

# **Real-time Decision-Making via Data-Driven Optimization**



**PRINCETON  
UNIVERSITY**

**ORFE**

**Bartolomeo Stellato – Politecnico di Milano, June 16 2022**

# New Capabilities for Autonomous Systems

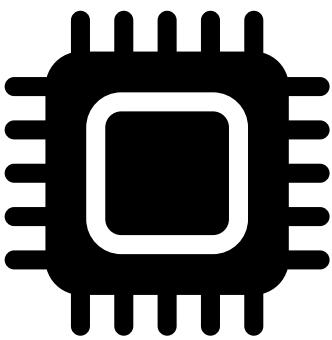
## Automation



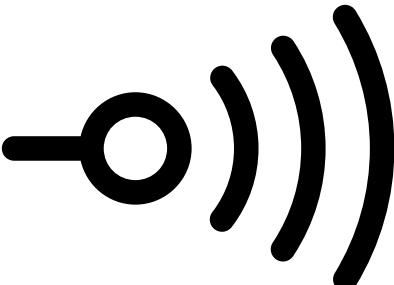
[Courtesy of EnterpriseTalk]

system limited to  
**pre-programmed behavior**

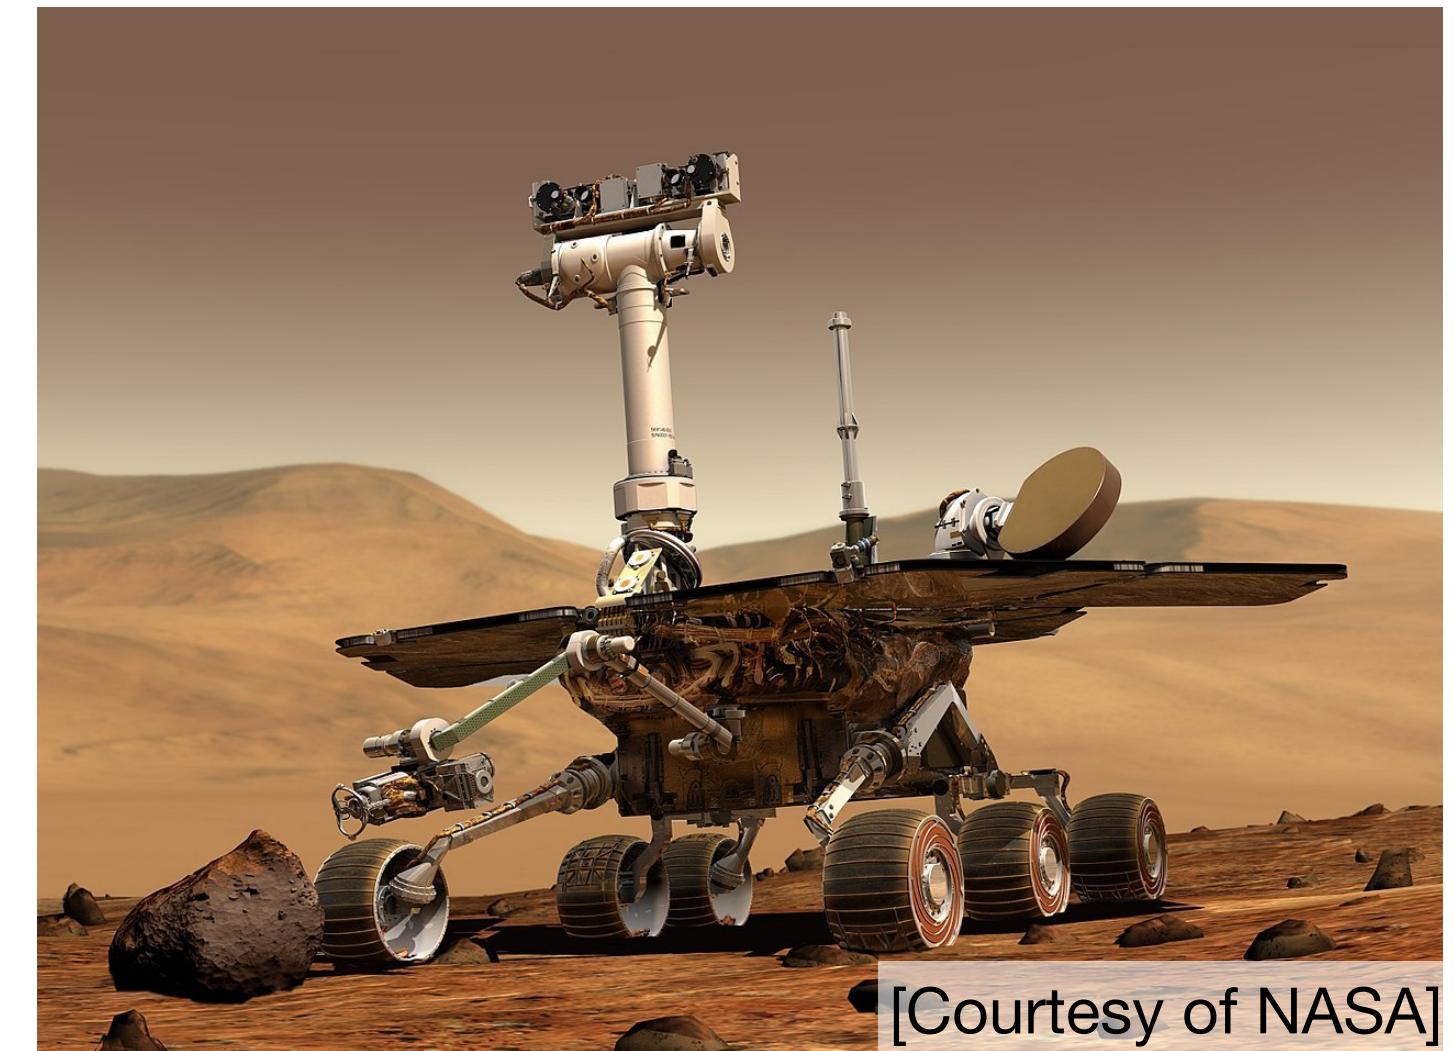
faster  
computing



accurate  
sensing



## Autonomy

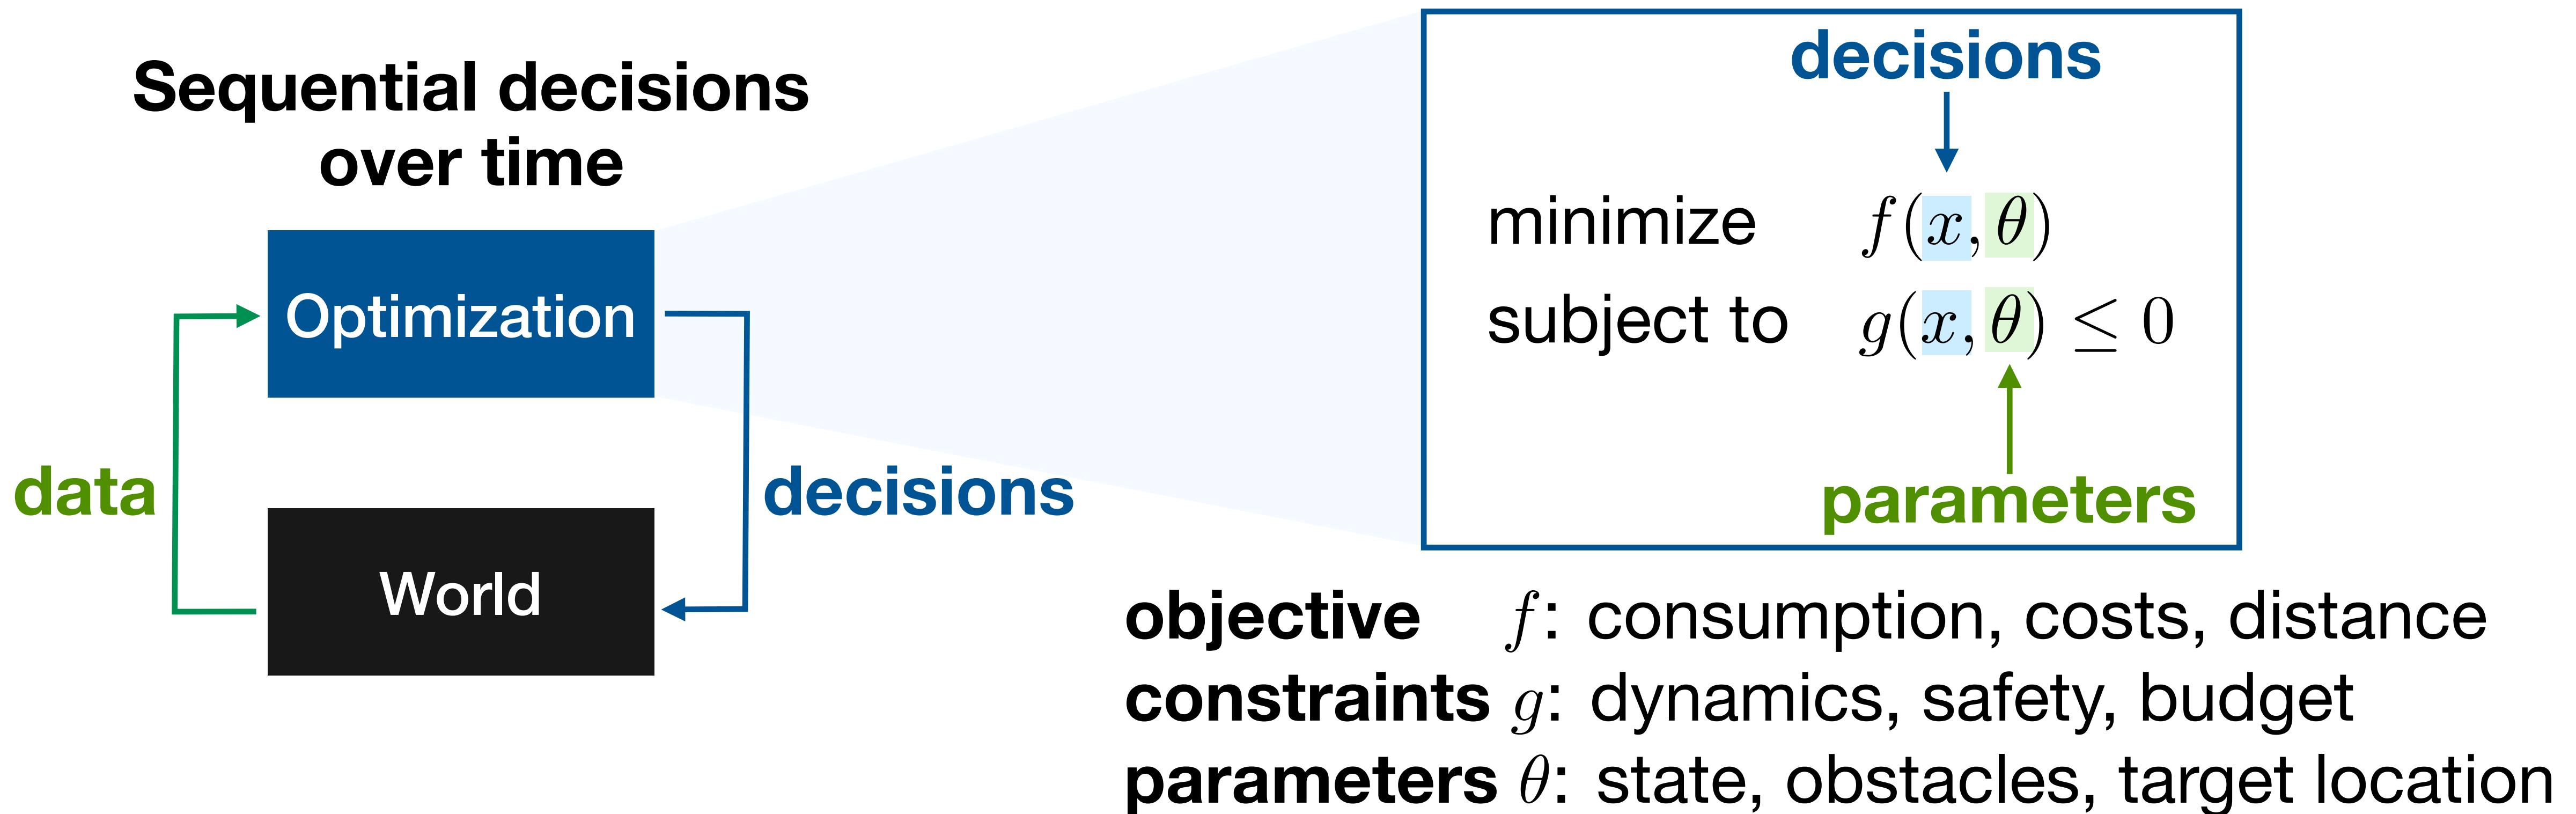


[Courtesy of NASA]

system reacts to  
**new not pre-programmed  
situations**

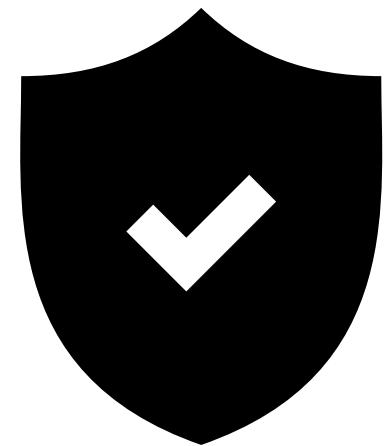
How can an autonomous system **quickly** and **safely** react to  
unexpected conditions?

# Optimization for real-time decision-making

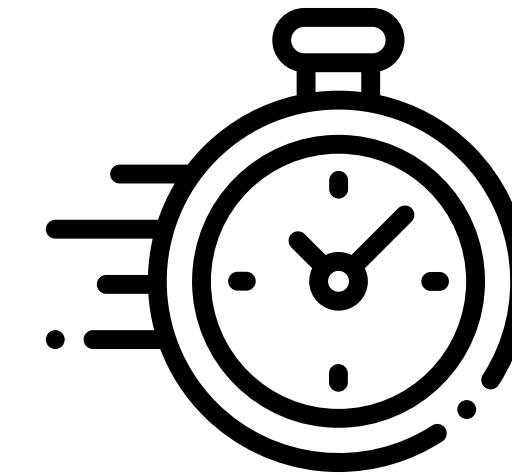


# Challenges in real-time optimization

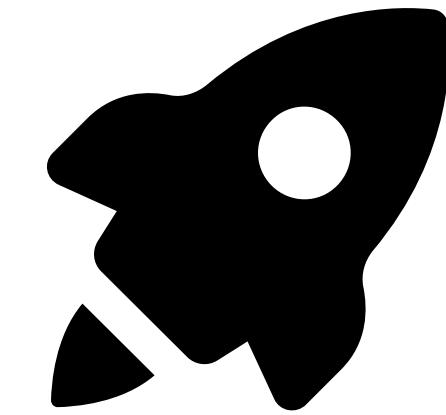
**Extreme reliability**



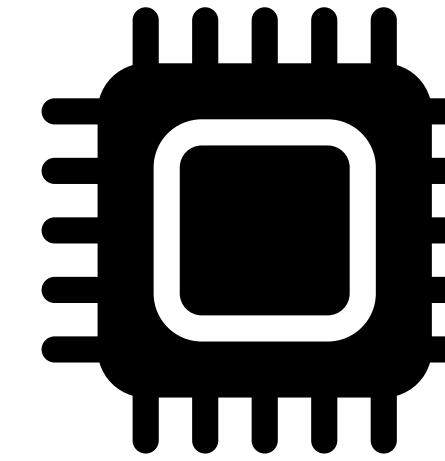
Real-Time



Performance



Limited resources

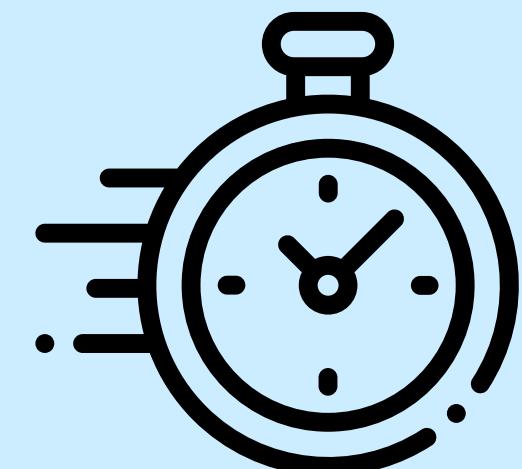


# Today's talk

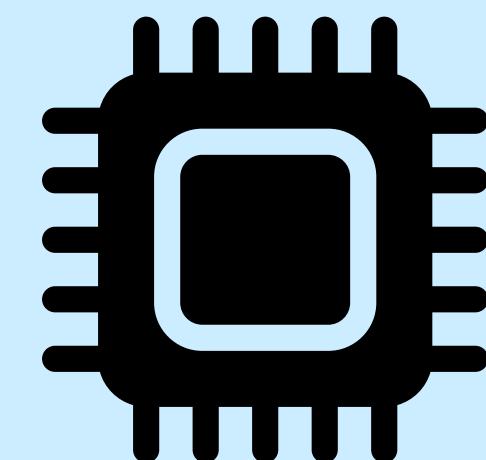
## Real-time Decision-Making via Data-Driven Optimization

