optimization layers

Akshay Agrawal, Brandon Amos, Shane Barratt, Stephen Boyd, Steven Diamond, Zico Kolter



Akshay Agrawal


@akshaykagrwal



Convex optimization


$$\begin{aligned} & \text{minimize} && f(x; \theta) \\ & \text{subject to} && g(x; \theta) \leq 0 \\ & && A(\theta)x = b(\theta) \end{aligned}$$

- $x \in \mathbf{R}^n$ is the variable
- $\theta \in \mathbf{R}^m$ is the parameter
- f and g are convex (curve upwards)
- find x that minimizes f while satisfying the constraints




Why convex optimization?

- Convex optimization problems can be solved quickly, reliably, and *exactly*



Why convex optimization?

- Convex optimization problems can be solved quickly, reliably, and *exactly*
- Software libraries like **CVXPY** make convex optimization easy




Why convex optimization?

- Convex optimization problems can be solved quickly, reliably, and *exactly*
- Software libraries like **CVXPY** make convex optimization easy

```
import cvxpy as cp

x = cp.Variable(n)
objective = cp.Minimize(cp.sum_squares(A @ x - b) + cp.pnorm(x, p=1))
constraints = [0 <= x, x <= 1]
prob = cp.Problem(objective, constraints)
result = prob.solve()
```



Why convex optimization?

- Convex optimization problems can be solved quickly, reliably, and *exactly*
- Software libraries like **CVXPY** make convex optimization easy
- **Tons of applications**




Controlling self-driving cars



Landing rockets




Designing airplanes



Until now ...

- Difficult to use convex optimization problems in TensorFlow pipelines
- Parameters θ were chosen and tuned by hand



CVXPY Layers

CVXPY

$$x^*(\theta) = \underset{\text{subject to}}{\operatorname{argmin}} \begin{cases} f(x; \theta) \\ g(x; \theta) \leq 0 \\ A(\theta)x = b(\theta) \end{cases}$$


```
0 import cvxpy as cp
1 import tensorflow as tf
2 from cvxpylayers.tensorflow import CvxpyLayer

3 n, m = 2, 3
4 x = cp.Variable(n)
5 A, b = cp.Parameter((m, n)), cp.Parameter(m)
6 constraints = [x >= 0]
7 objective = cp.Minimize(0.5 * cp.pnorm(A @ x - b, p=1))
8 problem = cp.Problem(objective, constraints)

9 cvxpylayer = CvxpyLayer(problem, parameters=[A, b], variables=[x])

10 A_tf = tf.Variable(tf.random.normal((m, n)))
11 b_tf = tf.Variable(tf.random.normal((m,)))
12 with tf.GradientTape() as tape:
13     solution, = cvxpylayer(A_tf, b_tf)
14     summed_solution = tf.math.reduce_sum(solution)
15 gradA, gradb = tape.gradient(summed_solution, [A_tf, b_tf])
```

```
0 import cvxpy as cp
1 import tensorflow as tf
2 from cvxpylayers.tensorflow import CvxpyLayer

3 n, m = 2, 3
4 x = cp.Variable(n)
5 A, b = cp.Parameter((m, n)), cp.Parameter(m)
6 constraints = [x >= 0]
7 objective = cp.Minimize(0.5 * cp.pnorm(A @ x - b, p=1))
8 problem = cp.Problem(objective, constraints)

9 cvxpylayer = CvxpyLayer(problem, parameters=[A, b], variables=[x])

10 A_tf = tf.Variable(tf.random.normal((m, n)))
11 b_tf = tf.Variable(tf.random.normal((m,)))
12 with tf.GradientTape() as tape:
13     solution, = cvxpylayer(A_tf, b_tf)
14     summed_solution = tf.math.reduce_sum(solution)
15 gradA, gradb = tape.gradient(summed_solution, [A_tf, b_tf])
```

```
0 import cvxpy as cp
1 import tensorflow as tf
2 from cvxpylayers.tensorflow import CvxpyLayer

3 n, m = 2, 3
4 x = cp.Variable(n)
5 A, b = cp.Parameter((m, n)), cp.Parameter(m)
6 constraints = [x >= 0]
7 objective = cp.Minimize(0.5 * cp.pnorm(A @ x - b, p=1))
8 problem = cp.Problem(objective, constraints)

9 cvxpylayer = CvxpyLayer(problem, parameters=[A, b], variables=[x])


10 A_tf = tf.Variable(tf.random.normal((m, n)))
11 b_tf = tf.Variable(tf.random.normal((m,)))
12 with tf.GradientTape() as tape:
13     solution, = cvxpylayer(A_tf, b_tf)
14     summed_solution = tf.math.reduce_sum(solution)
15 gradA, gradb = tape.gradient(summed_solution, [A_tf, b_tf])
```

```
0 import cvxpy as cp
1 import tensorflow as tf
2 from cvxpylayers.tensorflow import CvxpyLayer


3 n, m = 2, 3
4 x = cp.Variable(n)
5 A, b = cp.Parameter((m, n)), cp.Parameter(m)
6 constraints = [x >= 0]
7 objective = cp.Minimize(0.5 * cp.pnorm(A @ x - b, p=1))
8 problem = cp.Problem(objective, constraints)

9 cvxpylayer = CvxpyLayer(problem, parameters=[A, b], variables=[x])


10 A_tf = tf.Variable(tf.random.normal((m, n)))
11 b_tf = tf.Variable(tf.random.normal((m,)))
12 with tf.GradientTape() as tape:
13     solution, = cvxpylayer(A_tf, b_tf)
14     summed_solution = tf.math.reduce_sum(solution)
15 gradA, gradb = tape.gradient(summed_solution, [A_tf, b_tf])
```




Learning to control a car




Learning to control a car




iteration 0



Learning to control a car



iteration 0



iteration 100



For more ...

github.com/cvxgrp/cvxpylayers

NeurIPS paper



Learning control policies (L4DC)



with examples in

- controlling a car
- managing a supply chain
- allocating financial portfolios



Thank you!



TensorFlow

DEV SUMMIT 2020