

Accelerating Quadratic Optimization with Reinforcement Learning



PRINCETON
UNIVERSITY

QRFE

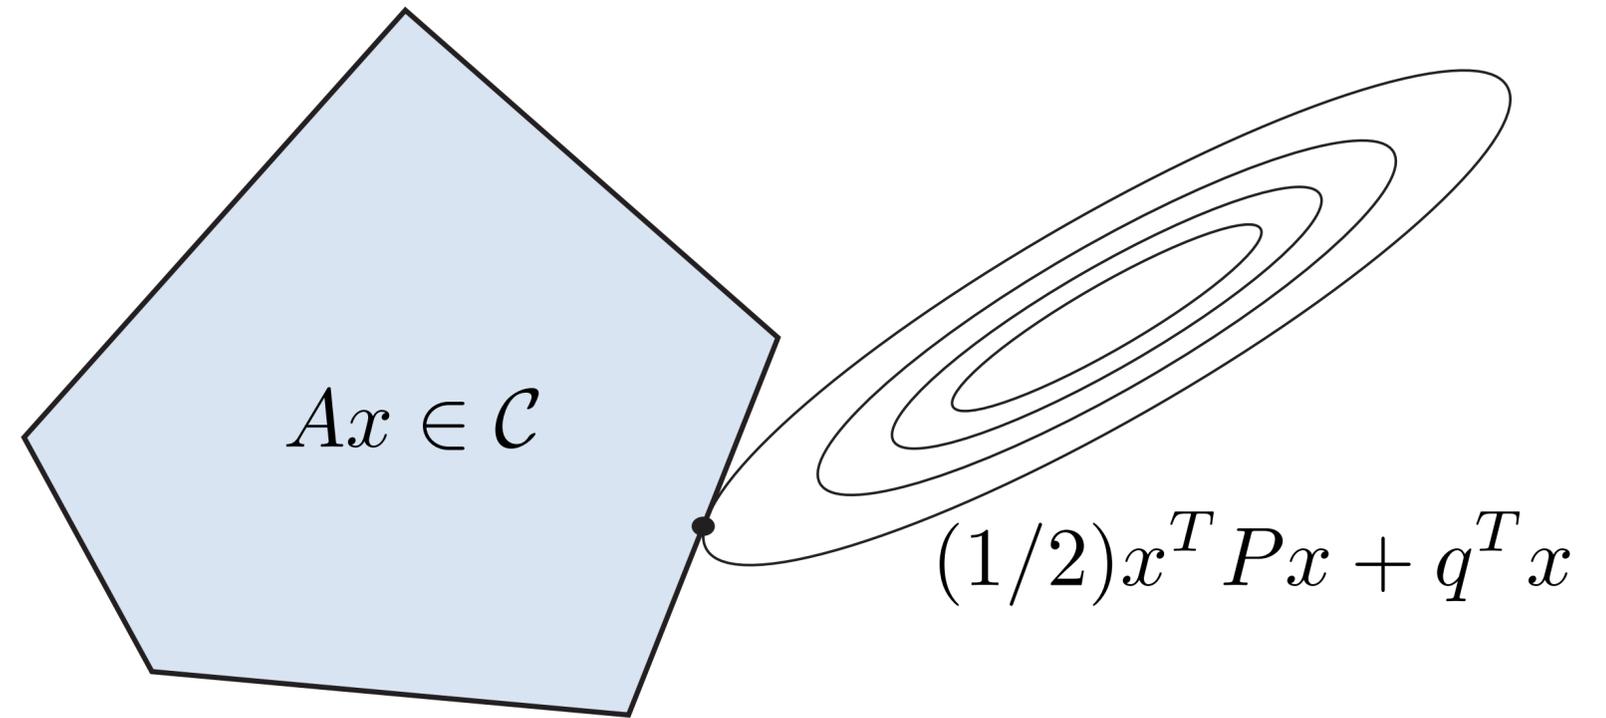
Bartolomeo Stellato – INFORMS Annual Meeting, October 24 2021