**osQP  
Solver**

Real-Time

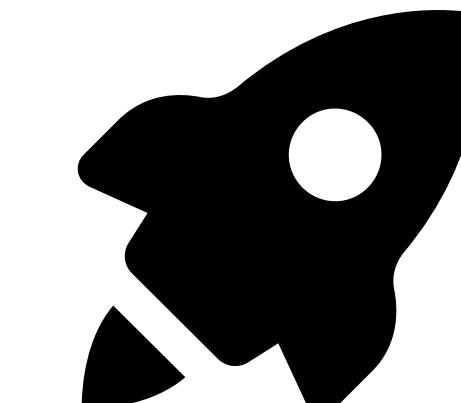


Limited resources



**Learning  
Convex Optimization  
Control Policies**

Performance





# First-order methods

Wide popularity

## Pros

Warm-starting

Large-scale problems

Embeddable

## Cons

Low quality solutions

Can't detect infeasibility

Problem data dependent

## OSQP

High-quality solutions

Detects infeasibility

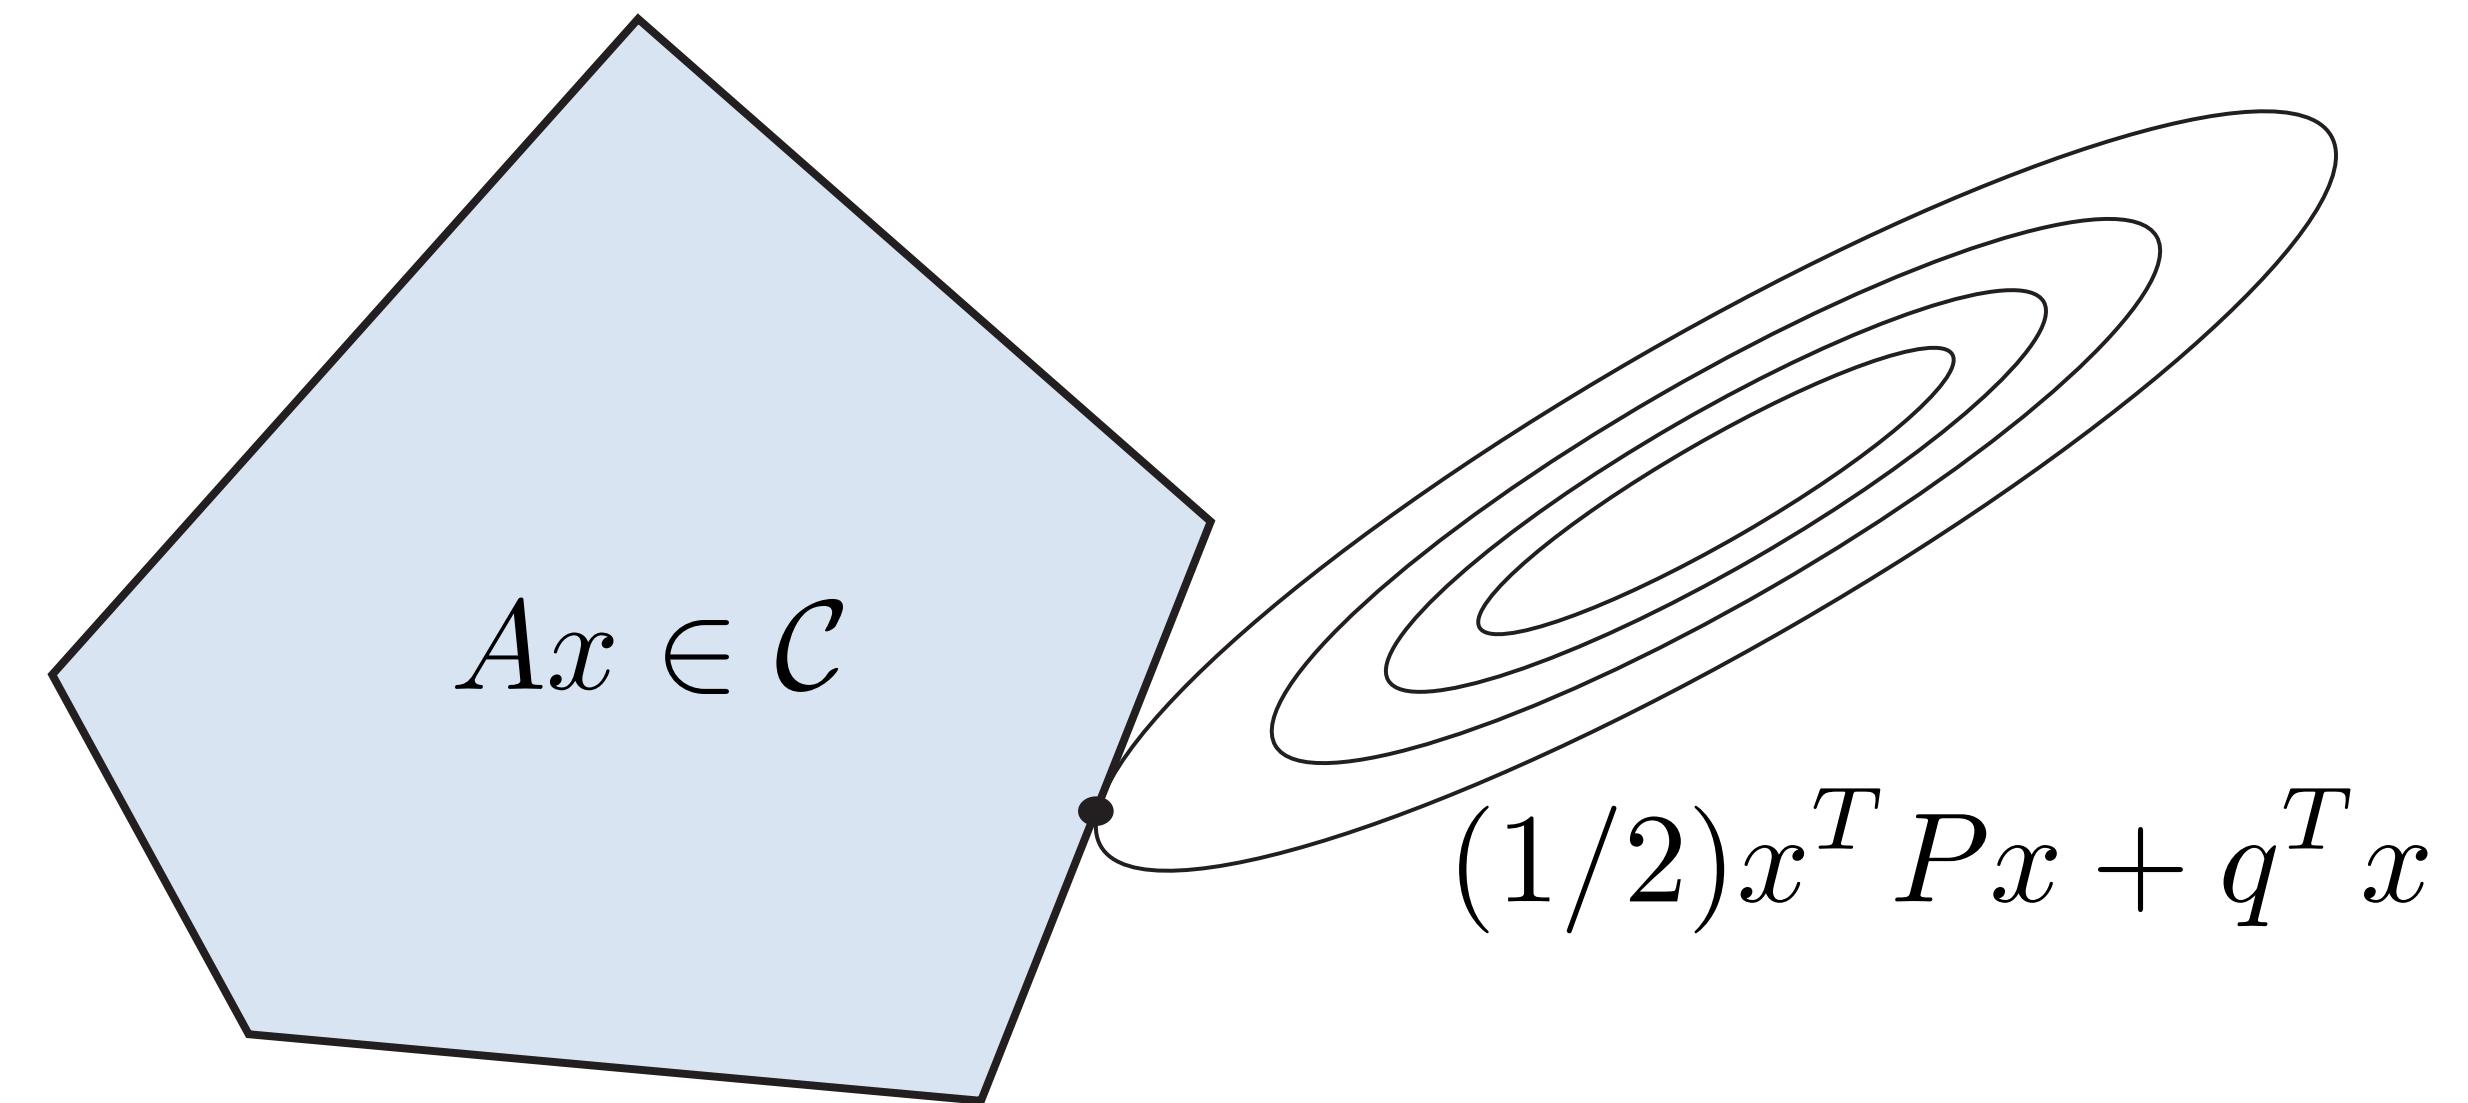
Robust



# The problem

$$\begin{aligned} & \text{minimize} && (1/2)x^T Px + q^T x \\ & \text{subject to} && Ax \in \mathcal{C} \end{aligned}$$

Quadratic program:  $\mathcal{C} = [l, u]$



# ADMM

## Alternating Direction Method of Multipliers

$$\text{minimize } f(x) + g(x)$$



**Splitting**

$$\begin{aligned} & \text{minimize} && f(\tilde{x}) + g(x) \\ & \text{subject to} && \tilde{x} = x \end{aligned}$$

## Iterations

$$\tilde{x}^{k+1} \leftarrow \operatorname{argmin}_{\tilde{x}} \left( f(\tilde{x}) + \rho/2 \left\| \tilde{x} - (x^k - y^k/\rho) \right\|^2 \right)$$

$$x^{k+1} \leftarrow \operatorname{argmin}_x \left( g(x) + \rho/2 \left\| x - (\tilde{x}^{k+1} + y^k/\rho) \right\|^2 \right)$$

$$y^{k+1} \leftarrow y^k + \rho (\tilde{x}^{k+1} - x^{k+1})$$

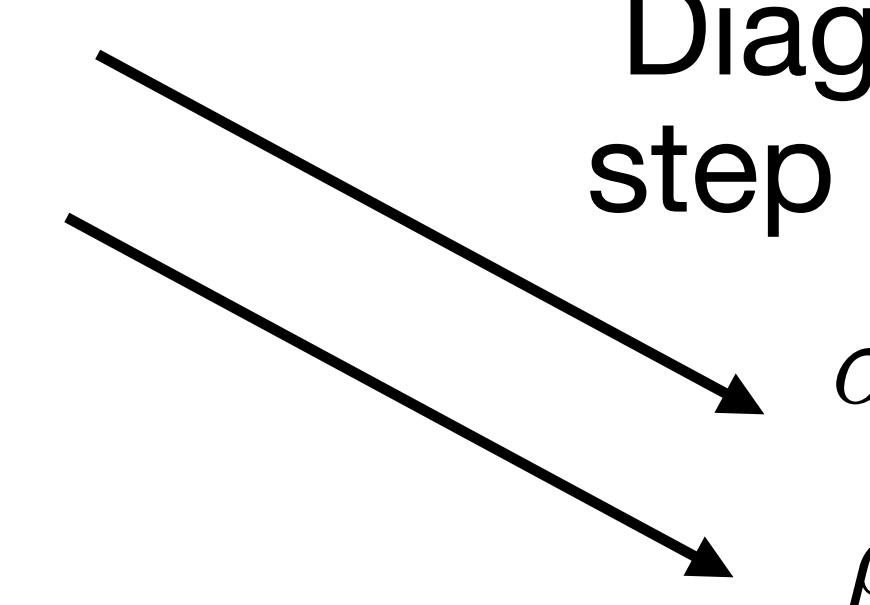
# How do we split the QP?

minimize  $(1/2)x^T Px + q^T x$   $f$   
subject to  $Ax = z$   
 $z \in \mathcal{C}$   $g$

## Splitting formulation

minimize  $(1/2)\tilde{x}^T P\tilde{x} + q^T \tilde{x} + \mathcal{I}_{Ax=z}(\tilde{x}, \tilde{z}) + \mathcal{I}_{\mathcal{C}}(z)$   $f$   $g$   
subject to  $\tilde{x} = x$   
 $\tilde{z} = z$

Diagonal  
step sizes


$$\sigma$$
$$\rho$$

# Complete algorithm

## Problem

$$\begin{aligned} & \text{minimize} && (1/2)x^T Px + q^T x \\ & \text{subject to} && l \leq Ax \leq u \end{aligned}$$

## Algorithm

Linear system  
solve

Easy  
operations

$$x^{k+1} \leftarrow \text{Solve } (P + \sigma I + \rho A^T A)x = \sigma x^k - q + A^T(\rho z^k - y^k)$$

$$z^{k+1} \leftarrow \Pi(Ax^{k+1} + \rho^{-1}y^k)$$

$$y^{k+1} \leftarrow y^k + \rho(Ax^{k+1} - z^{k+1})$$

always solvable!

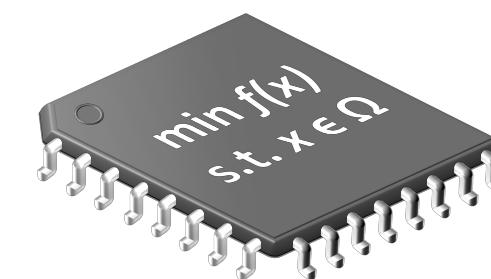
# Code generation with OSQP

Optimized C code

```
# Create OSQP object  
m = osqp.OSQP()  
  
# Initialize solver  
m.setup(P, q, A, l, u,  
        settings)  
  
# Generate C code  
m.codegen('folder_name')
```

```
/ Main ADMM algorithm  
for (iter = 1; iter <= work->settings->max_iter; iter++) {  
    /* U */  
    swap /* Main ADMM algorithm */  
    swap /* Update x_prev, z_prev (preallocated, no malloc) */  
    /* A */  
    swap /* Main ADMM algorithm */  
    swap /* Update x_prev, z_prev (preallocated, no malloc) */  
    /* C */  
    update /* Main ADMM algorithm */  
    /* C */  
    update /* Compute |tilde(x)|^{k+1}, |tilde(z)|^{k+1} */  
    update_xz_tilde(work);  
    /* C */  
    update /* Compute x^{k+1} */  
    update_x(work);  
    /* C */  
    update /* Compute z^{k+1} */  
    update_z(work);  
    /* C */  
    update /* Compute y^{k+1} */  
    update_y(work);  
    /* End of ADMM Steps */  
    #ifdef CTRLC  
    /* Check the interrupt signal */  
    if (isInterrupted()) {  
        update_status(work->info, OSQP_SIGINT);  
        c_print("Solver interrupted\n");  
        endInterruptListener();  
        return 1; // exitflag  
    }  
    #endif
```

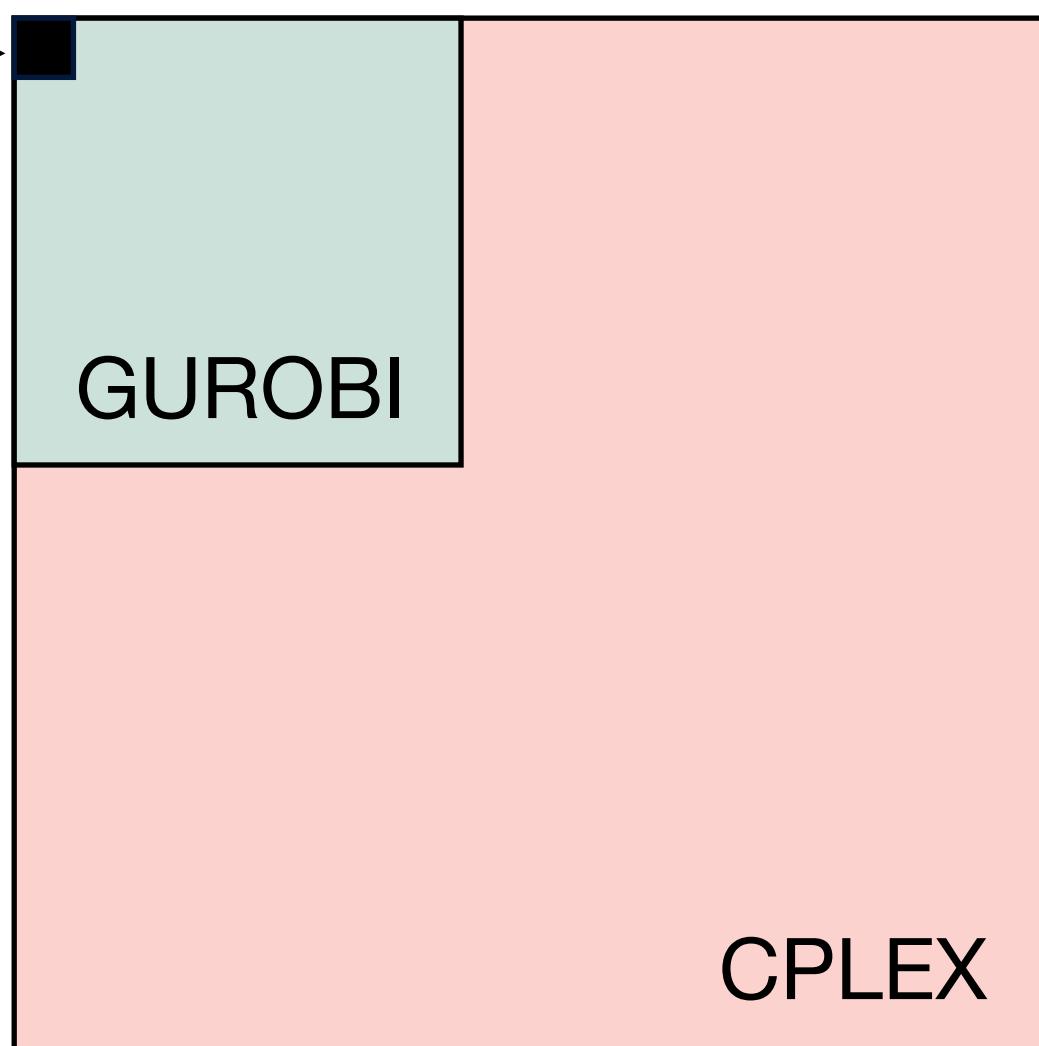
Embedded  
Hardware



It can be  
compiled into  
division-free!

Compiled code size ~80kb (low footprint)

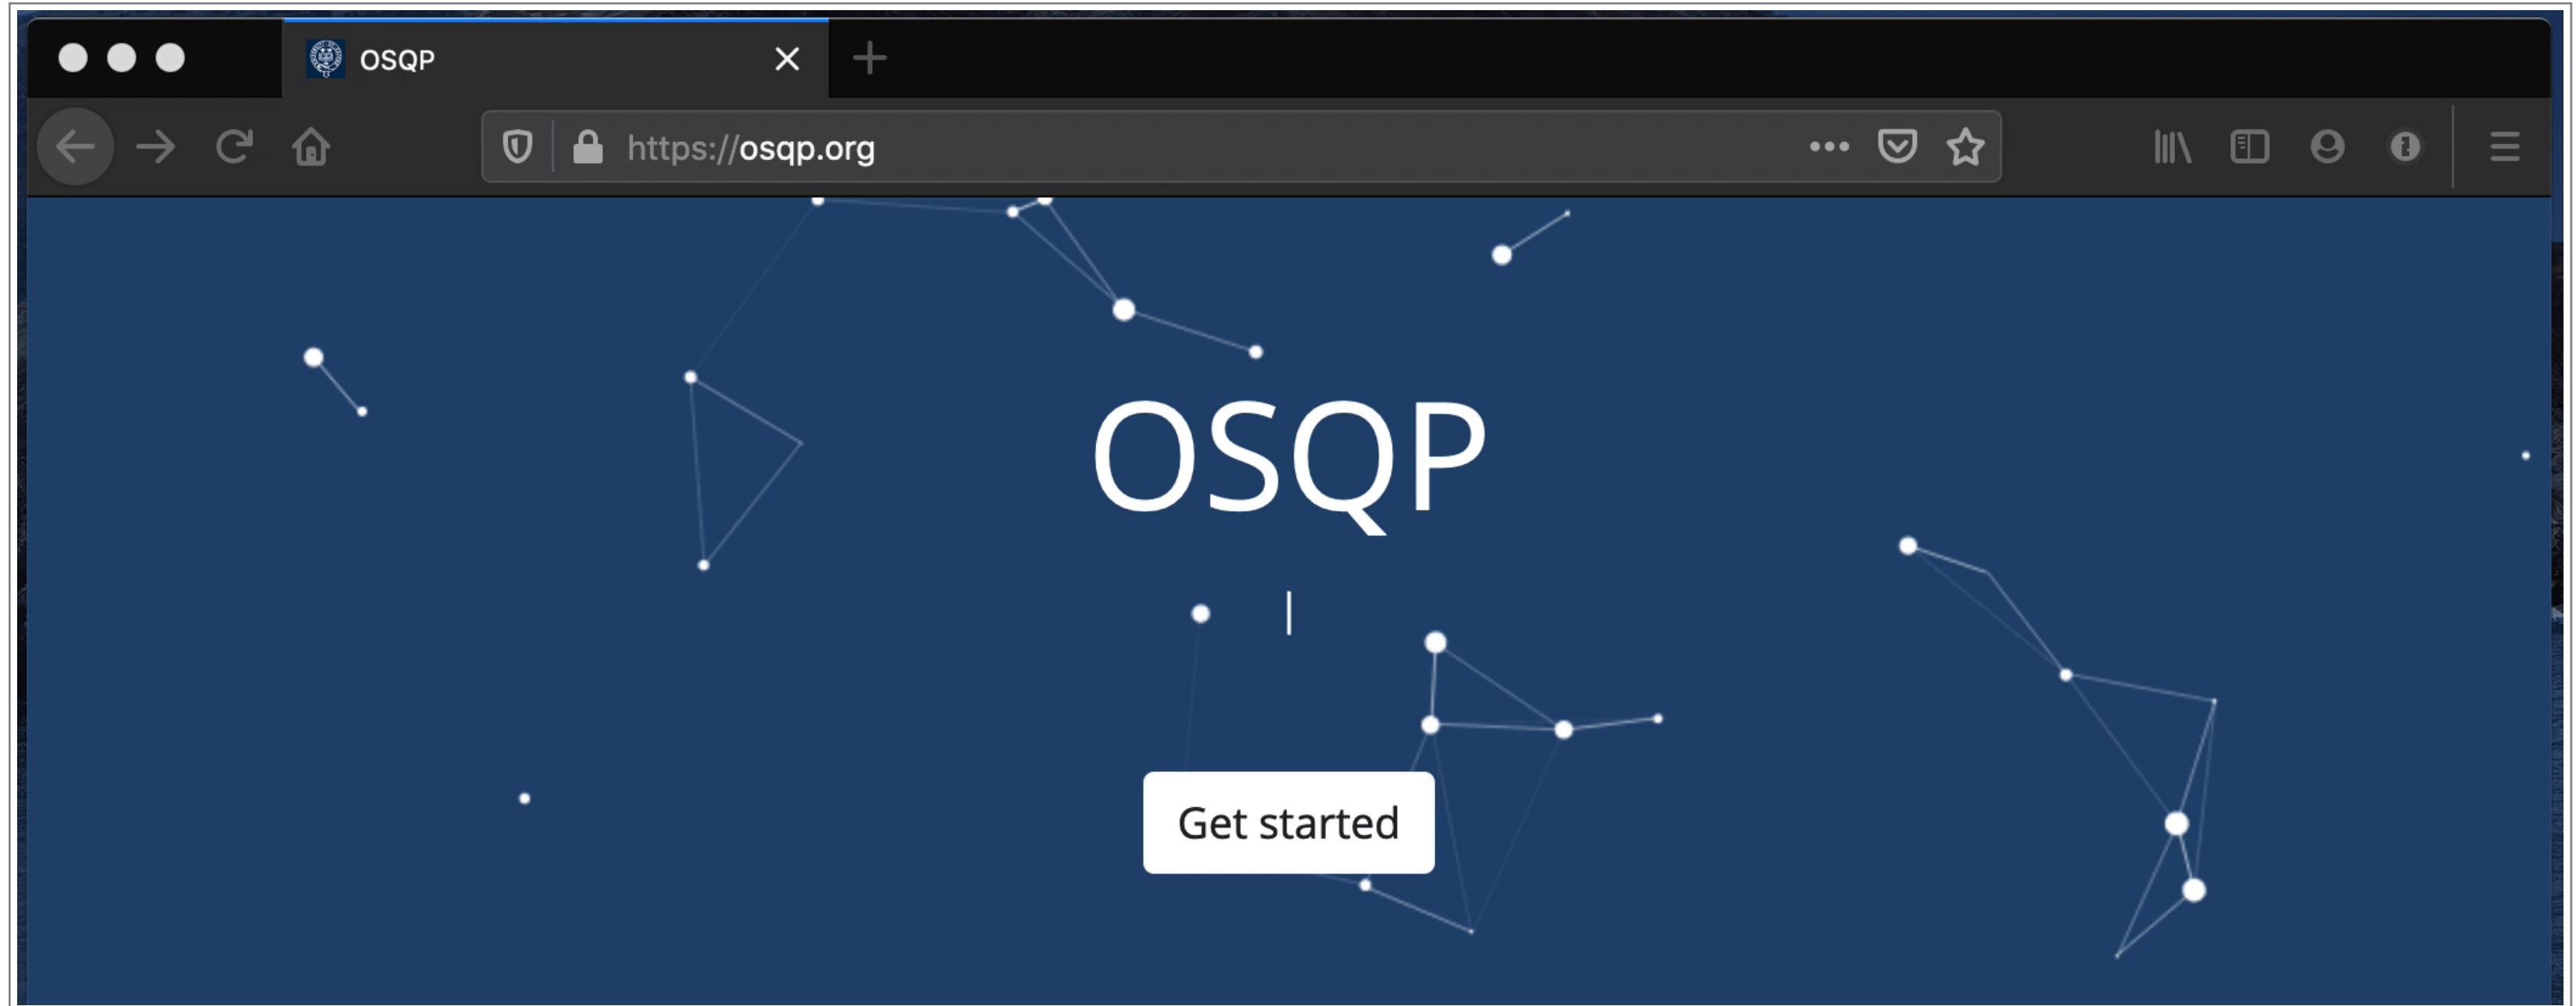
OSQP



300x  
Reduction!

# OSQP

**Operator Splitting solver for Quadratic Programs**



**Embeddable  
(can be division free!)**

**Supports  
warm-starting**

**Detects  
infeasibility**

**Solves large-scale  
problems**

# Users

More than 16 million downloads!



• A P T I V •



# How do we ensure fast convergence?

$$\begin{array}{ll}\text{minimize} & (1/2)x^T Px + q^T x \\ \text{subject to} & Ax = z \\ & l \leq z \leq u\end{array}$$

**Primal residual**

$$r_{\text{prim}}^k = Ax^k - z^k$$

**Dual residual**

$$r_{\text{dual}}^k = Px^k + q + A^T y^k$$

**Intuition**

linear system solution with  $\sigma = 0$

$$(P + \rho A^T A) x^{k+1} = -q + A^T (\rho z^k - y^k)$$

$$\rho = \infty$$

$$A^T A x^{k+1} = A^T z^k$$

$$\rho = 0$$

$$Px^{k+1} = -q - A^T y^k$$

Small primal residual

Small dual residual

What's the optimal  $\rho$ ?

# Constraint-wise step size

$$\text{minimize} \quad (1/2)x^T Px + q^T x$$

$$\text{subject to} \quad l \leq Ax \leq u$$

**Tight constraints**

$$l_i = (Ax^*)_i \text{ or } (Ax^*)_i = u_i$$

**Diagonal step size**

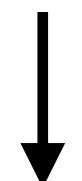
$$\rho = (\rho_1, \dots, \rho_m)$$

**Never tight**

$$l_i = -\infty \text{ and } u_i = \infty$$

$$\downarrow$$
  
$$\rho_i = 0$$

**Otherwise**



**Balance residuals**

$$\rho_i^{k+1} \leftarrow \rho_i^k \sqrt{\|r_{\text{prim}}\| / \|r_{\text{dual}}\|}$$

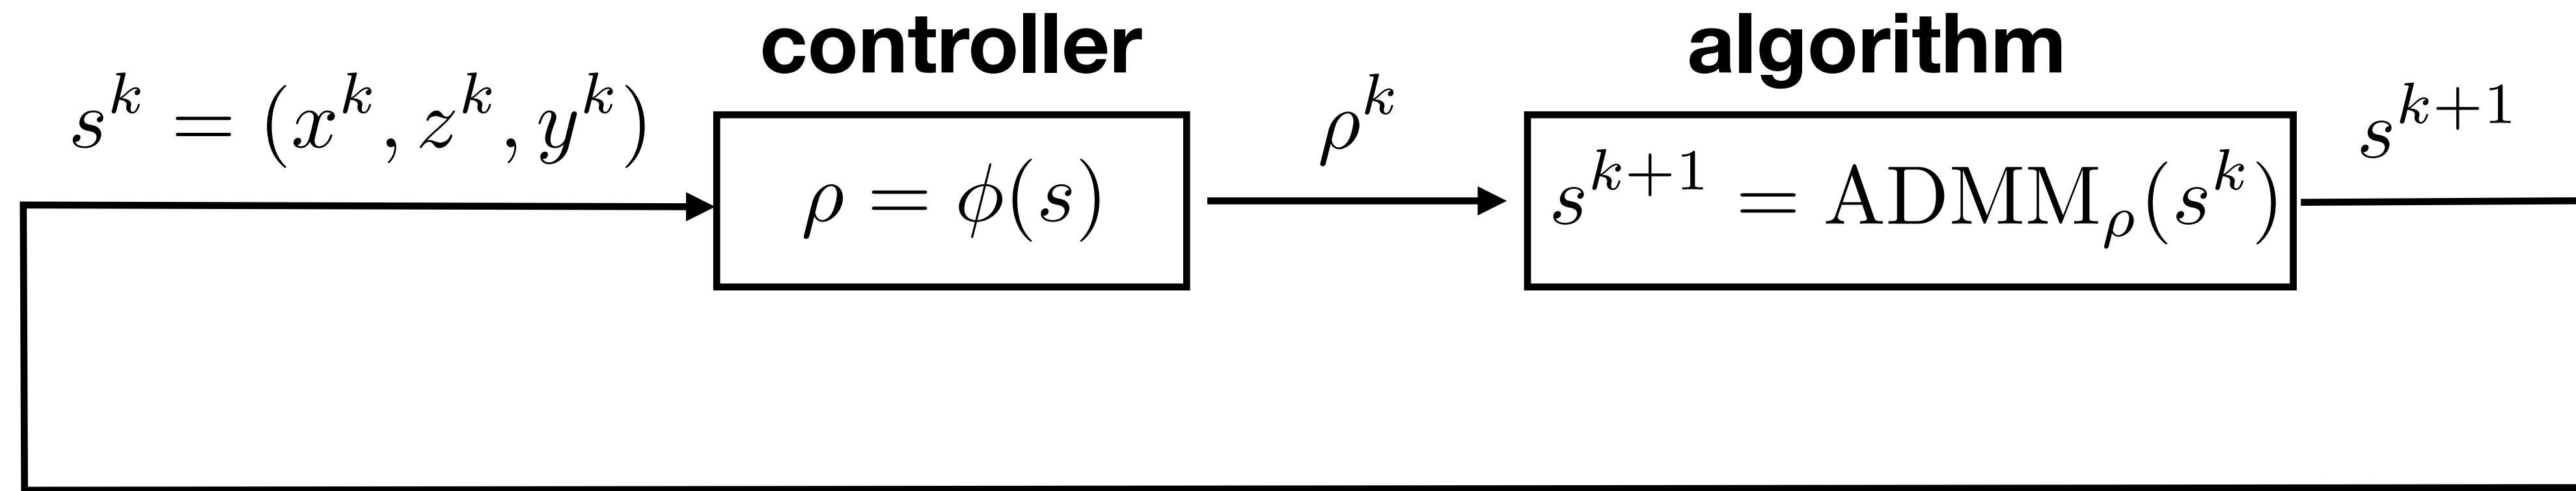
**Always tight**

$$l_i = u_i \neq \infty$$

$$\downarrow$$
  
$$\rho_i = \infty$$

Can we learn a  
better update rule  
from data?

# Step size choice as a control problem



**Stage cost**

$$\ell(s) = \begin{cases} 1 & \text{if not converged} \\ 0 & \text{if converged} \end{cases}$$

**Cumulative cost**

$$J = \mathbf{E} \sum_{k=1}^{\infty} \gamma^k \ell(s^k)$$

**Train with  
Deep Policy Gradient methods (TD3)**

# Constraint-wise control policy

**Per-constraint update rule**

$$\rho_i = \phi_c(s_i)$$

**Per-constraint state**

$$s_i = \begin{bmatrix} \min(z_i - l_i, u_i - z_i) \\ (Ax)_i - z_i \\ y_i \end{bmatrix}$$

**slack**  
**infeasibility**  
**dual variable**

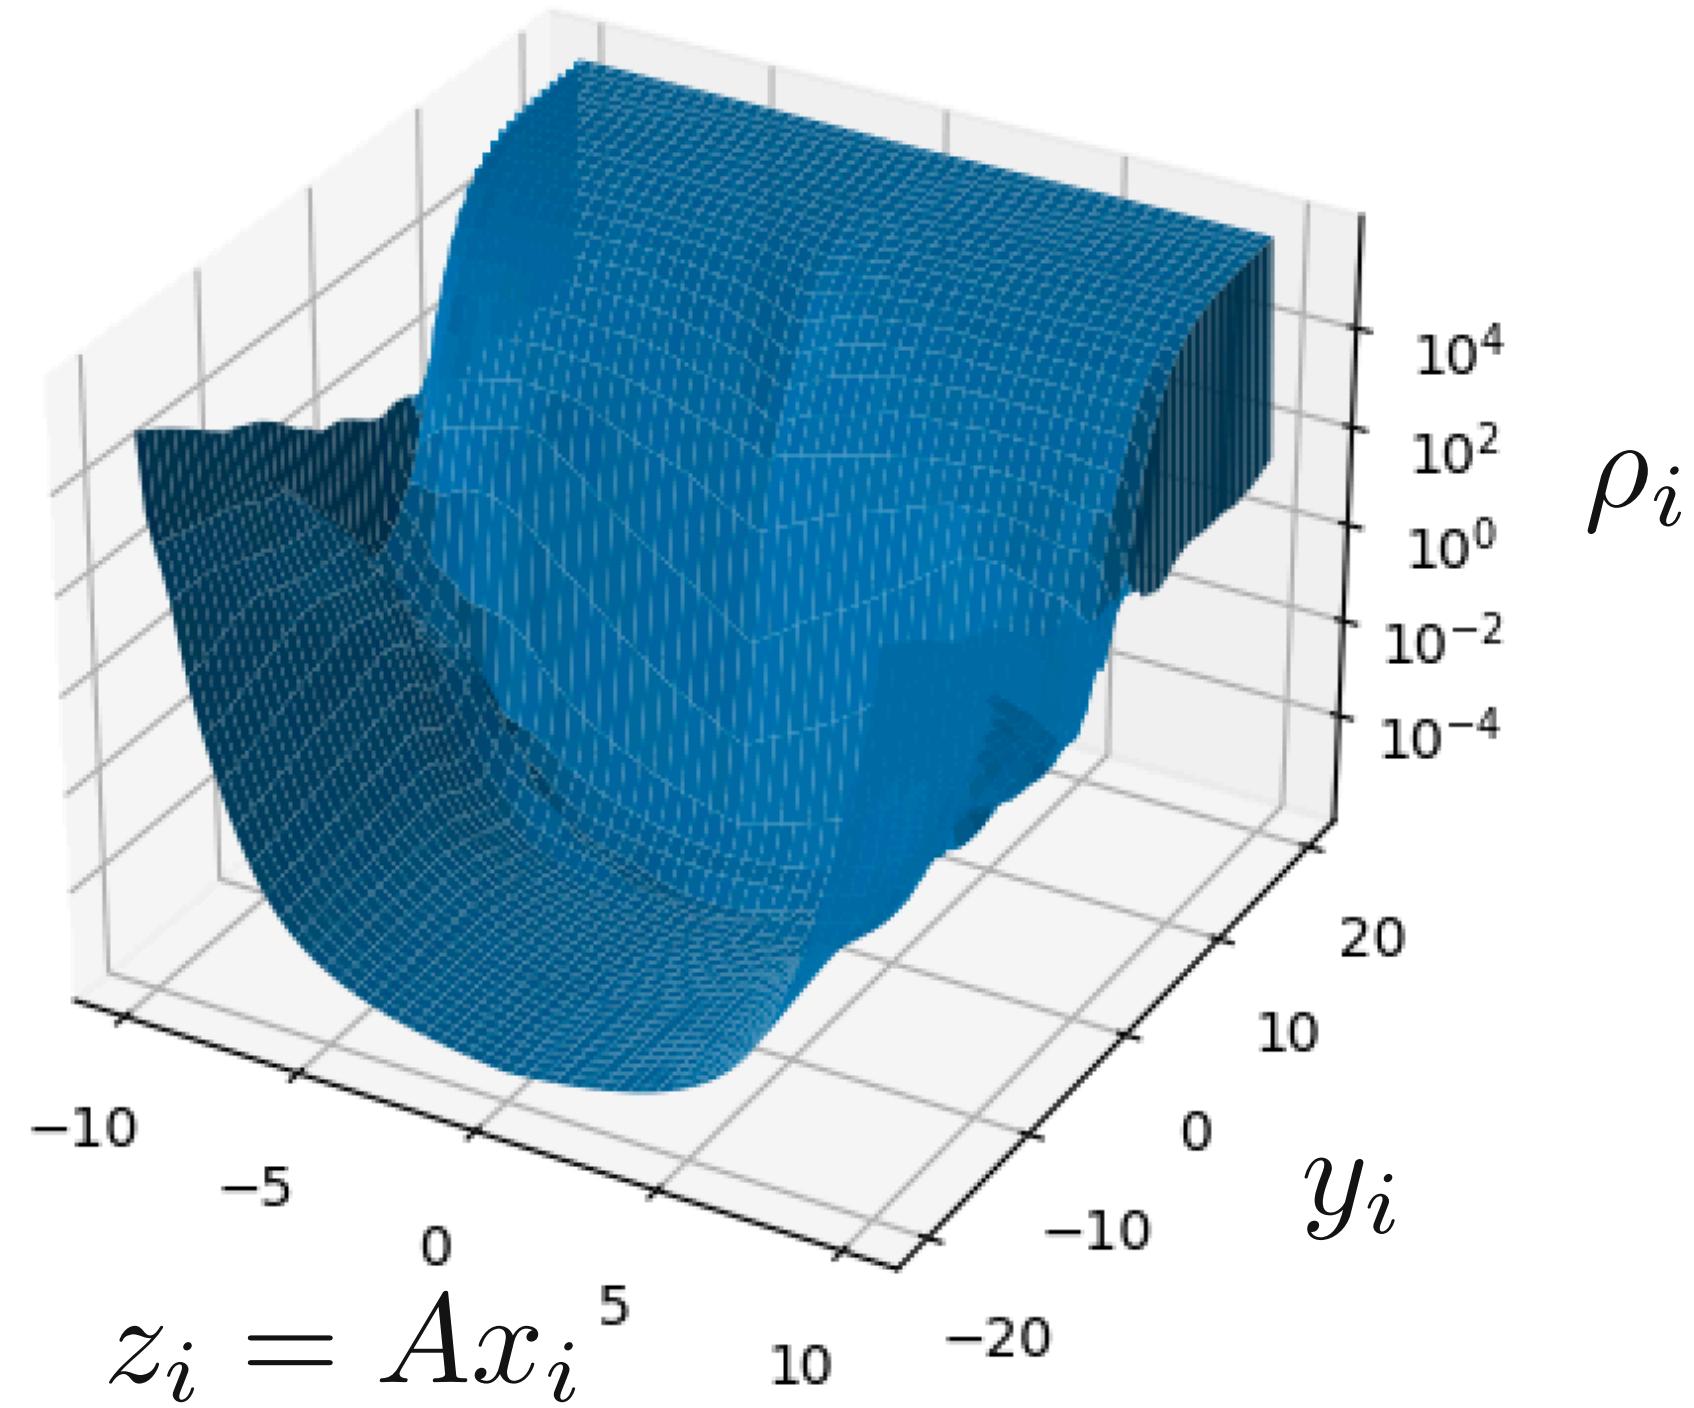
$$\phi(s) = \begin{bmatrix} \phi_c(s_1) \\ \phi_c(s_2) \\ \vdots \\ \phi_c(s_m) \end{bmatrix}$$