The problem

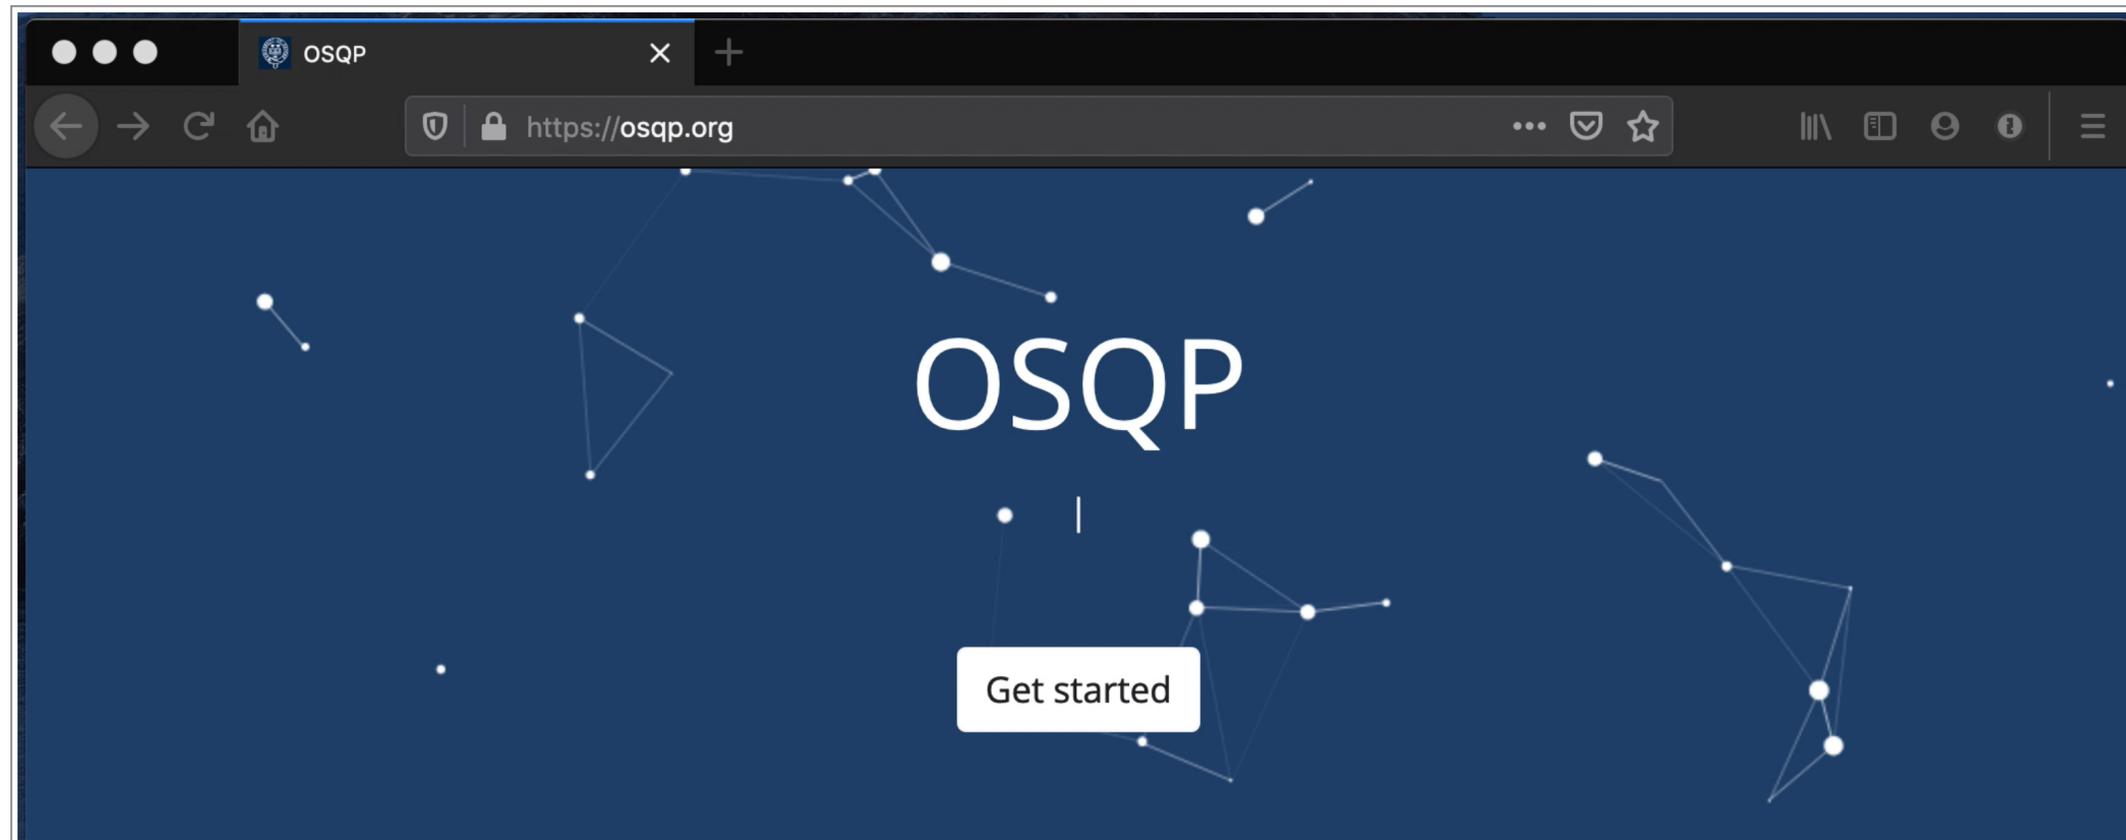
$$\begin{array}{ll} \text{minimize} & (1/2)x^T P x + q^T x \\ \text{subject to} & Ax \in \mathcal{C} \end{array}$$