Generalize to  
different  
dimensions

Low-dimensional  
state per  
constraint

Small NN  
policy  
 $\phi_c(s_i)$

# Visualize learned policy

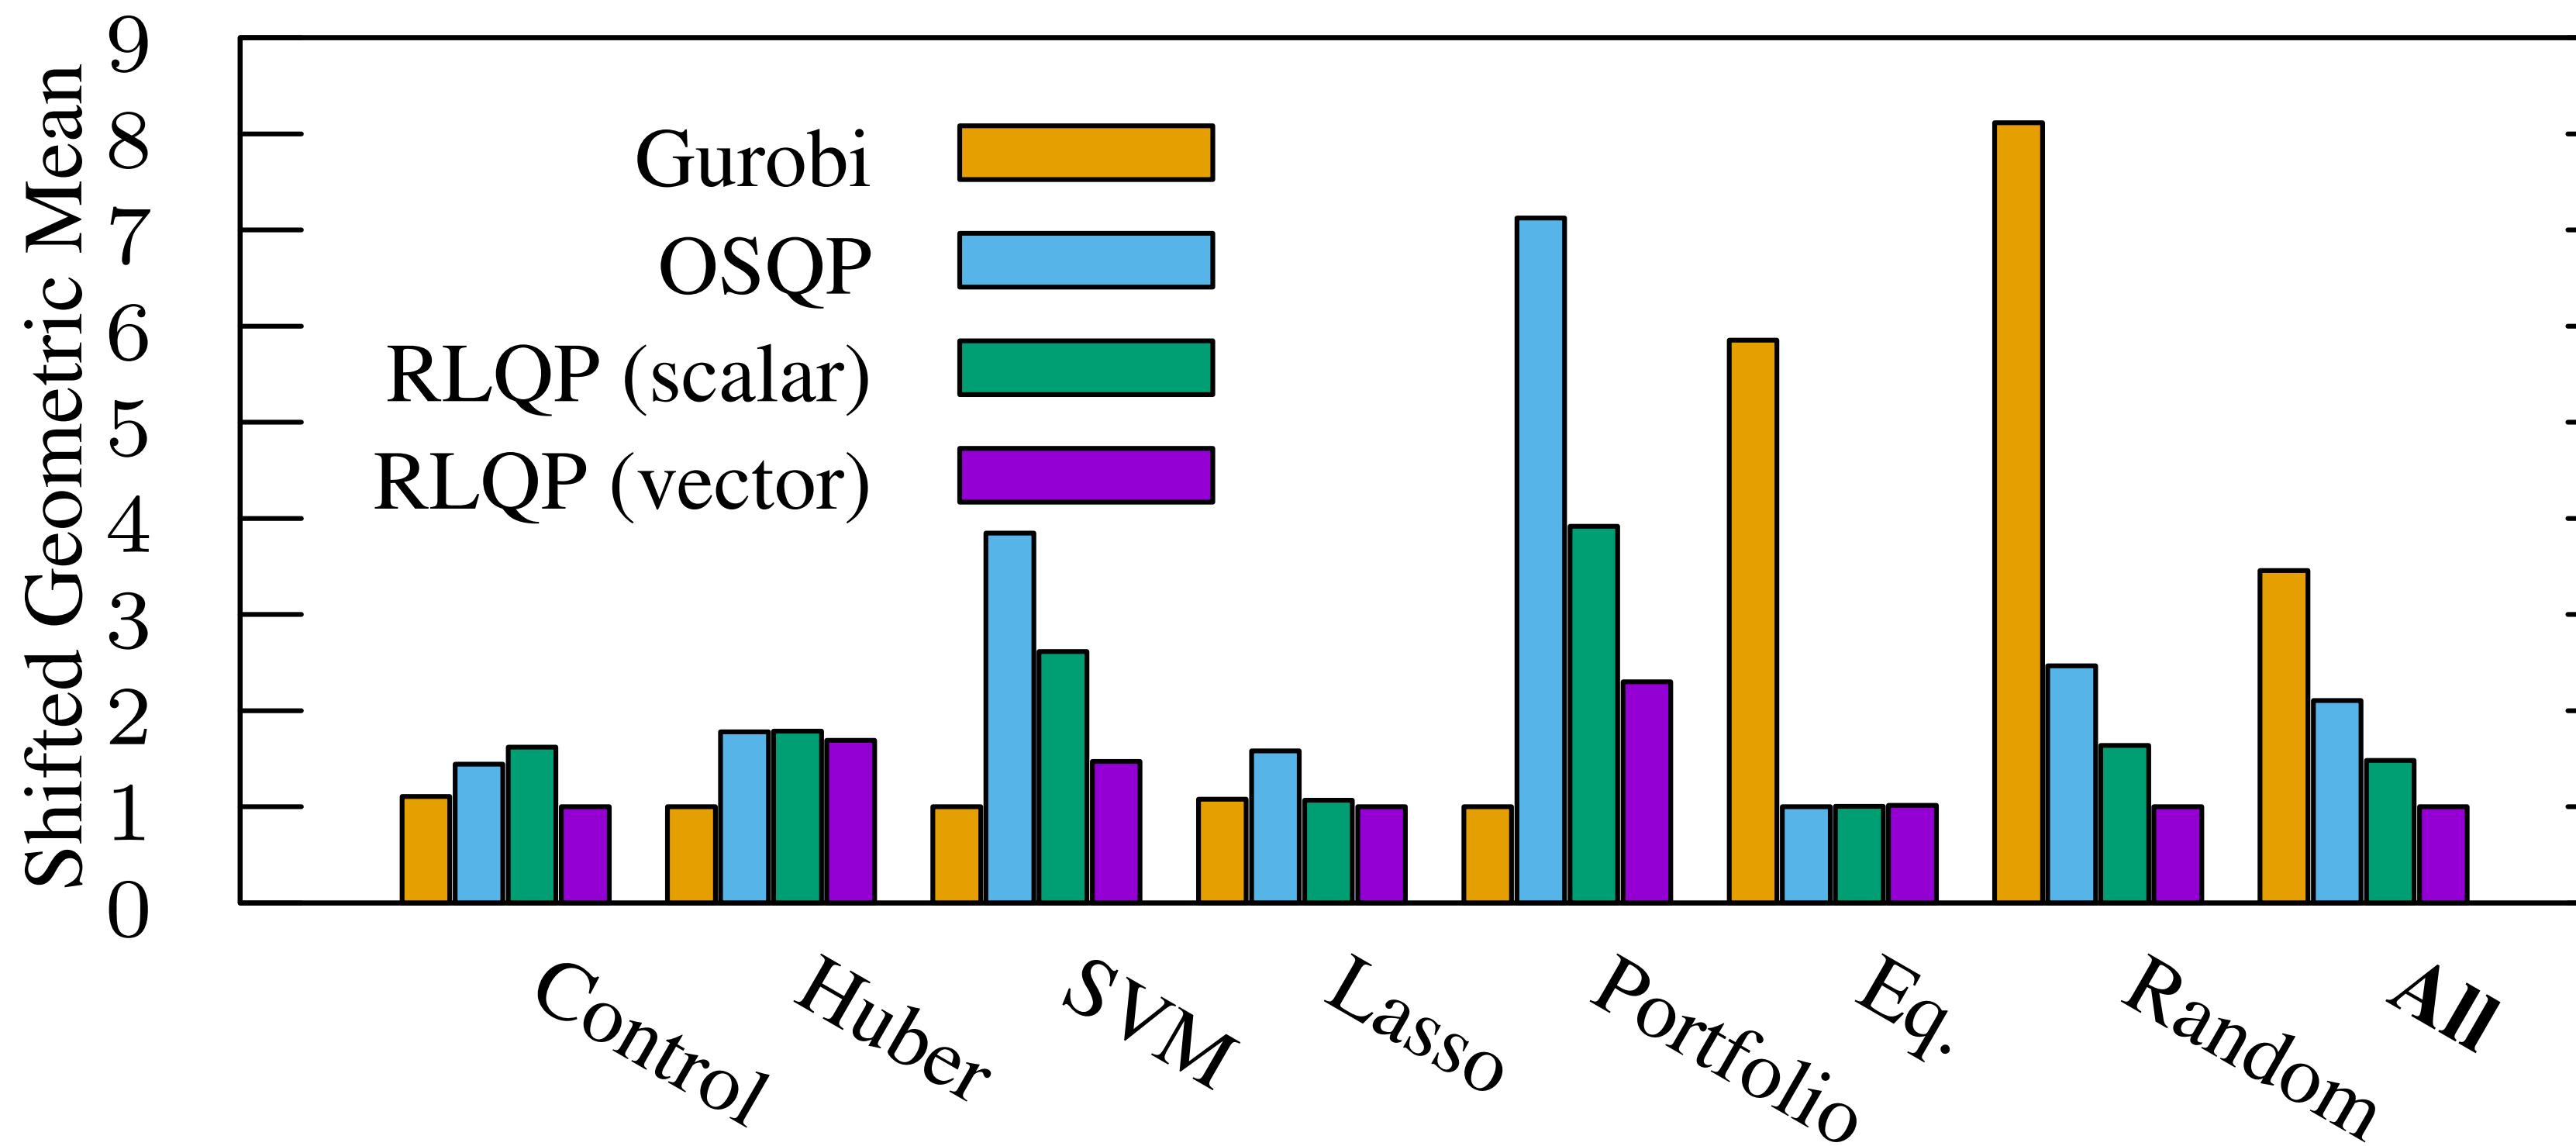


Interpretable policy

**High step size**  $\rho_i$   
when we reach bounds  
 $z_i \approx l_i$  and  $z_i \approx u_i$ .

# Performance with step size learning

Timings for high-accuracy convergence criteria



Up to 3x faster  
than Gurobi

# OSQP 1.0 (this summer!)

# Improved embedded code generation

```
# Create OSQP object
m = osqp.OSQP()

# Initialize solver
m.setup(P, q, A, l, u,
        settings)

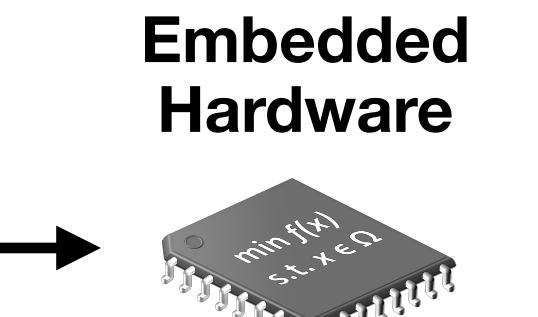
# Generate C code
m.codegen('folder_name')
```

```

Main ADMM algorithm
for (iter = 1; iter <= work->settings->max_iter; iter++) {
    // Main ADMM algorithm
    swap_vectors(&(work->x), &(work->z));
    swap_vectors(&(work->z), &(work->x));
    // Update x_prev, z_prev (preallocated, no malloc)
    swap_vectors(&(work->x), &(work->x_prev));
    swap_vectors(&(work->z), &(work->z_prev));
    update_xz_tilde(work);
    /* ADMM STEPS */
    /* Compute |tilde{x}|^{k+1}, |tilde{z}|^{k+1} */
    update_xz_tilde(work);
    /* Compute x^{k+1} */
    update_x(work);
    /* Compute z^{k+1} */
    update_z(work);
    /* Compute y^{k+1} */
    update_y(work);
    /* End of ADMM Steps */
}

#endif
// C
if (0) {
    /* Check the interrupt signal
    if (isInterrupted()) {
        update_status(work->info, OSQP_SIGINT);
        c_pprint("Solver interrupted\n");
        endInterruptListener();
        return 1; // exitflag
    }
#endif
}

```



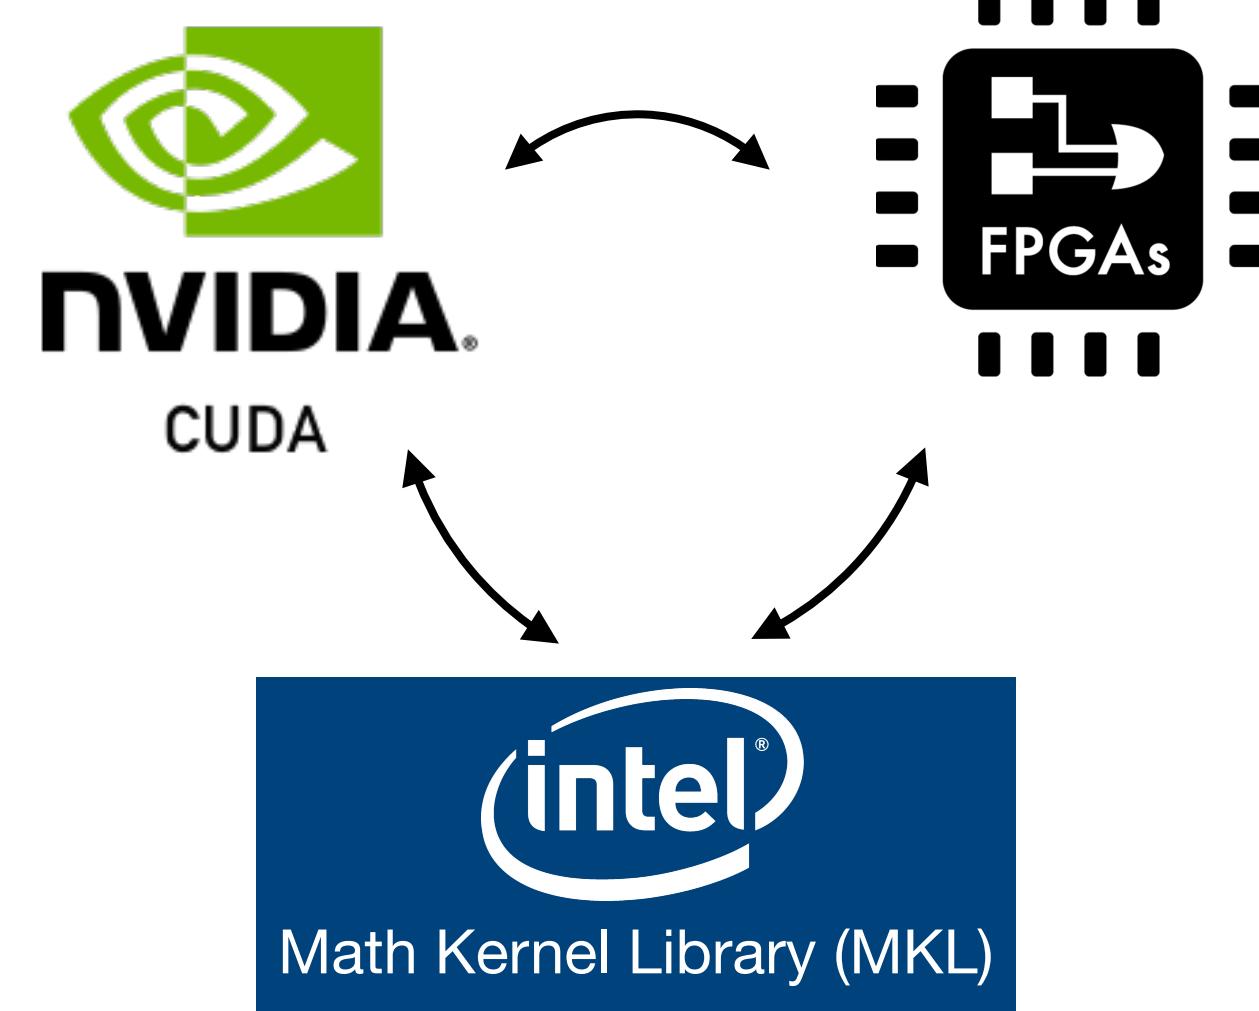
# Code generation from C to C

# Differentiable layers



# PyTorch

# Modular linear algebra



# OSQP features

## Features

Robust

Embeddable  
(can be division free!)

Detects  
infeasibility

Supports  
warm-starting

## Learning for optimization

Integrate RL and ADMM to dynamically  
tune parameters

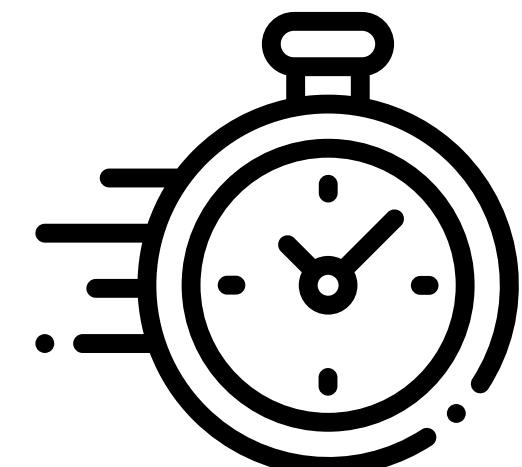
- Faster convergence
- Very low overhead
- Interpretable policy

# Today's talk

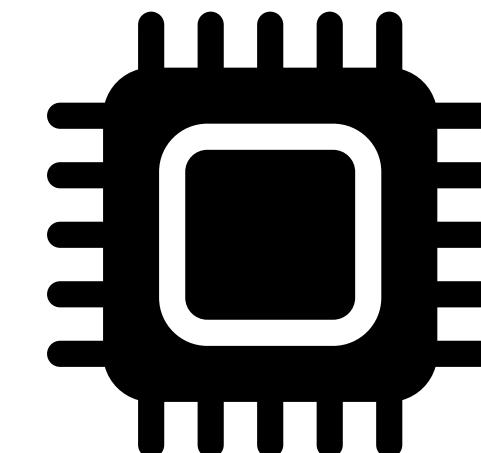
## Real-time Decision-Making via Data-Driven Optimization

**osQP  
Solver**

Real-Time

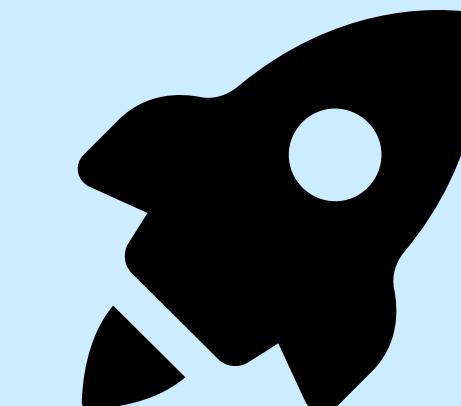


Limited resources

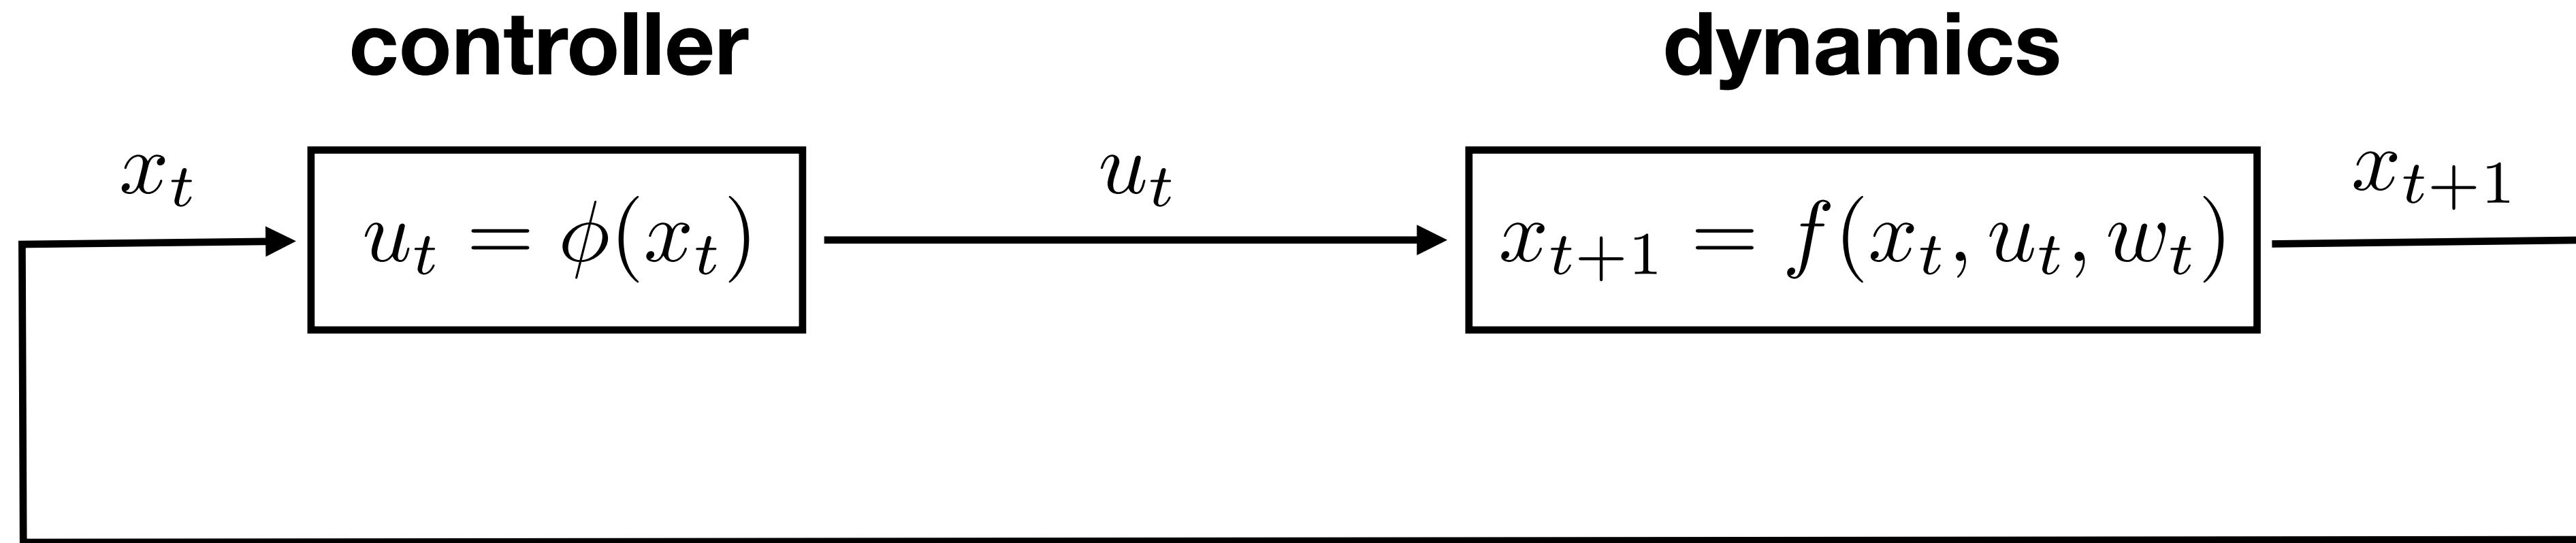


**Learning  
Convex Optimization  
Control Policies**

Performance



# Control loop



$x_t$  state  
 $u_t$  input  
 $w_t$  (random) disturbance

$\phi(x_t)$  control policy

# Explicit vs implicit control policies

## Explicit

Complete control specification

## Implicit (optimization-based)

Designer specifies  
goal and  
requirements



**Optimizer**  
computes  
the action

### Example: PI Controller

$$u_t = -K_P e_t - K_I \sum_{\tau=0}^t e_{\tau}$$

### Example: LQR Controller

dynamics:  $x_{k+1} = Ax_t + Bu_t + w_t$

stage cost:  $x^T Q x + u^T R u$



$$\begin{aligned} u_t &= \underset{u}{\operatorname{argmin}} u^T R u + (Ax_t + Bu)^T P (Ax_t + Bu) \\ &= K x_t \end{aligned}$$

# Convex optimization control policies (COCPs)

$$\begin{aligned} u_t = \operatorname{argmin}_u \quad & f(x_t, u, \theta) \\ \text{subject to} \quad & g(x_t, u, \theta) \leq 0 \\ & A(x_t, \theta)u = b(x_t, \theta) \end{aligned}$$

$x_t$  state  
 $\theta$  parameters to tune  
 $f, g$  convex functions

# Many control policies are COCPs

## Examples

- Linear Quadratic Regulator (LQR)
- Model predictive control (MPC)
- Actuator allocation
- Resource allocation
- Portfolio trading

## Advantages

Interpretable

Satisfy  
constraints

Handle varying  
dynamics

Efficient and reliable  
(even division-free: OSQP)

# Judging COCPs

Given a policy, state and input trajectories form a ***stochastic process***

## Trajectories

$$X = (x_0, \dots, x_{T-1}, x_T)$$

$$U = (u_0, \dots, u_{T-1})$$

$$W = (w_0, \dots, w_{T-1})$$



## Policy cost

$$J(\theta) = \mathbf{E} \psi(X, U, W)$$

**Approximate  $J(\theta)$  from data (monte carlo simulation)**

$$\hat{J}(\theta) = \frac{1}{K} \sum_{i=1}^K \psi(X^i, U^i, W^i)$$

# COCP Example: dynamic programming

**Time-separable cost**

$$\psi(X, U, W) = \sum_{t=0}^{T-1} g(x_t, u_t, w_t)$$

**Optimal policy as  $T \rightarrow \infty$**

$$\phi(x_t) = \underset{u}{\operatorname{argmin}} \mathbf{E} (g(x_t, u, w_t) + V(f(x_t, u, w_t)))$$



**Value function**

**COCP if**

- $f$  affine in  $x$  and  $u$
- $g$  convex in  $x$  and  $u$
- $V$  is convex

# COCP Example: approximate dynamic programming

$$\phi(x_t) = \underset{u}{\operatorname{argmin}} \mathbf{E} \left( g(x_t, u, w_t) + \hat{V}(f(x_t, u, w_t)) \right)$$

Approximate  
value function

(even when  $V$  is not)

COCP if

- $f$  affine in  $x$  and  $u$
- $g$  convex in  $x$  and  $u$
- $\hat{V}$  is convex

# Controller tuning problem

**Goal**

$$\text{minimize } J(\theta)$$

**Nonconvex  
and difficult  
to solve**

**Traditional approaches**

- **Hand-tuning** (few parameters, simple dependencies)
- **Derivative-free method** (very slow)

# Learning scheme

## Auto-tuning

### Stochastic gradient descent

$$\theta^{k+1} = \theta^k - t^k \nabla_{\theta} \hat{J}(\theta^k)$$

↑  
stochastic gradient  
from simulation

step size →

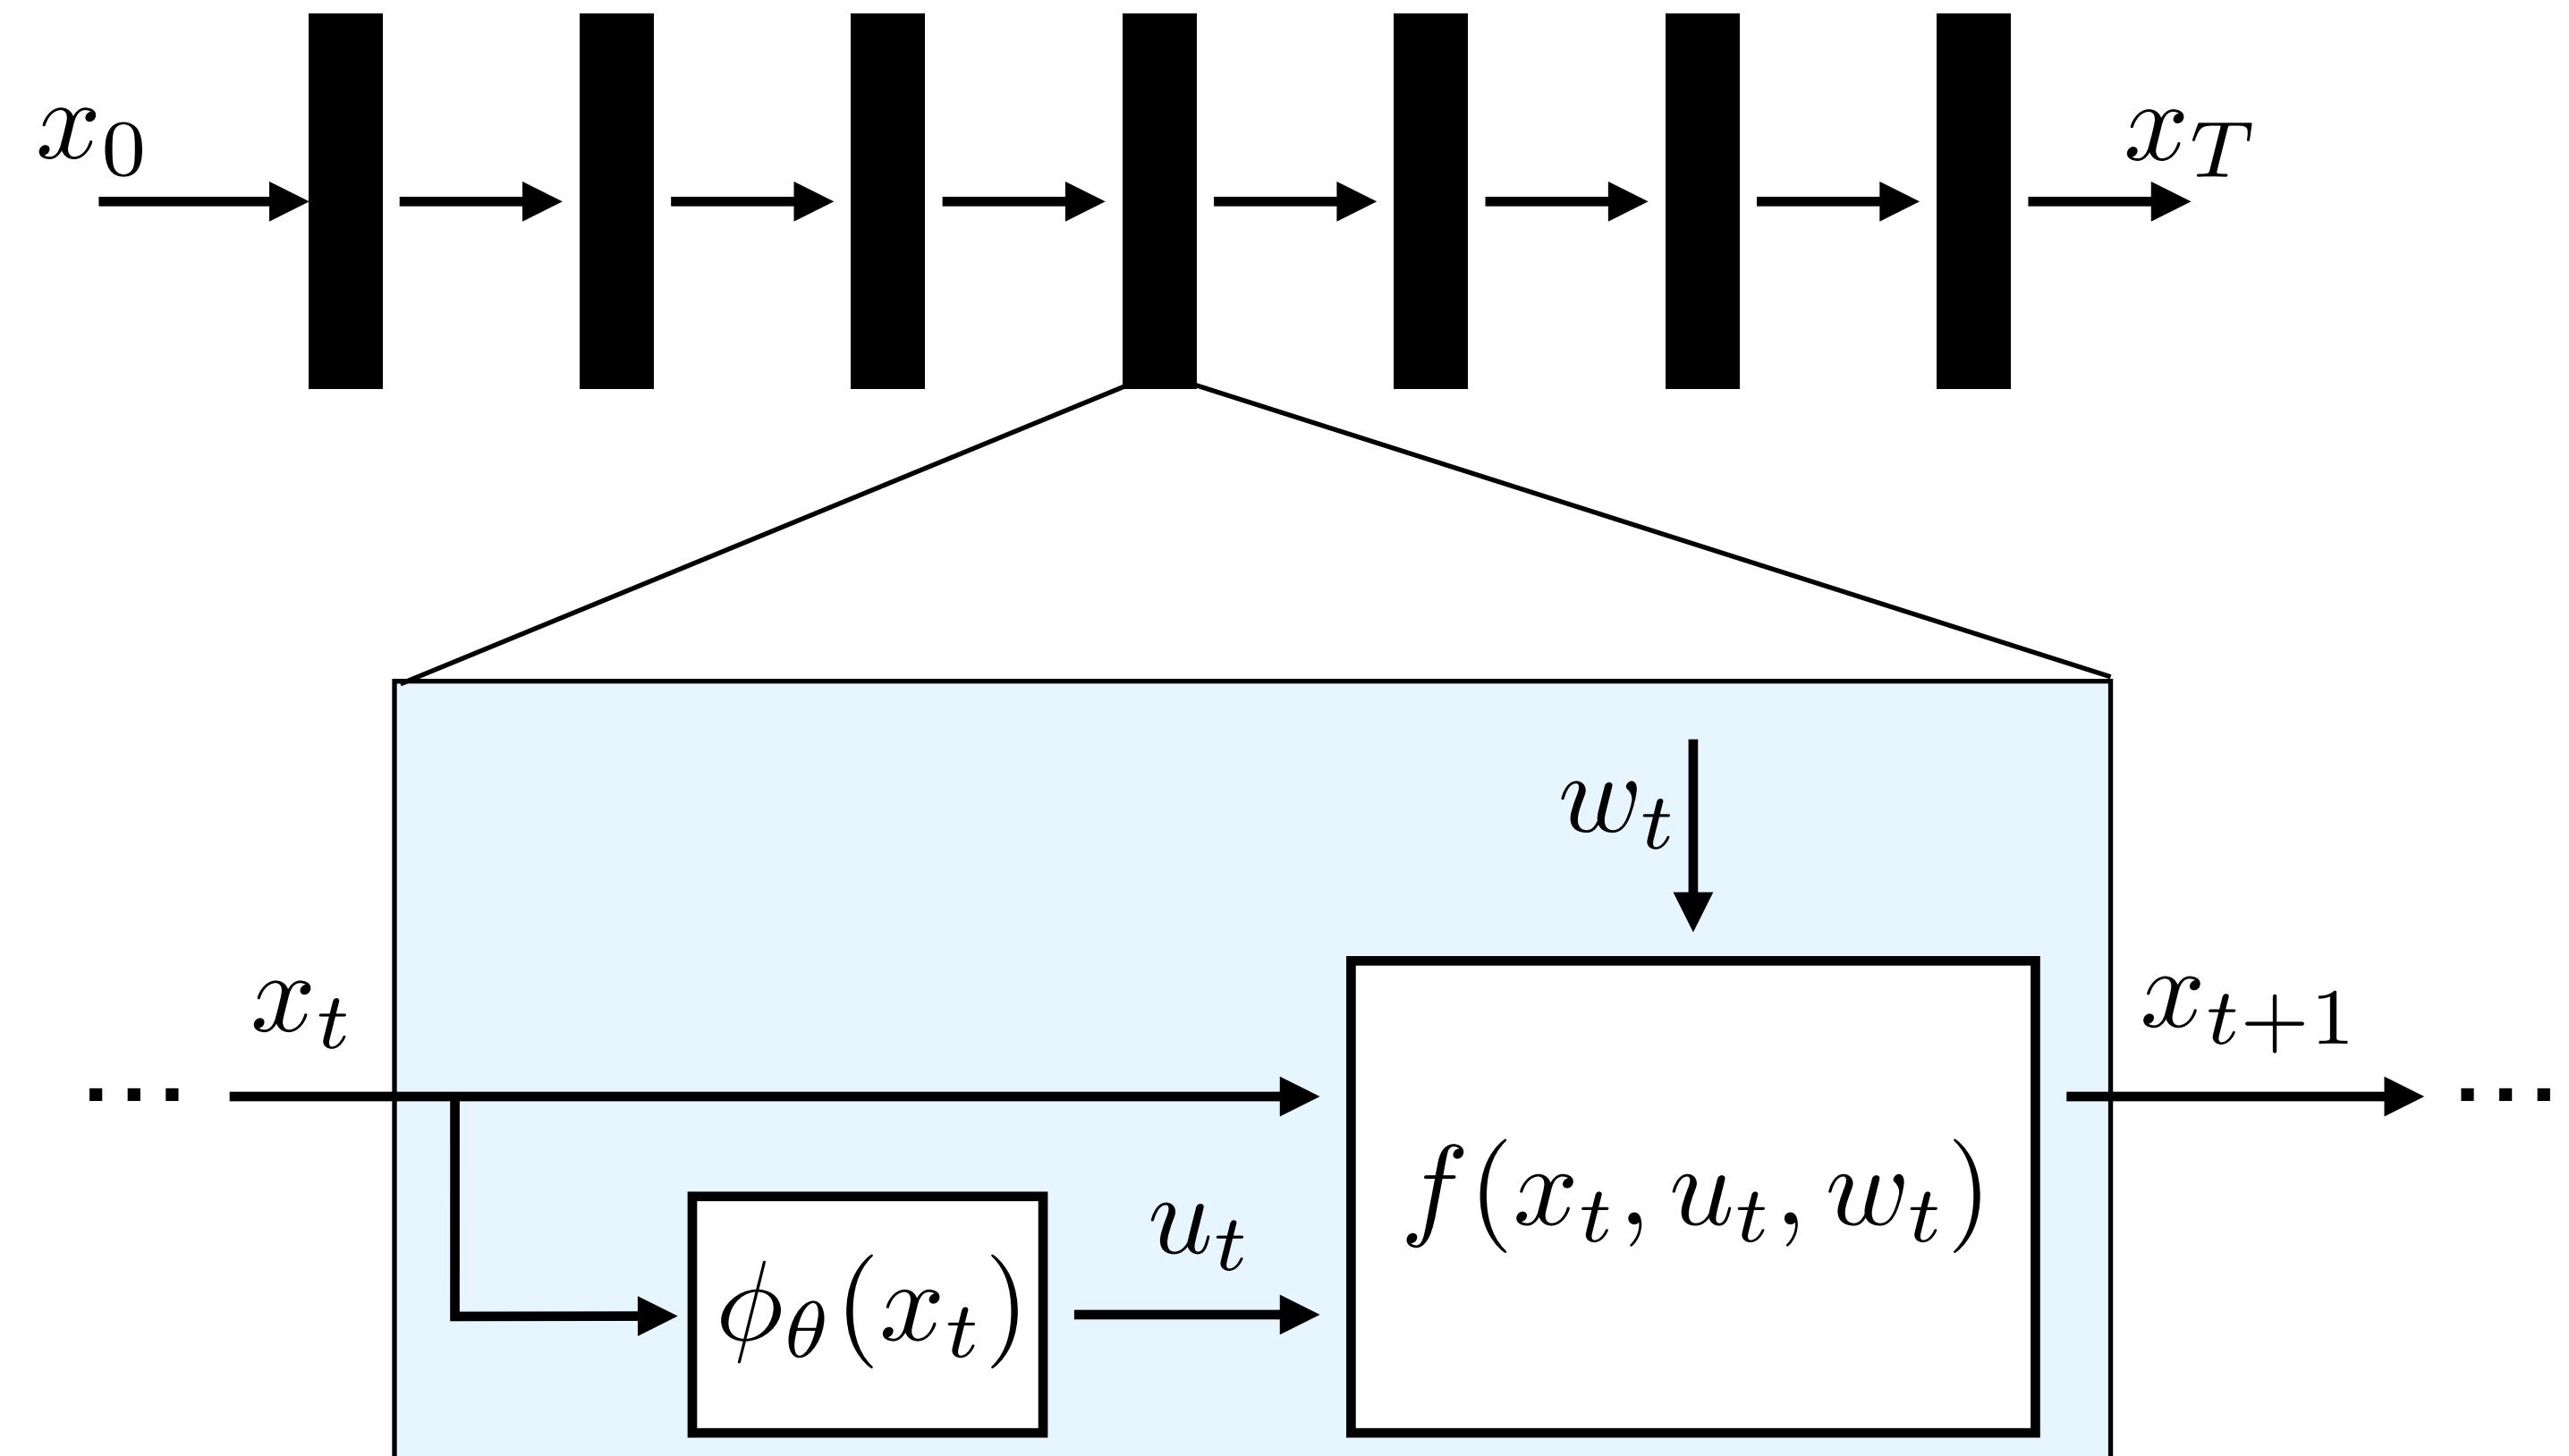
### Generalization

Split simulation data in training, validation and testing

### Non differentiable $\hat{J}(\theta)$ ?

Still get a descent direction (common in NN community)

# Implementation



## Automatic differentiation

- **Build** computation graph  
(simulate forward in time)
- **Backpropagate** using PyTorch

How do you backpropagate through  $\phi_\theta$  ?