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



OSQP

Operator Splitting solver for Quadratic Programs



Embeddable
(can be division free!)

Supports
warm-starting

Detects
infeasibility

Solves large-scale
problems

OSQP algorithm

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



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

KKT System

Algorithm

Solve

$$\begin{bmatrix} P & A^T \\ A & -\text{diag}(\rho)^{-1} \end{bmatrix} \begin{bmatrix} x^{k+1} \\ \nu^{k+1} \end{bmatrix} = \begin{bmatrix} -q \\ z^k - \text{diag}(\rho)^{-1} y^k \end{bmatrix}$$

$$\tilde{z}^{k+1} \leftarrow z^k + \text{diag}(\rho)^{-1} (\nu^{k+1} - y^k)$$

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

$$y^{k+1} \leftarrow y^k + \text{diag}(\rho) (\tilde{z}^{k+1} - z^{k+1})$$

**Critical
step size**
 $\rho = (\rho_1, \dots, \rho_m)$

Hand-tuned step size

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

Vector step size

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

Tight constraints

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

Never tight

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

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

Otherwise



Balance residuals

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

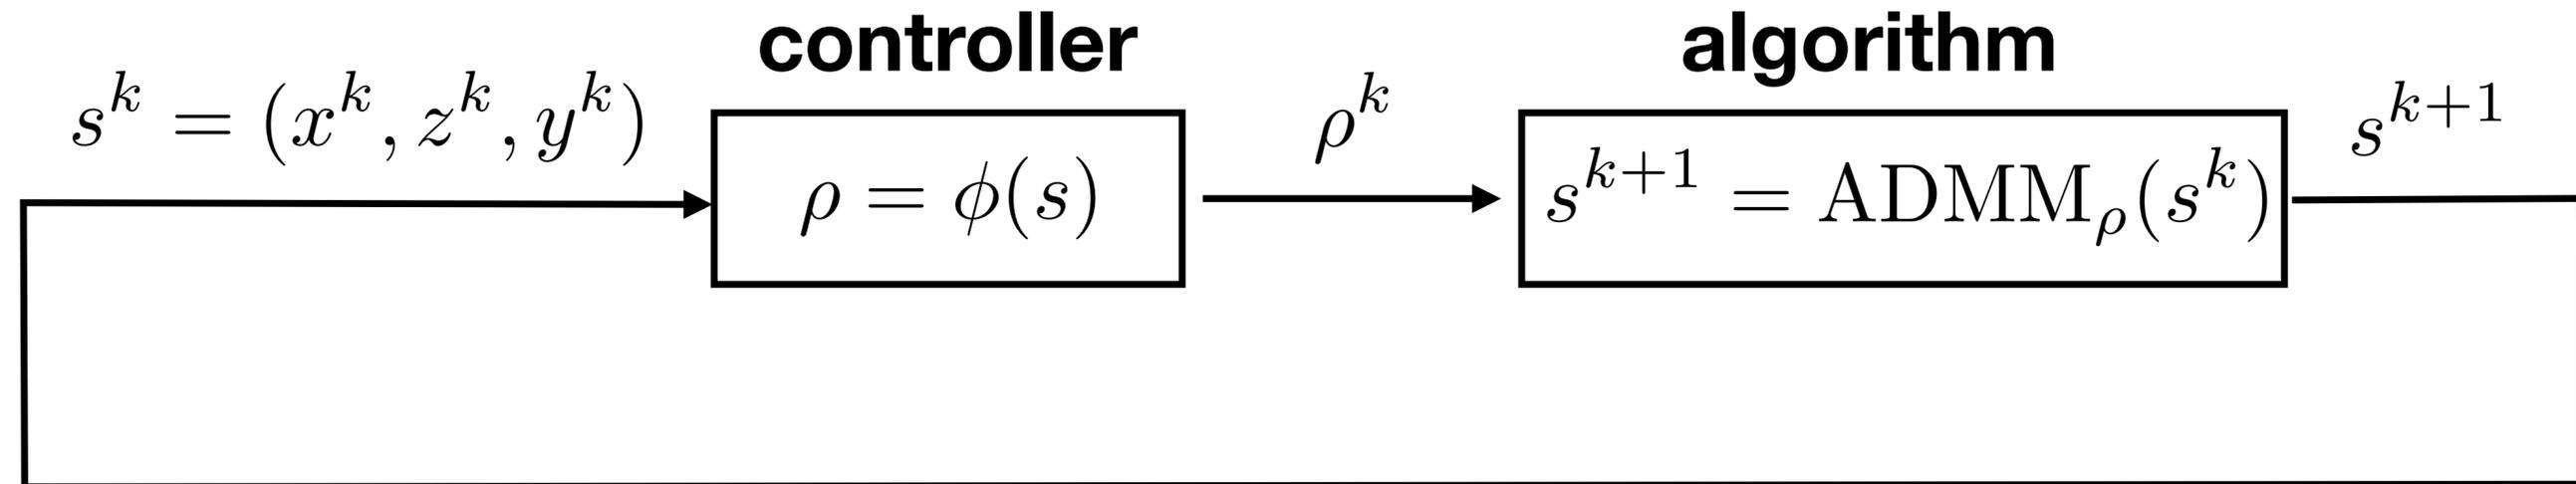
Always tight

$$l_i = u_i \neq \infty$$