# Differentiating through convex optimization problems

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

- $x$  variable
- $\theta$  parameter

## Optimality conditions

$$\begin{array}{c} \text{stationarity} \longrightarrow \boxed{\nabla_x f(x, \theta) + D_x g(x, \theta)^T y = 0} \\ \text{complementary slackness} \longrightarrow \boxed{\begin{matrix} \text{diag}(y)g(x, \theta) = 0 \\ g(x^*, \theta) \leq 0 \\ y^* \geq 0 \end{matrix}} \longrightarrow F(z^*(\theta), \theta) = 0 \end{array}$$

primal/dual  
solution

$$z^* = (x^*(\theta), y^*(\theta))$$

## Goal

Compute  $Dz^*(\theta)$

# Differentiating through convex optimization problems

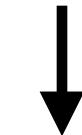
$$F(z^*(\theta), \theta) = 0$$

primal/dual  
solution

$$z^* = (x^*(\theta), y^*(\theta))$$

## Implicit function theorem

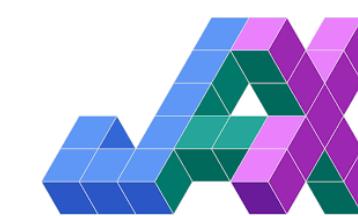
$$D_z F(z^*, \theta) Dz^*(\theta) + D_\theta F(z^*, \theta) = 0$$



$$Dz^*(\theta) = - (D_z F(z^*, \theta))^{-1} D_\theta F(z^*, \theta) \quad (D_z F(z^*, \theta) \text{ must be invertible})$$

one linear  
system  
solution

We plug  $Dz^*(\theta)$  in AD  
(automatic differentiation)  PyTorch



# Box-constrained LQR

## Problem setup

- dynamics:  $x_{t+1} = Ax_t + Bu_t + w_t$
- actuator limit:  $\|u_t\|_\infty \leq 1$
- stage cost:  $x_t^T Q x_t + u_t^T R u_t$

## COCP Policy (QP)

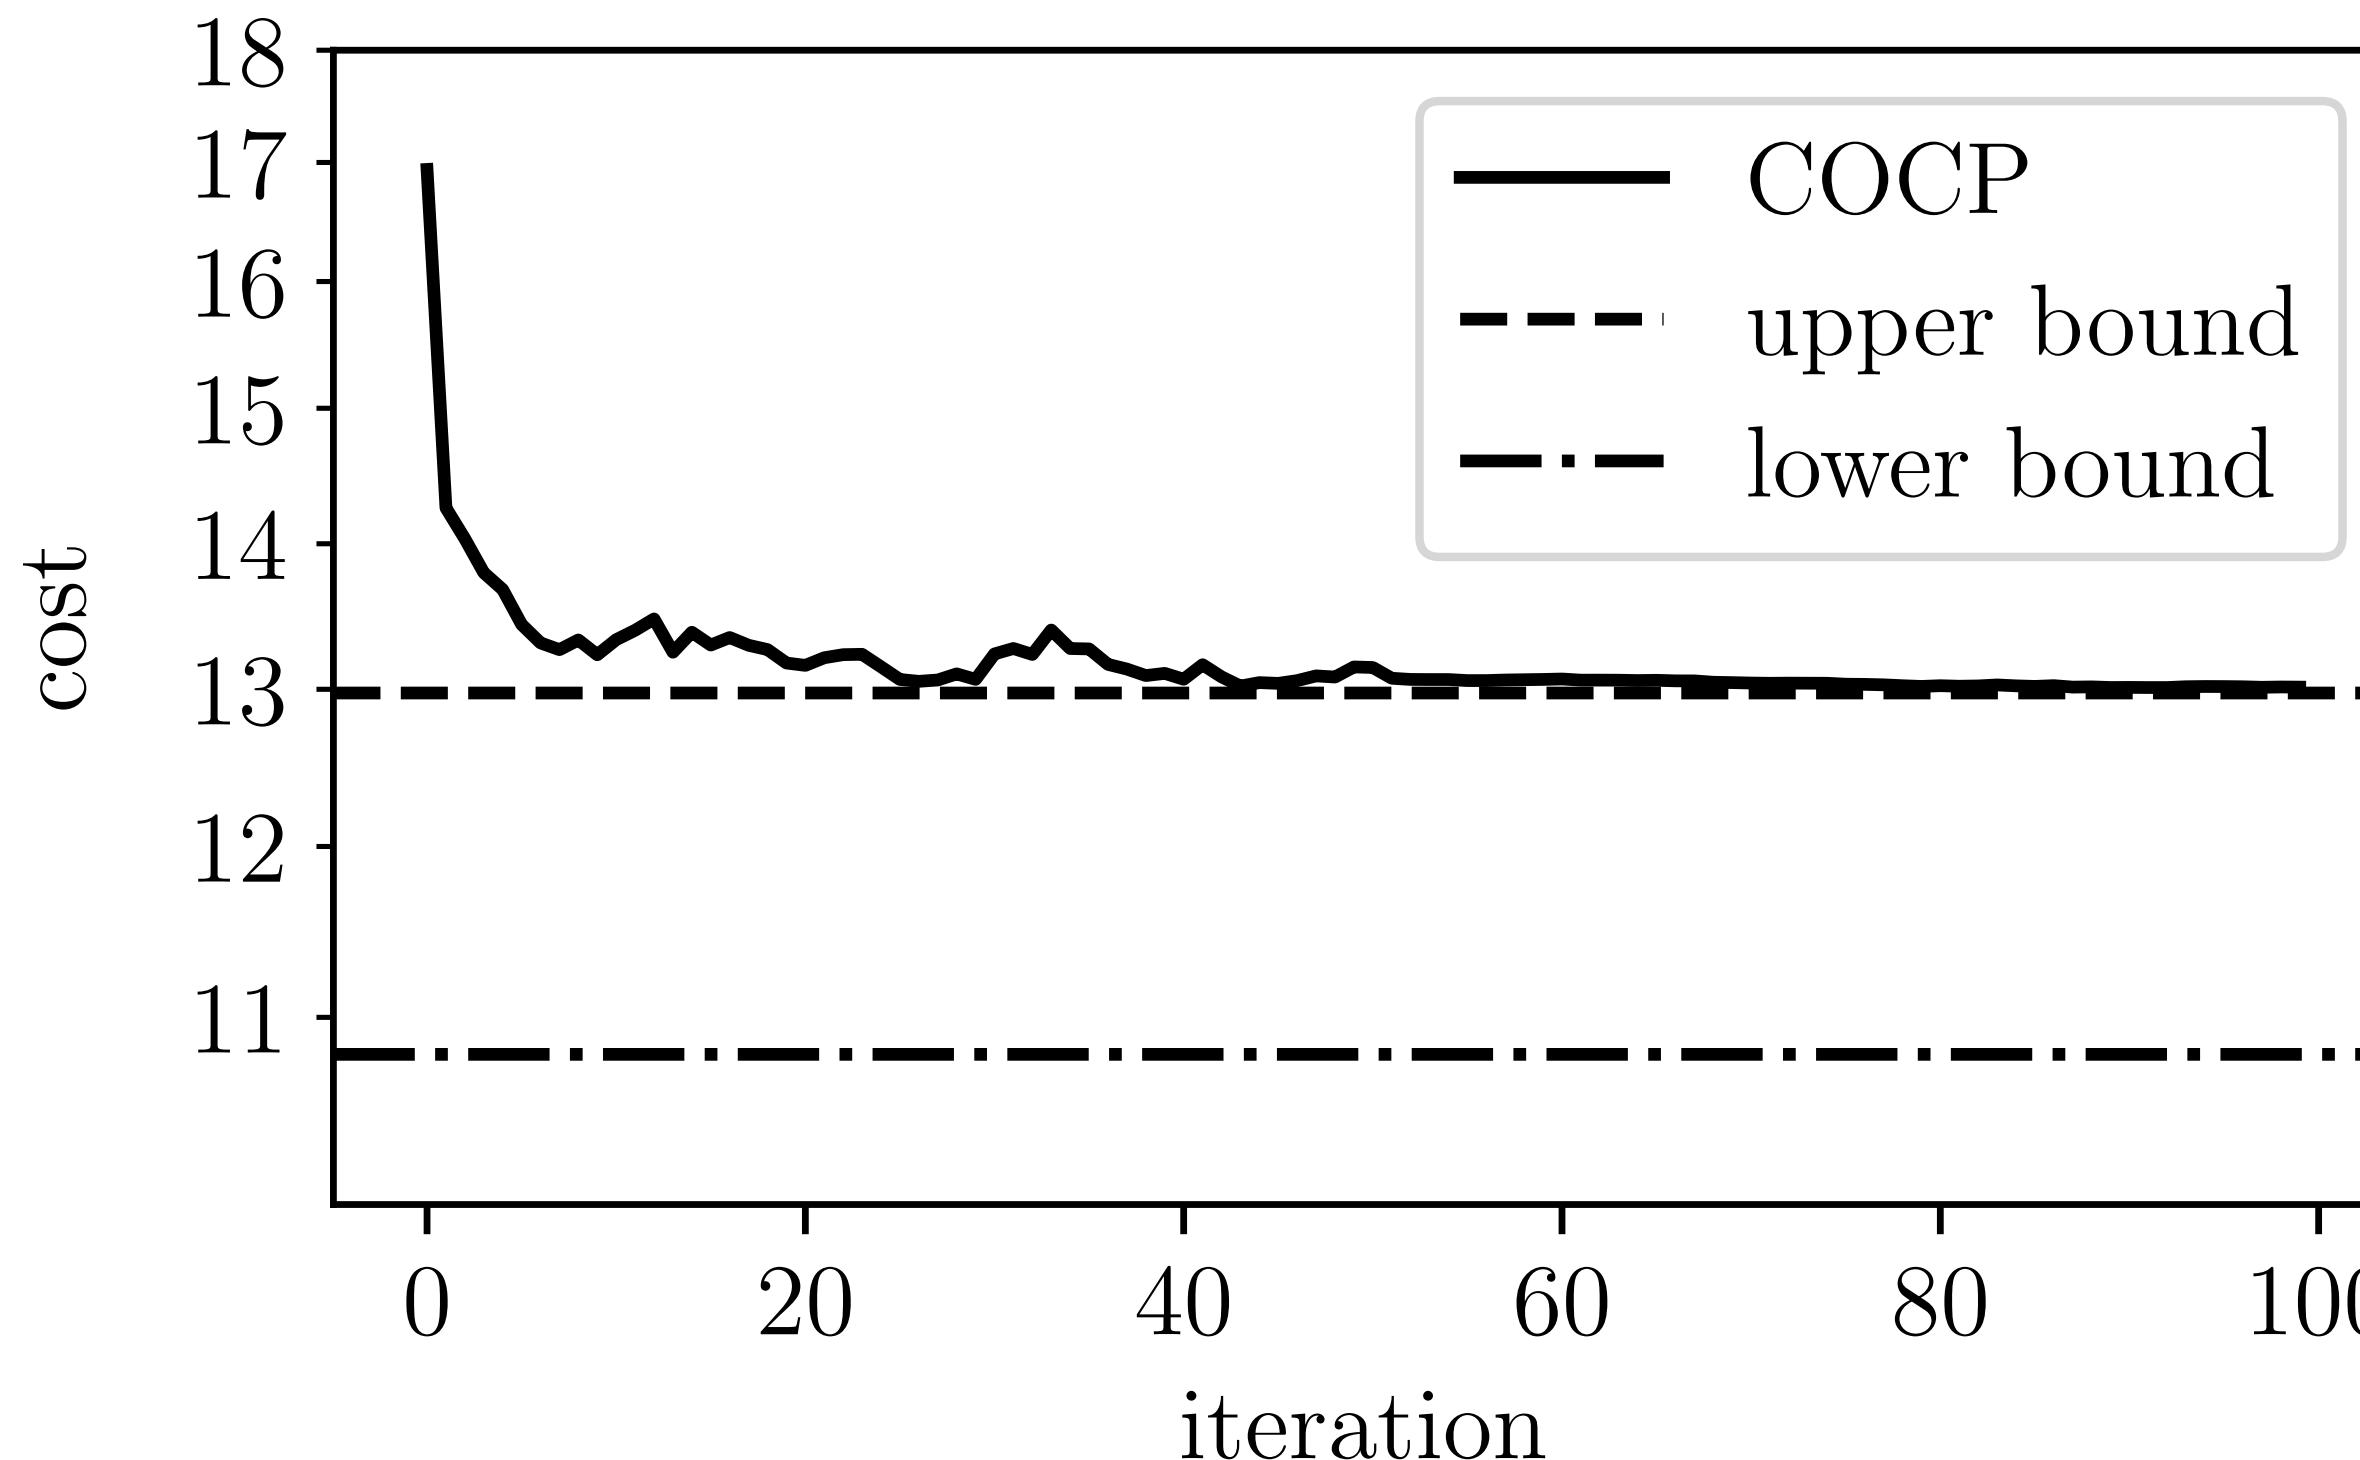
$$u_t = \underset{u}{\operatorname{argmin}} \quad u^T R u + \|\theta(Ax_t + Bu)\|_2^2$$

subject to  $\|u\|_\infty \leq 1$

parameters

# Box-constrained LQR

## Performance



**Standard  
upper/lower bounds  
from SDPs**

↓

**Hard to generalize**  
(other dynamics,  
disturbances, etc)

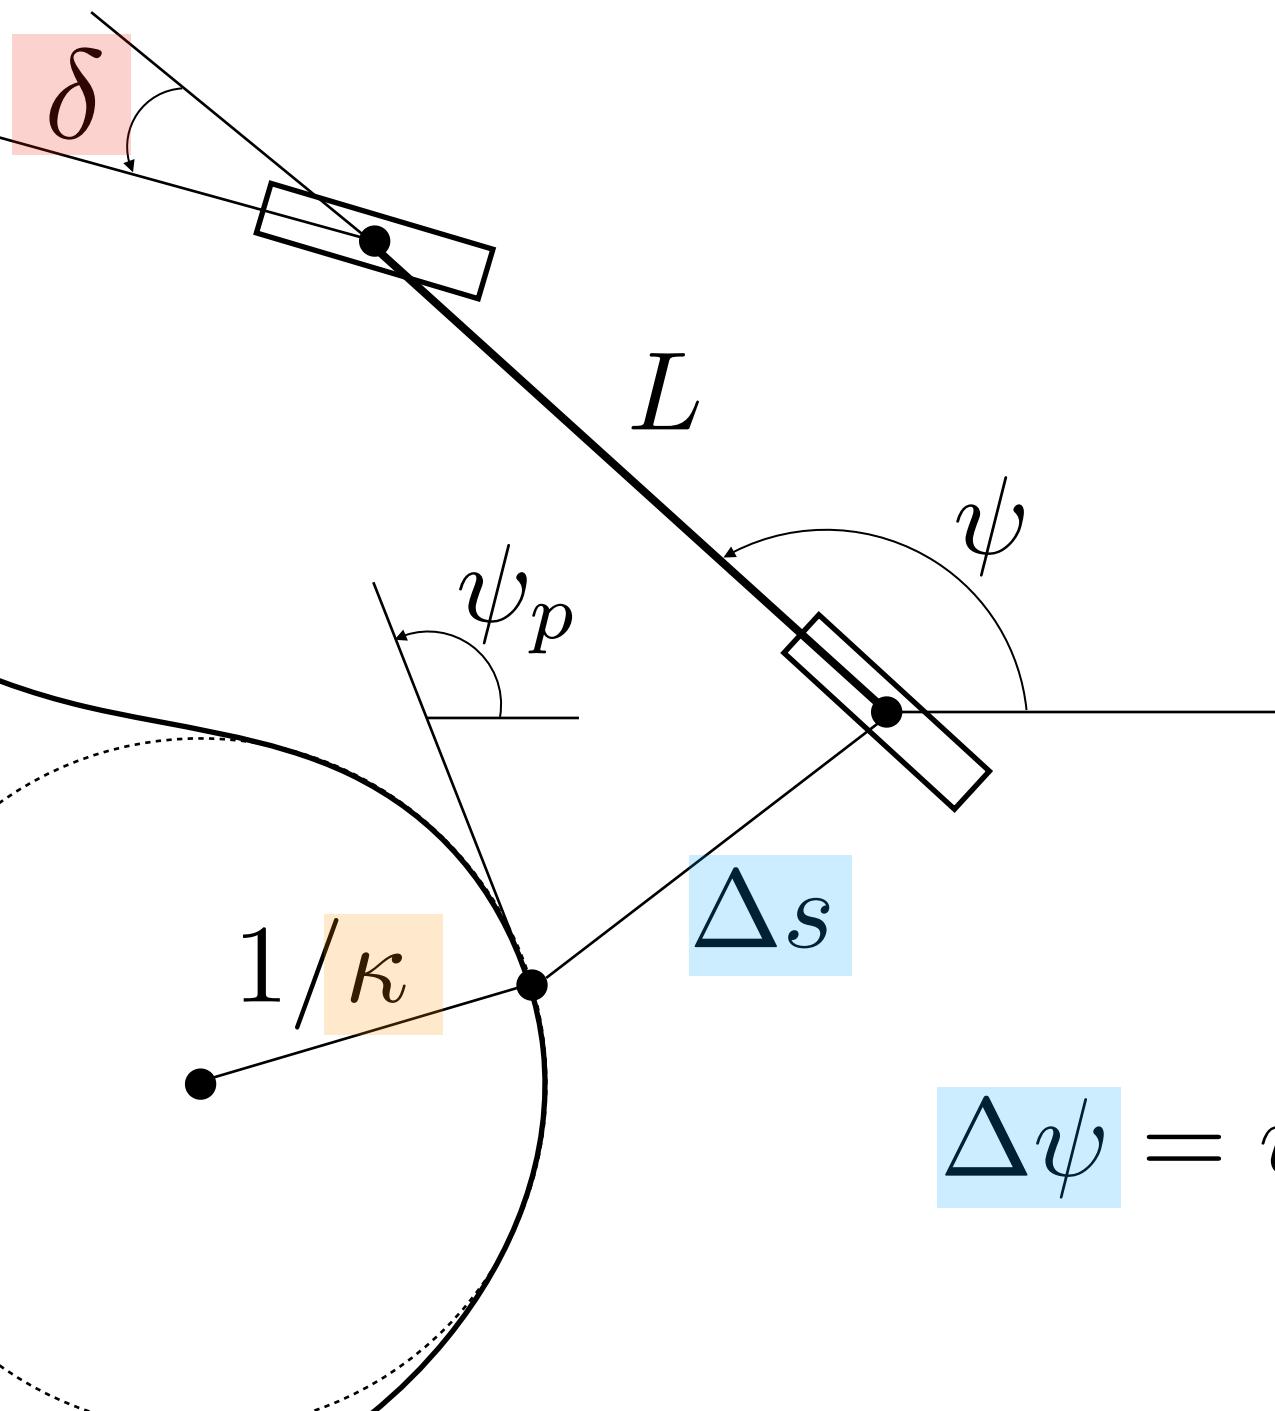
# Vehicle tracking curved paths

lateral deviation      heading deviation

**State:**  $x_t = (\Delta s_t, \Delta \psi_t, v_t, v_t^{\text{des}}, \kappa_t)$

path curvature

true and desired  
velocities



$$\Delta \psi = \psi - \psi_p$$

acceleration

**Input:**  $u_t = (a_t, z_t)$

turn

$$z_t = \tan(\delta_t) - L\kappa_t$$

## Dynamics

$$x_{t+1} = f(x_t, u_t, w_t)$$

# Vehicle tracking curved paths

## Cost and constraints

track velocity

Stage cost

small control effort

( $v_t - v_t^{\text{des}})^2 + \Delta s_t^2 + \Delta \psi_t^2 + \lambda(|a_t| + z_t^2)$ )

↑

track path

maximum acceleration

$|a_t| \leq a_{\max}$

maximum turn

$|z + L\kappa_t| \leq \tan(\delta_{\max})$

The diagram illustrates the cost and constraints for vehicle tracking. At the top center, the text "Stage cost" is displayed above the formula  $(v_t - v_t^{\text{des}})^2 + \Delta s_t^2 + \Delta \psi_t^2 + \lambda(|a_t| + z_t^2)$ . To the left of the formula, an arrow points from the text "track velocity". To the right, another arrow points from the text "small control effort". Below the stage cost formula, an upward-pointing arrow is labeled "track path". At the bottom, two pairs of text and equations represent constraints. On the left, "maximum acceleration" is shown with the equation  $|a_t| \leq a_{\max}$ . On the right, "maximum turn" is shown with the equation  $|z + L\kappa_t| \leq \tan(\delta_{\max})$ .

# Vehicle tracking curved paths

## COCP as a Convex QP

$$u_t = (a_t, z_t) = \phi(x_t) = \operatorname{argmin}_u |a| + z^2 + \|S\mathbf{y}\|_2^2 + q^T \mathbf{y}$$

subject to  $x_t = (\Delta s_t, \Delta \psi_t, v_t, v_t^{\text{des}}, \kappa_t)$

$u = (a, z)$

**next state**  $\longrightarrow \mathbf{y} = f(x_t, u, 0)$

$|a| \leq a_{\max}$

$|z + L\kappa_t| \leq \tan(\delta_{\max})$

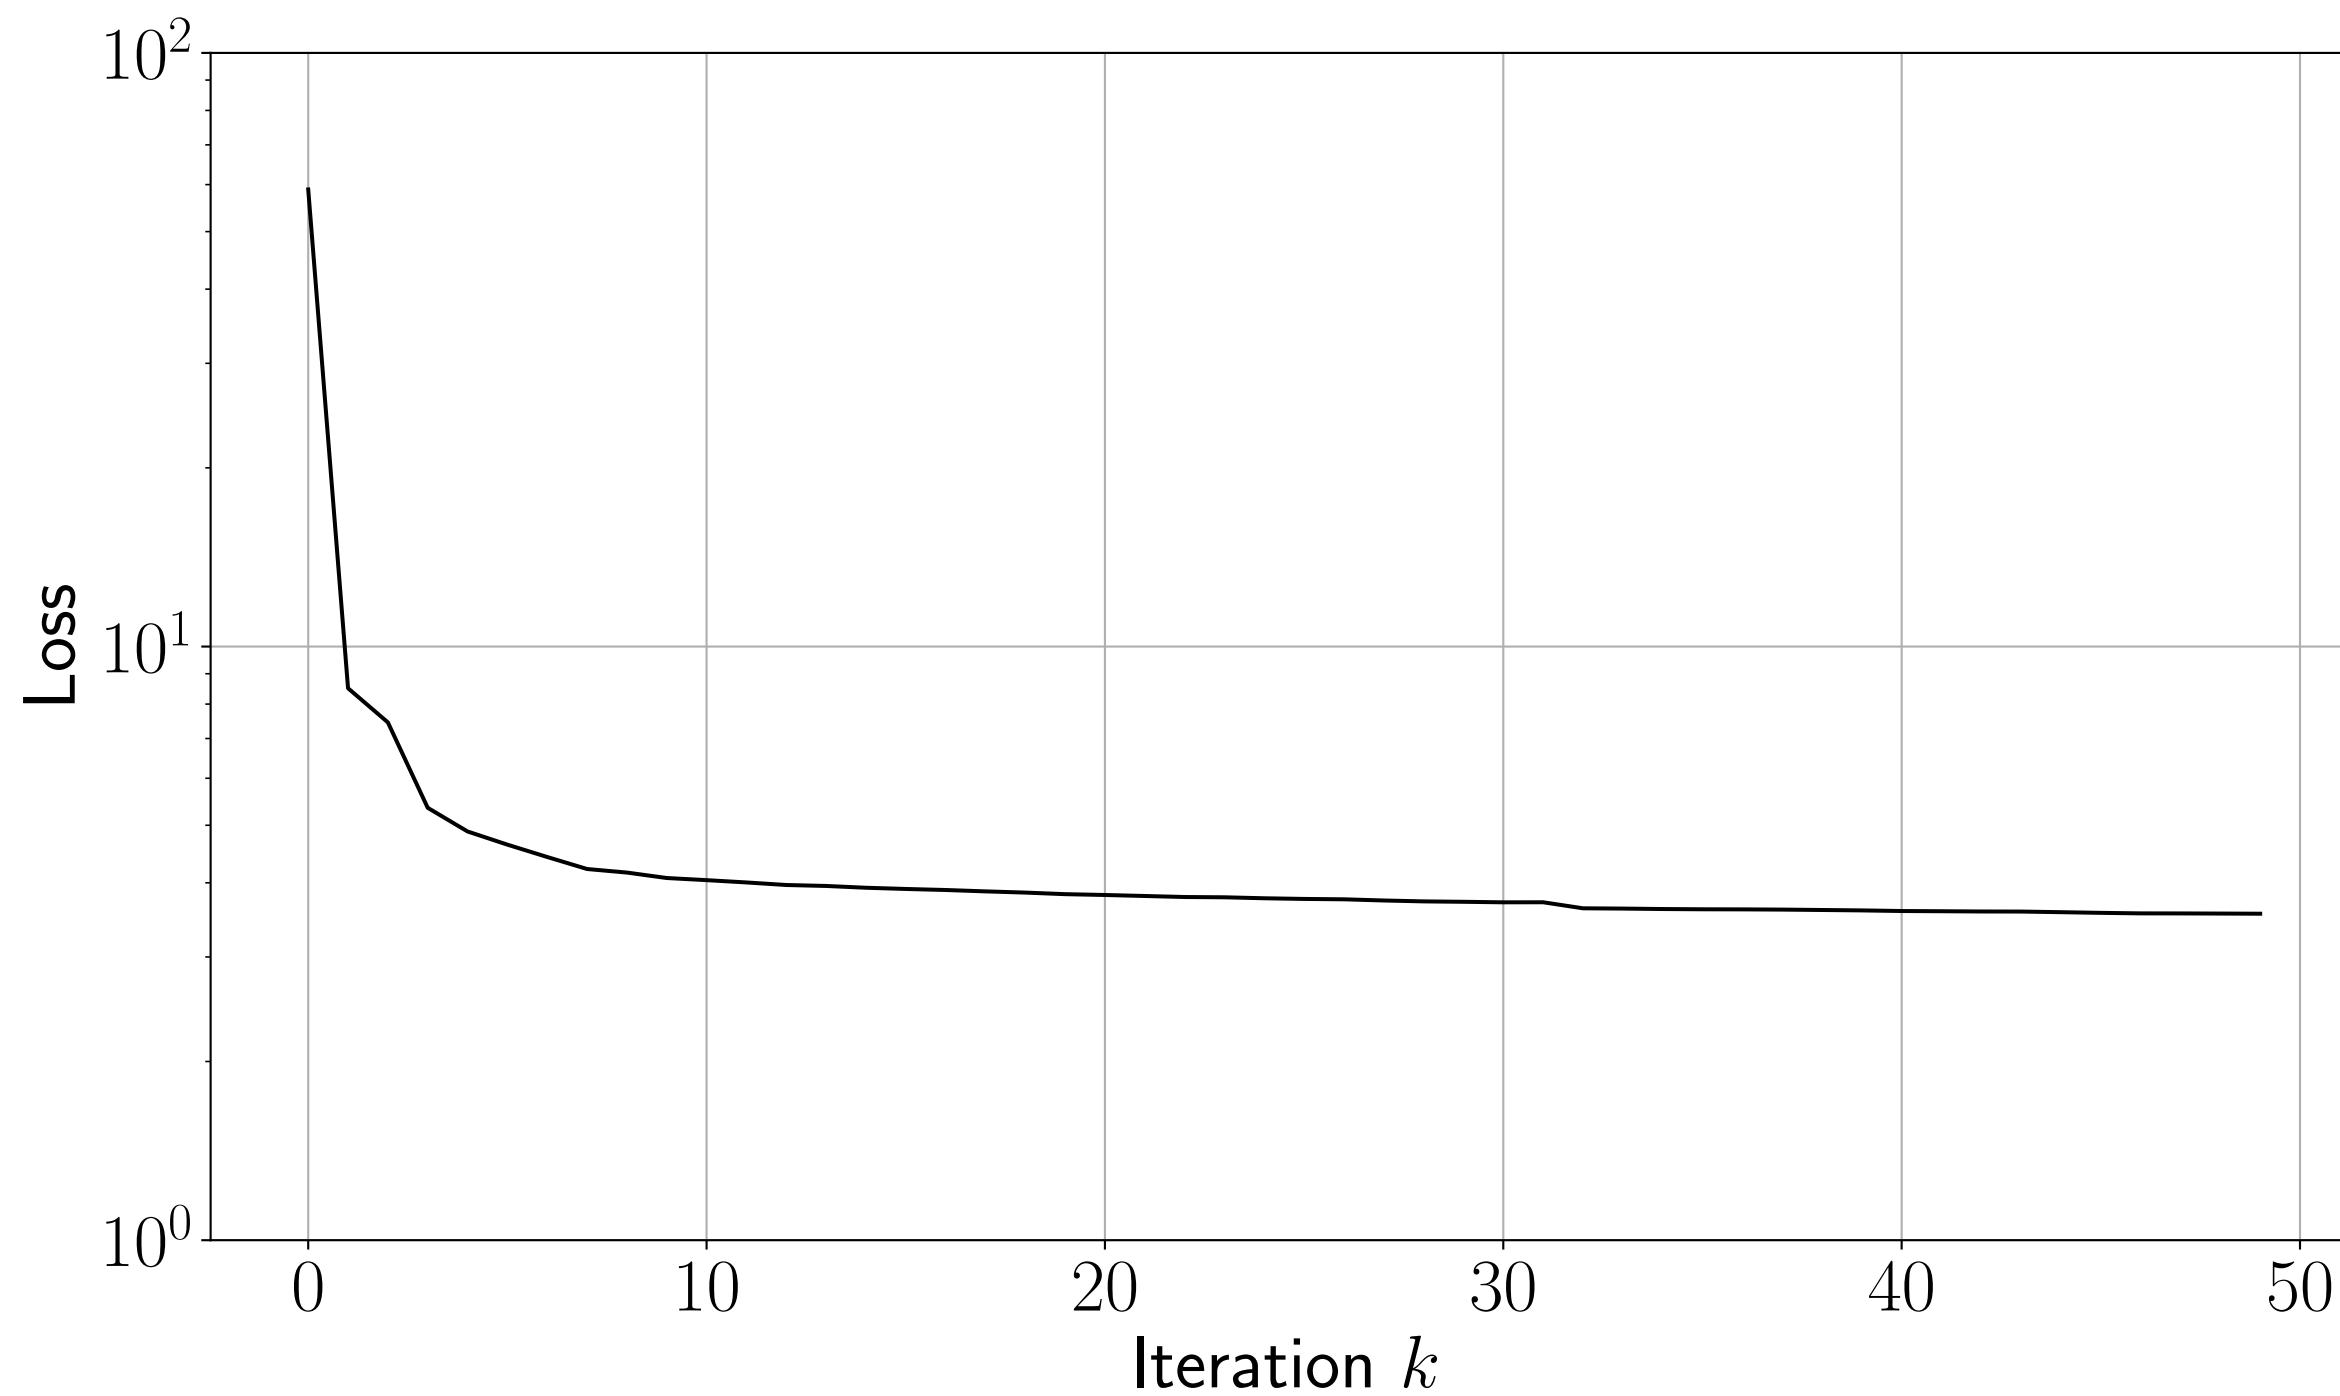
**parameters**

```
graph TD; A[parameters] --> B[|a|]; A --> C[z^2]; A --> D["\|S\mathbf{y}\|_2^2"]; A --> E["q^T \mathbf{y}"]
```