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

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

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

Per-constraint update rule

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

Per-constraint state

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

slacks

infeasibility

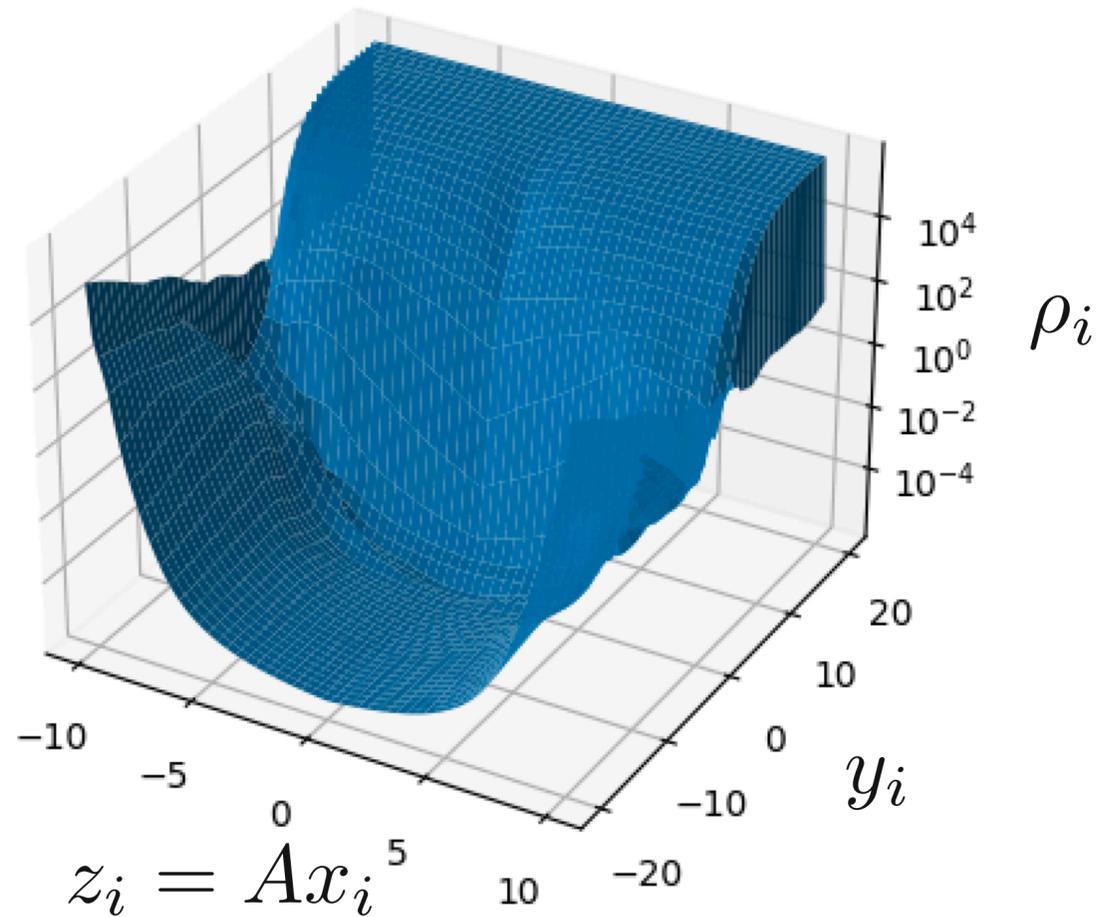
dual variable

Generalize to
different
dimensions

Low-dimensional
state per
constraint

Small NN
policy
 $\phi_c(s_i)$

Visualize learned policy



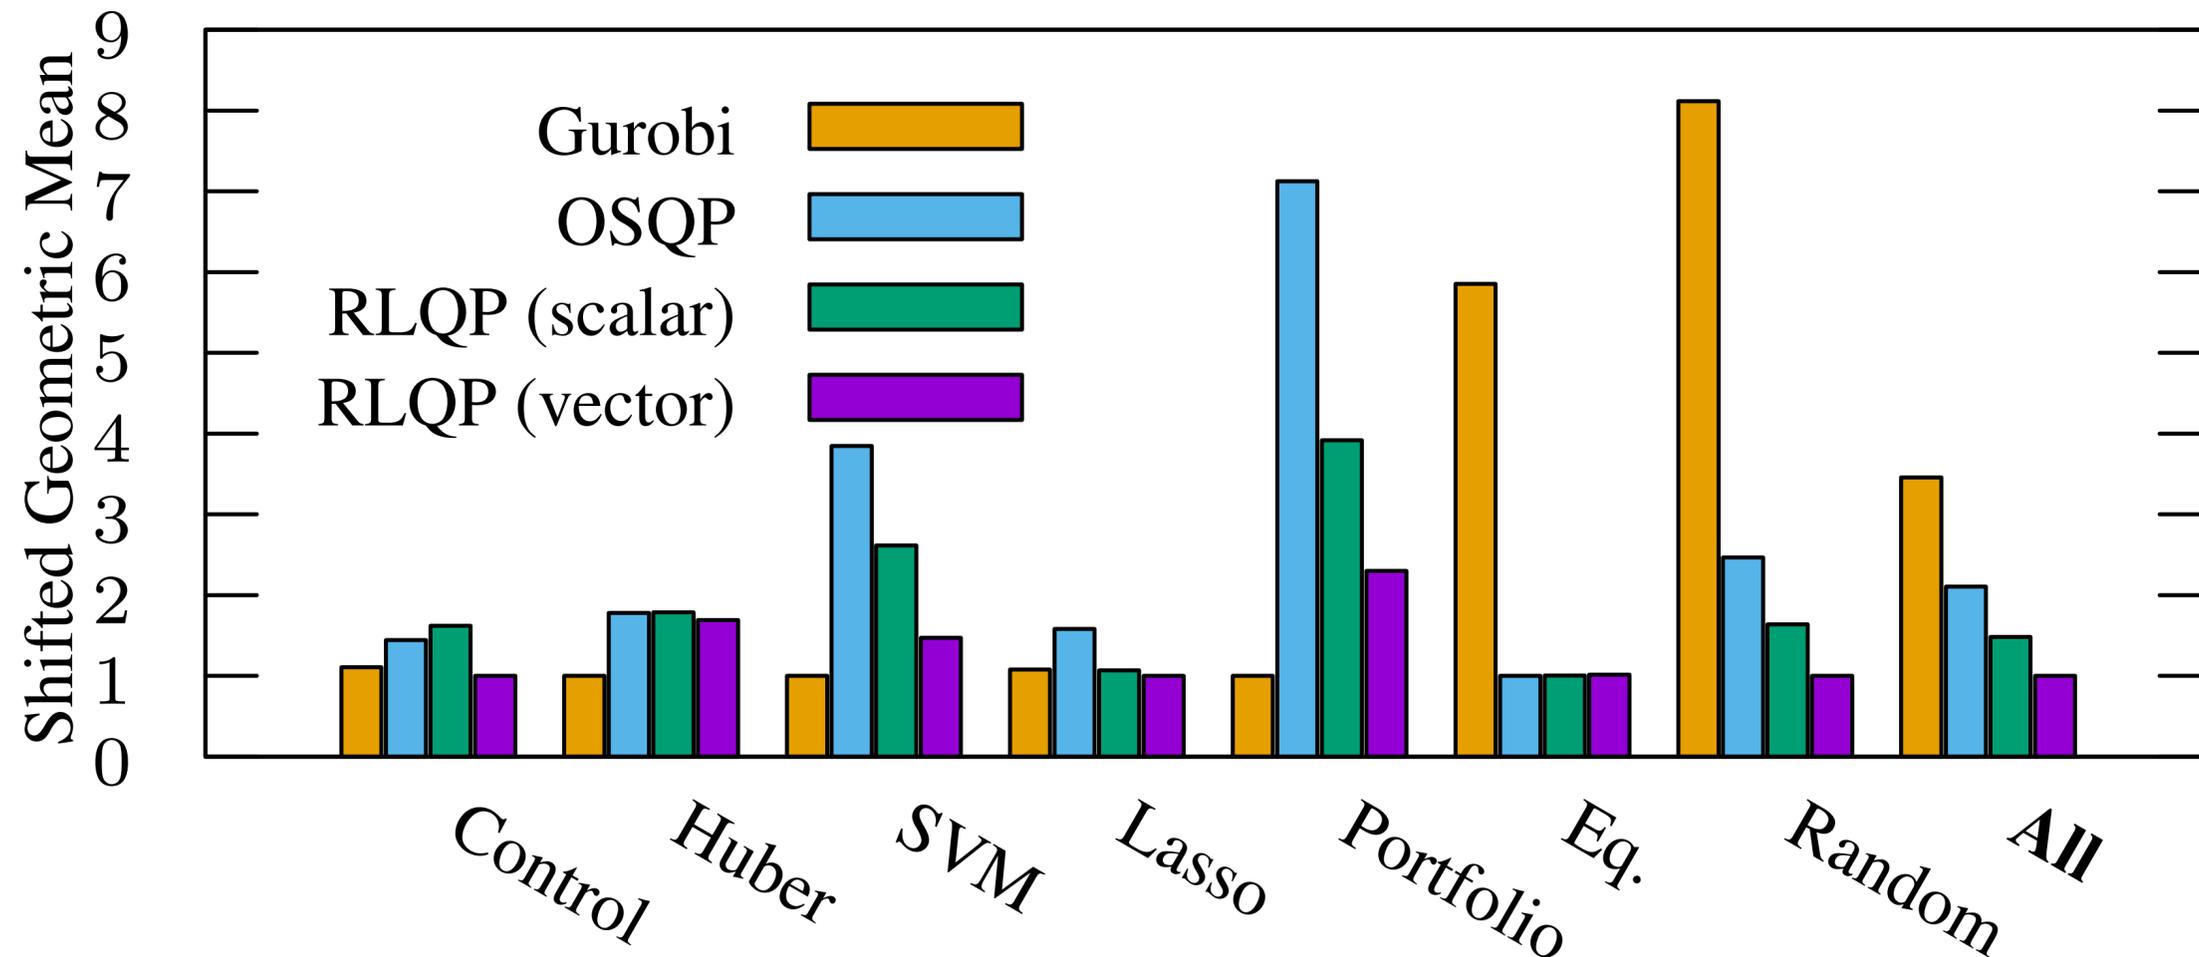
Interpretable policy

High step size ρ_i
when we reach bounds

$$z_i \approx l_i \text{ and } z_i \approx u_i.$$

Performance with step size learning

Timings for high-accuracy convergence criteria



Up to 3x faster than Gurobi

Conclusions

Integrate RL and ADMM to dynamically tune parameters

- Faster convergence
- Very low overhead
- Interpretable policy



stellato.io



bstellato@princeton.edu



github.com/bstellato



@b_stellato

References

[[OSQP: An Operator Splitting Solver for Quadratic Programs.](#)

Stellato, Banjac, Goulart, Bemporad, and Boyd. Mathematical Programming Computation 2020]

[[Accelerating Quadratic Optimization with Reinforcement Learning.](#)

Ichnowski, Jain, Stellato, Banjac, Luo, Gonzalez, Stoica, Borrelli, and Goldberg. NeurIPS 2021]