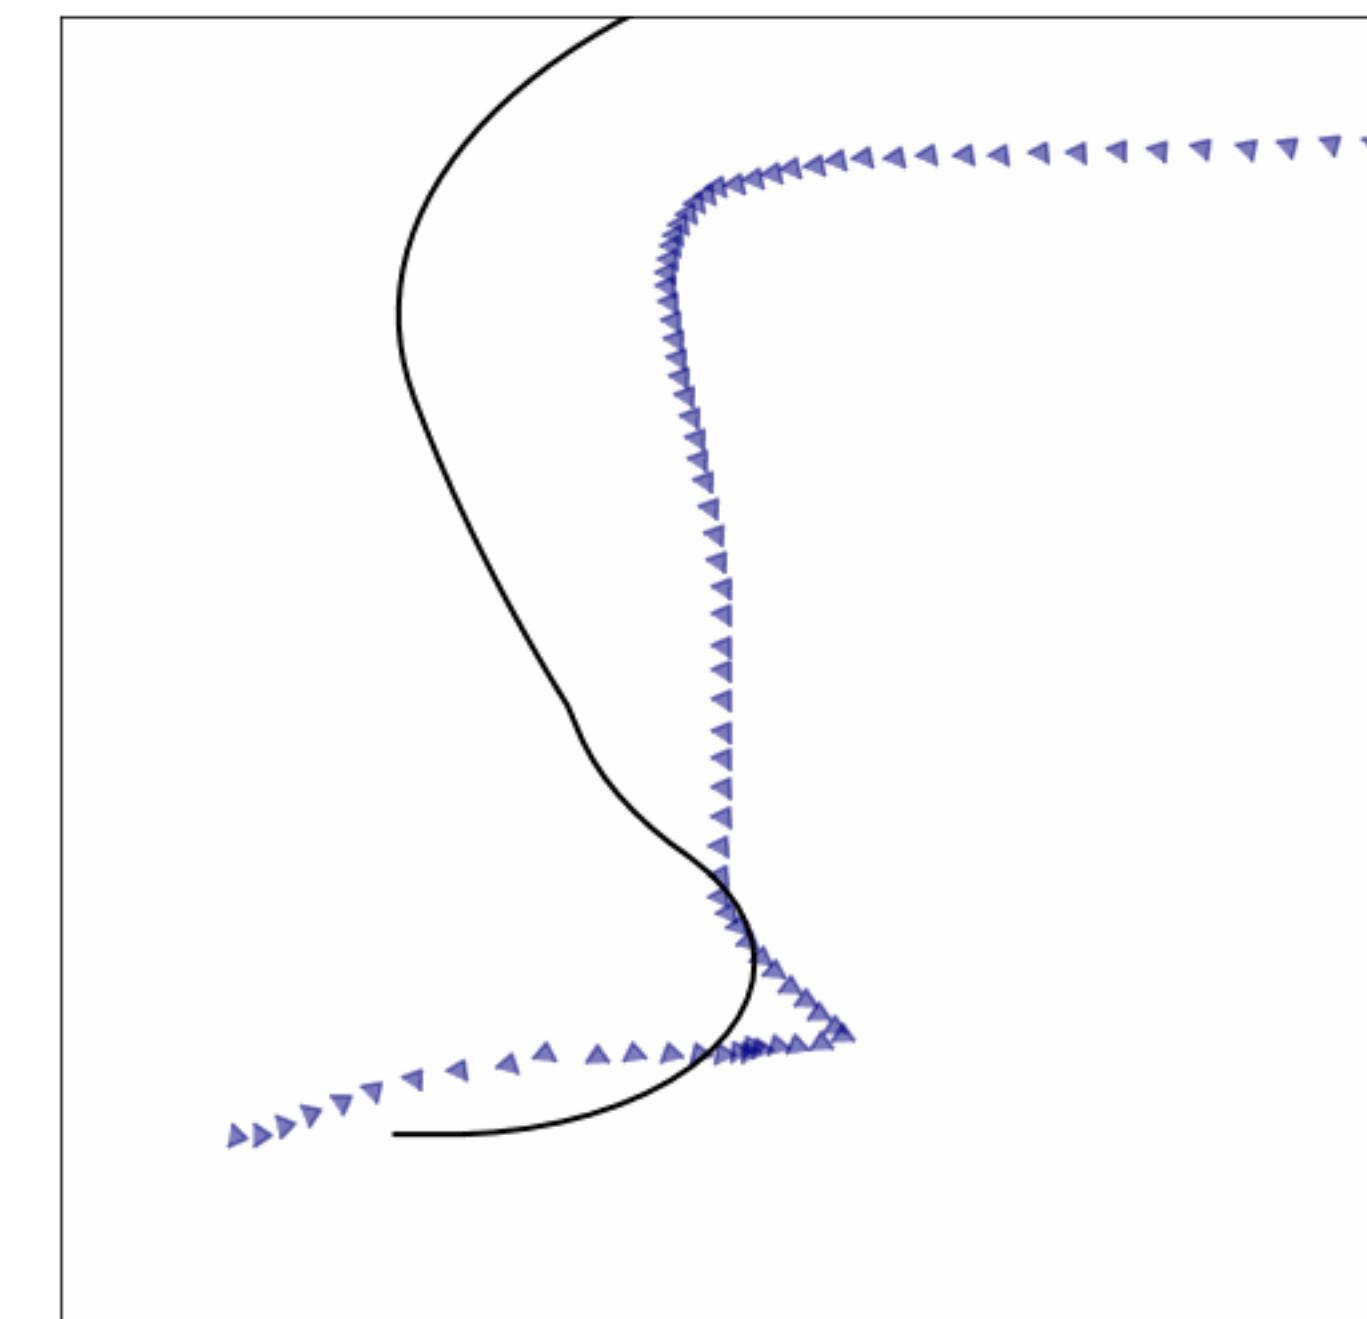
# Vehicle tracking curved paths

## Results

**Validation loss**



**Tracking behavior**



Good trajectory in very few iterations

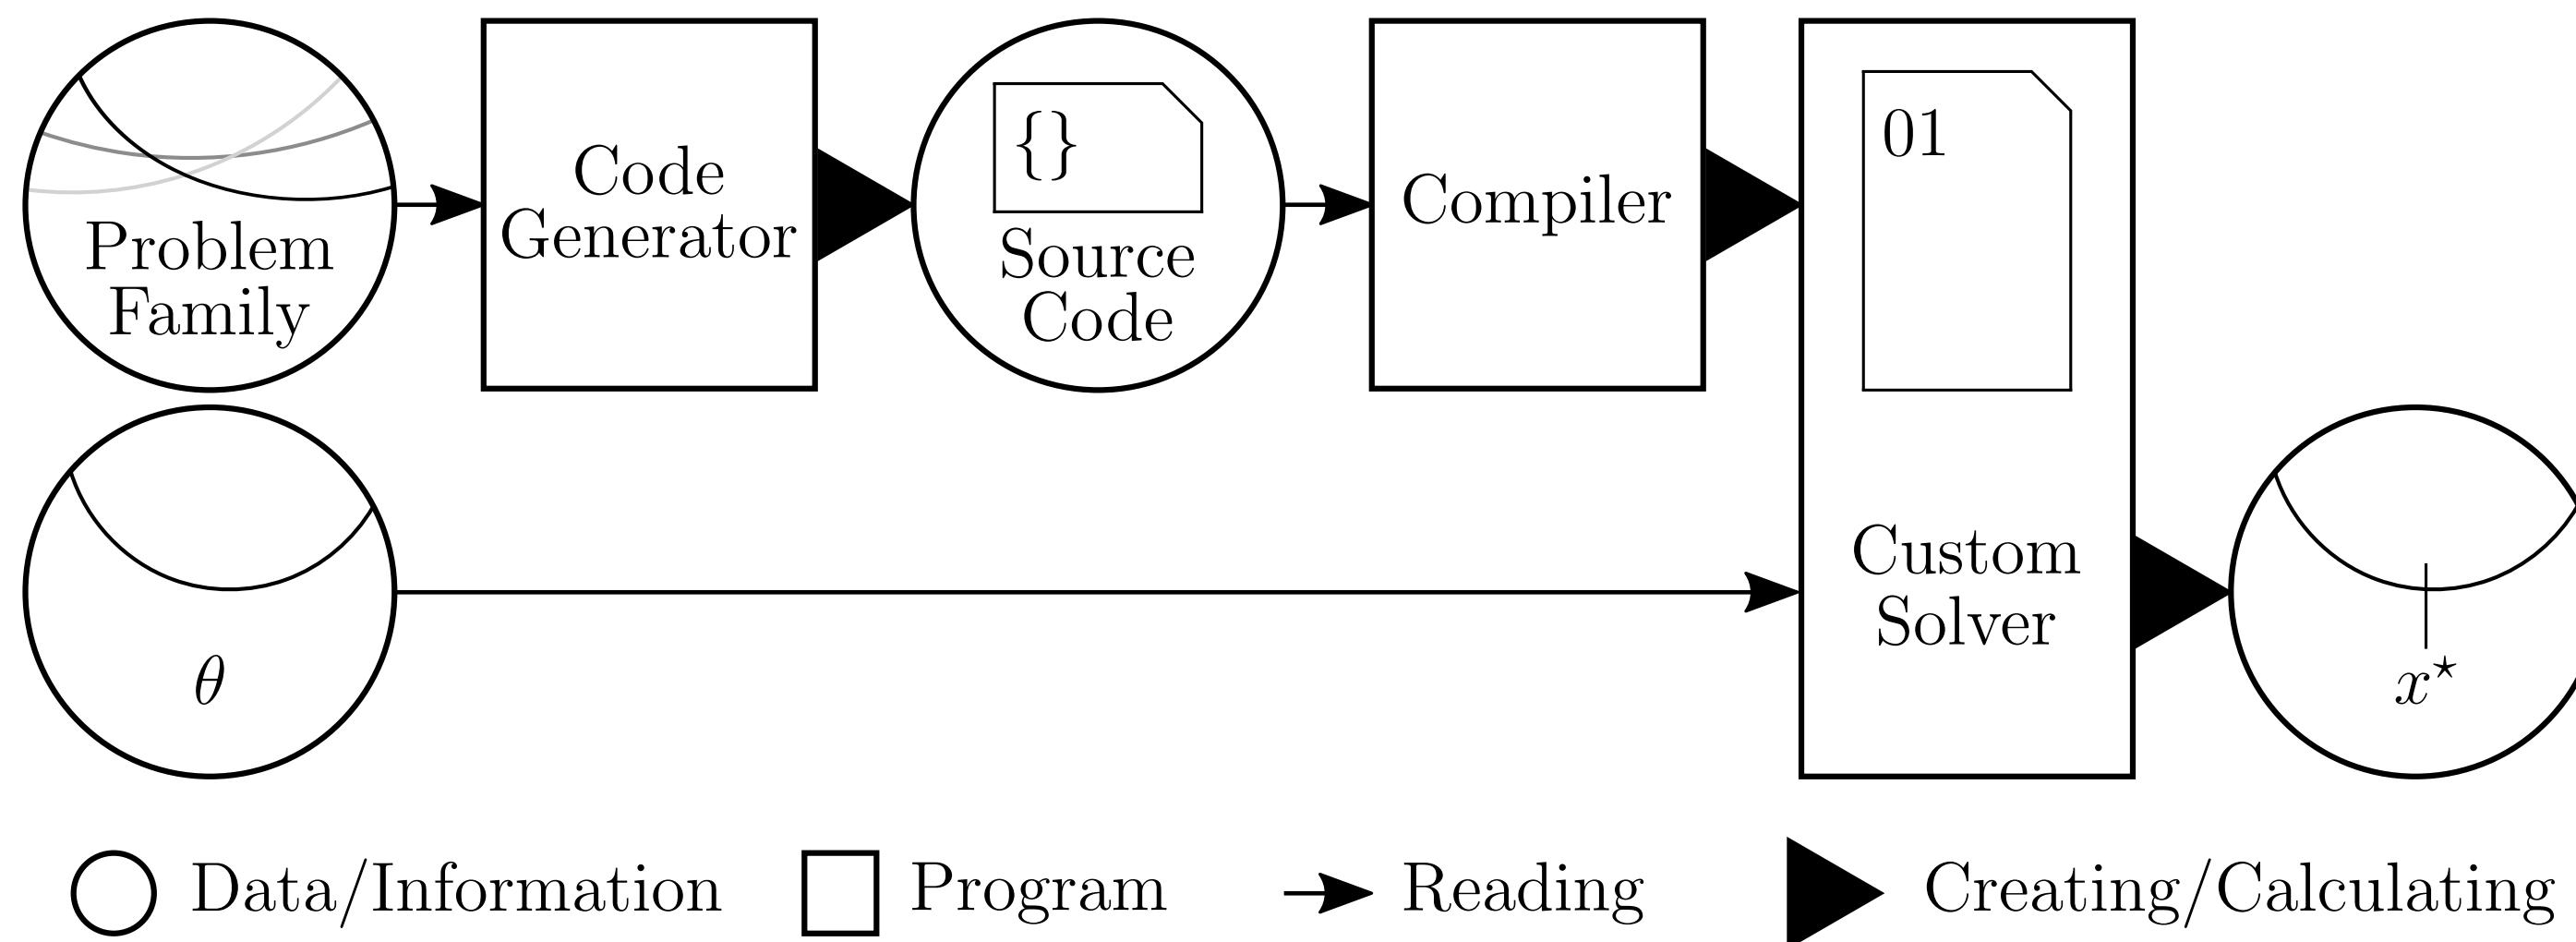
# Code generation for parametric convex optimization

## Example

$$\begin{aligned} & \text{minimize} && \|Gx - h\|^2 \\ & \text{subject to} && x \geq 0 \end{aligned}$$

### CVXPYgen

<https://pypi.org/project/cvxpygen/>



**It supports convex quadratic (OSQP),  
and conic (ECOS, SCS) optimization.**

```
import cvxpy as cp
from cvxpygen import cpg

# model problem
x = cp.Variable(n, name='x')
G = cp.Parameter((m,n), name='G')
h = cp.Parameter(m, name='h')
p = cp.Problem(cp.Minimize(cp.sum_squares(G@x-h)),
               [x>=0])
# generate code
cpg.generate_code(p)
```

# Learning COCPs summary

Interpretable

Satisfy  
constraints

Handle varying  
dynamics

Efficient and reliable  
(code generation)

Easy to learn  
from data

## Future work

- Hybrid (mixed-integer) control policies
- Stochastic policies with safety-guarantees

# **Conclusions**

# Acknowledgements

Goran Banjac



Paul Goulart



Alberto Bemporad



Ken Goldberg



Stephen Boyd



Akshay Agrawal



Shane Barratt



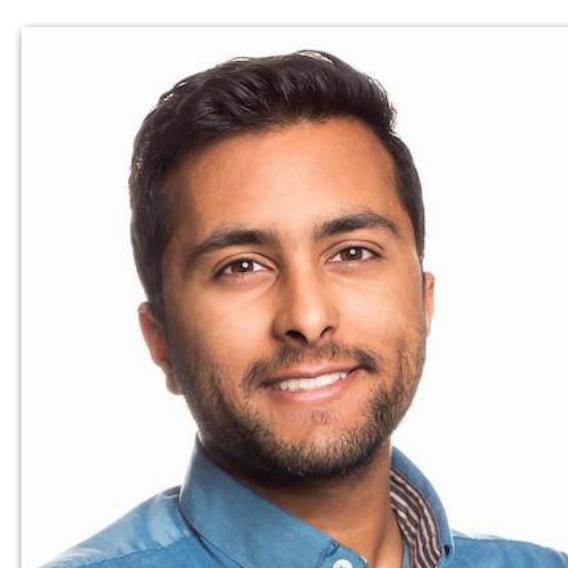
Jeff Ichnowski



Max Schaller



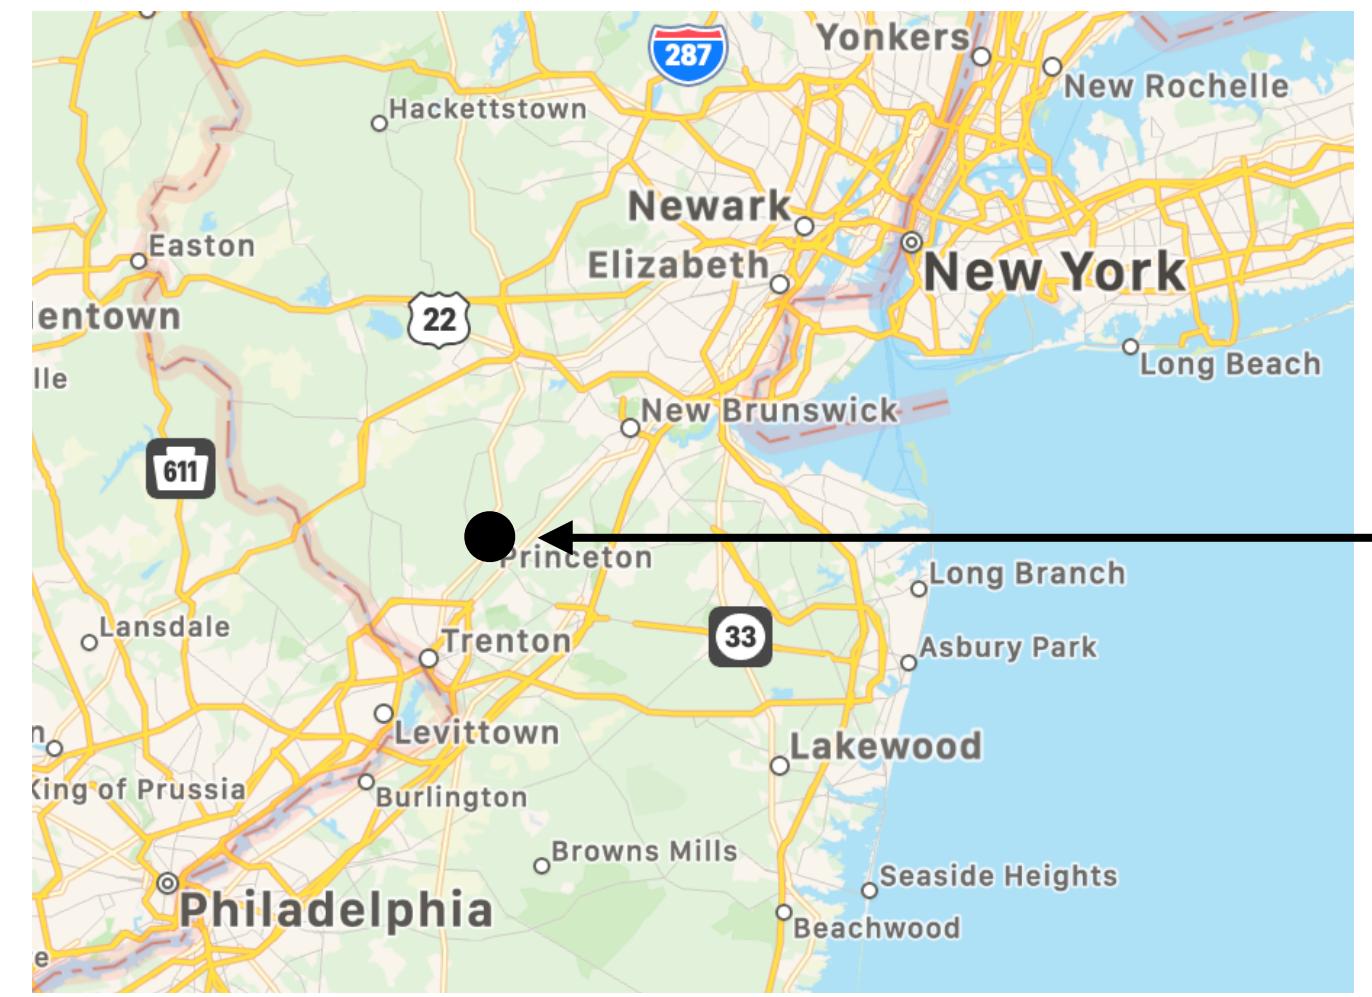
Paras Jain



Francesco Borrelli



# Come for a visit!



Princeton

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

49

50

51

52

53

54

55

56

57

58

59

60

61

62

63

64

65

66

67

68

69

70

71



PRINCETON  
UNIVERSITY

# References

## OSQP ([osqp.org](https://osqp.org))

[Accelerating Quadratic Optimization with Reinforcement Learning. Ichnowski, Jain, *Stellato*, Banjac, Luo, Gonzalez, Stoica, Borrelli, and Goldberg. NeurIPS 2021]

[OSQP: An Operator Splitting Solver for Quadratic Programs. *Stellato*, Banjac, Goulart, Bemporad, and Boyd. Mathematical Programming Computation 2020]

[Infeasibility detection in the alternating direction method of multipliers for convex optimization. Banjac, Goulart, *Stellato*, and Boyd. Journal of Optimization Theory and Applications 2019]

[Embedded code generation using the OSQP solver. *Stellato*, Banjac, Stellato, Moehle, Goulart, Bemporad, and Boyd. IEEE Conf. on Decision and Control 2017]

## Learning COCPs

[Embedded Code Generation with CVXPY. Schaller, Banjac, Diamond, Agrawal, *Stellato*, and Boyd. IEEE Control Systems Letters 2022]

[Learning Convex Optimization Control Policies. Agrawal, Amos, Barratt, Boyd, and Stellato. Learning for Dynamics and Control (L4DC) 2020]

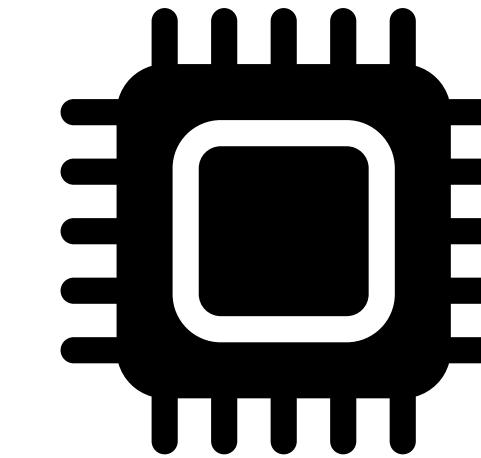
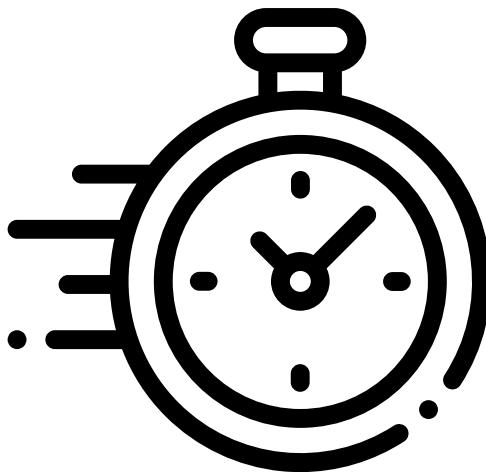
[Differentiable convex optimization layers. Agrawal, Amos, Barratt, Boyd, Diamond, and Kolter. NeurIPS 2019]

(<https://pypi.org/project/cvxpygen/>)

(<https://github.com/cvxgrp/cocp>)

# Conclusions

Real-time and embedded optimization



will soon be applied everywhere

Thanks to

Efficient and  
reliable optimizers

Data-driven  
control policies



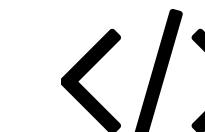
stellato.io



@b\_stellato



bstellato@princeton.edu



github.com/bstellato