TRON

A Fast Solver for Trajectory Optimization with Non-Smooth Cost Functions

Anirudh Vemula **J. Andrew Bagnell**
Robotics Institute, Carnegie Mellon University

Success of Trajectory Optimization

Video from [Williams et. al. 2017]

Trajectory Optimization

$$\min_{x_{0:T}, u_{0:T-1}} \quad \ell_T(x_T) + \sum_{t=0}^{T-1} \ell_t(x_t, u_t)$$

$$\text{subject to} \quad x_{t+1} = \kappa(x_t, u_t)$$

- Exploit smoothness of $\{\ell_t\}_{t=1}^{T}$
- Compute gradients and use efficient nonlinear programming methods
- E.g. Newton's method which has quadratic convergence!

But what if the cost functions are not smooth?

Collision Avoidance

$$\ell_t(x_t, u_t) = \max(0, -d(x_t))$$

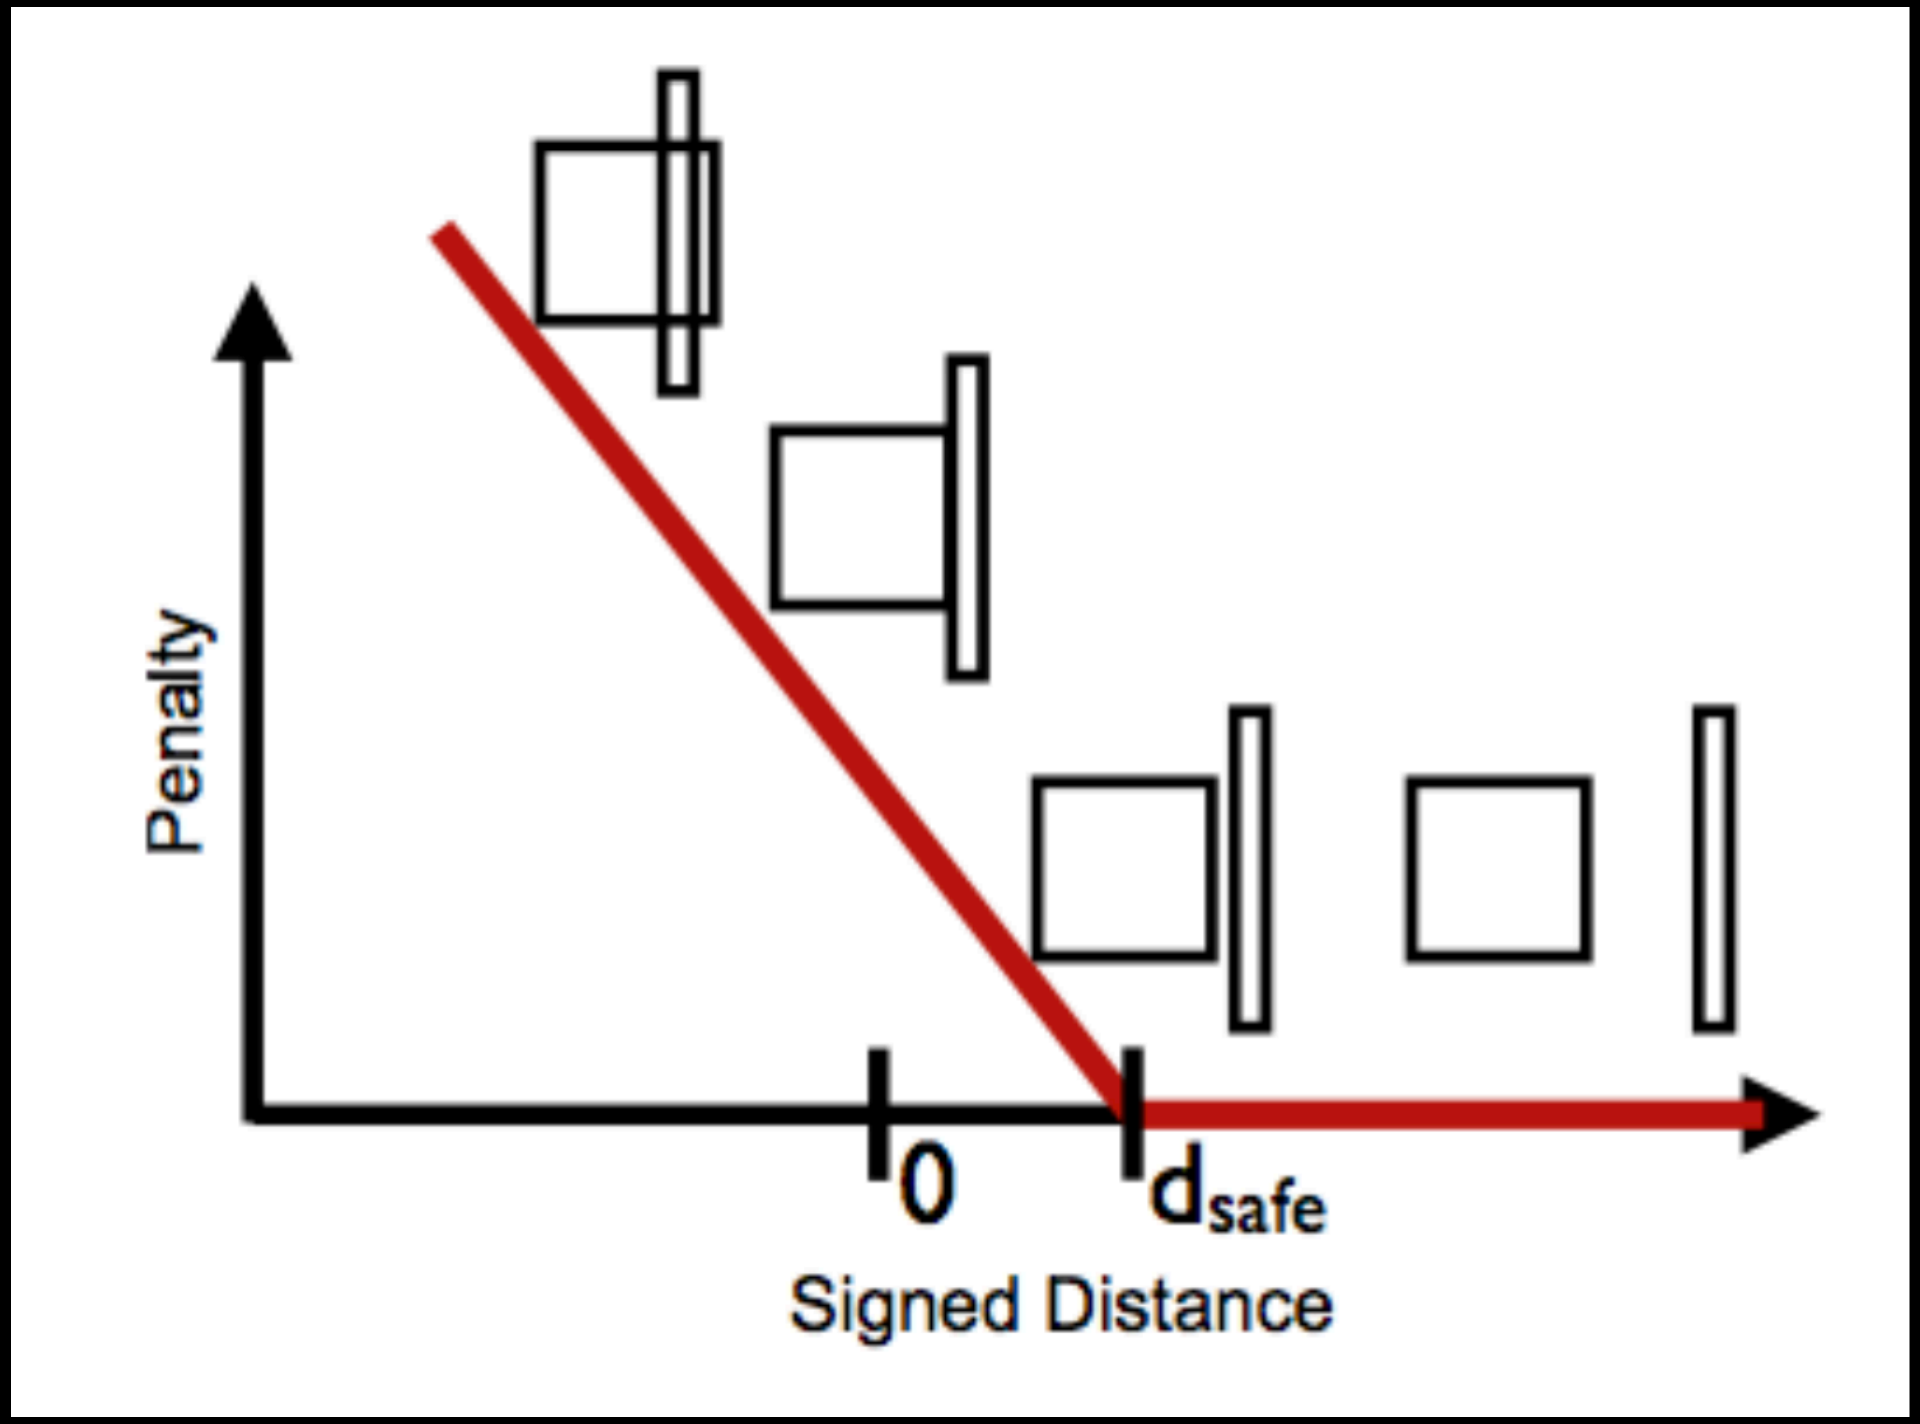

Sparsity in Control

$$\ell_t(x_t, u_t) = \|u_t\|_1$$

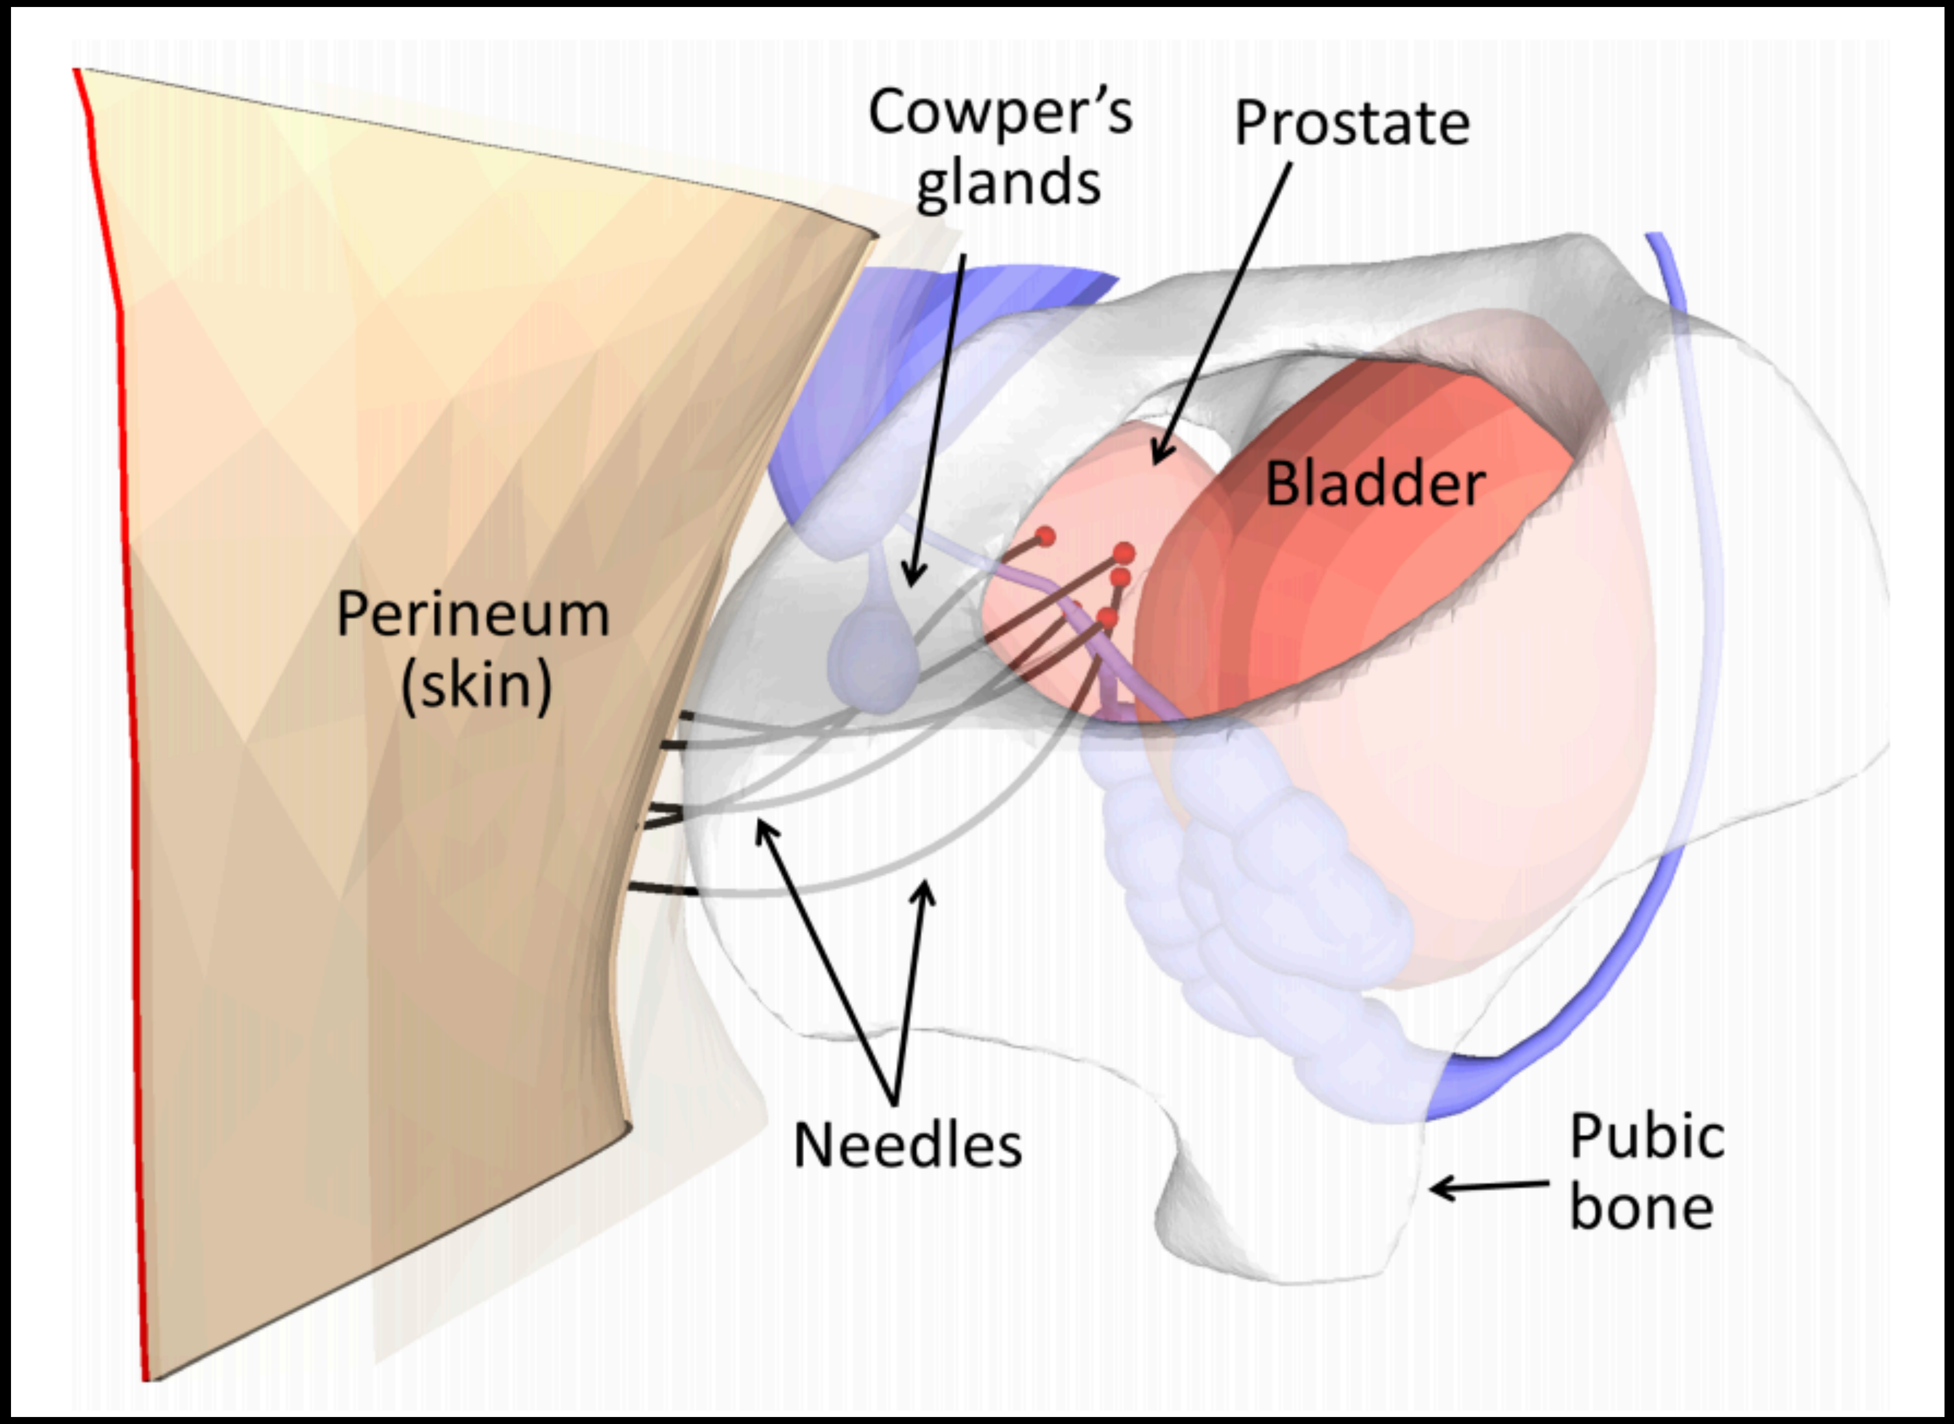

Bang-Off-Bang Control for Satellite Rendezvous

$$\ell_t(x_t, u_t) = \|u_t\|_1$$

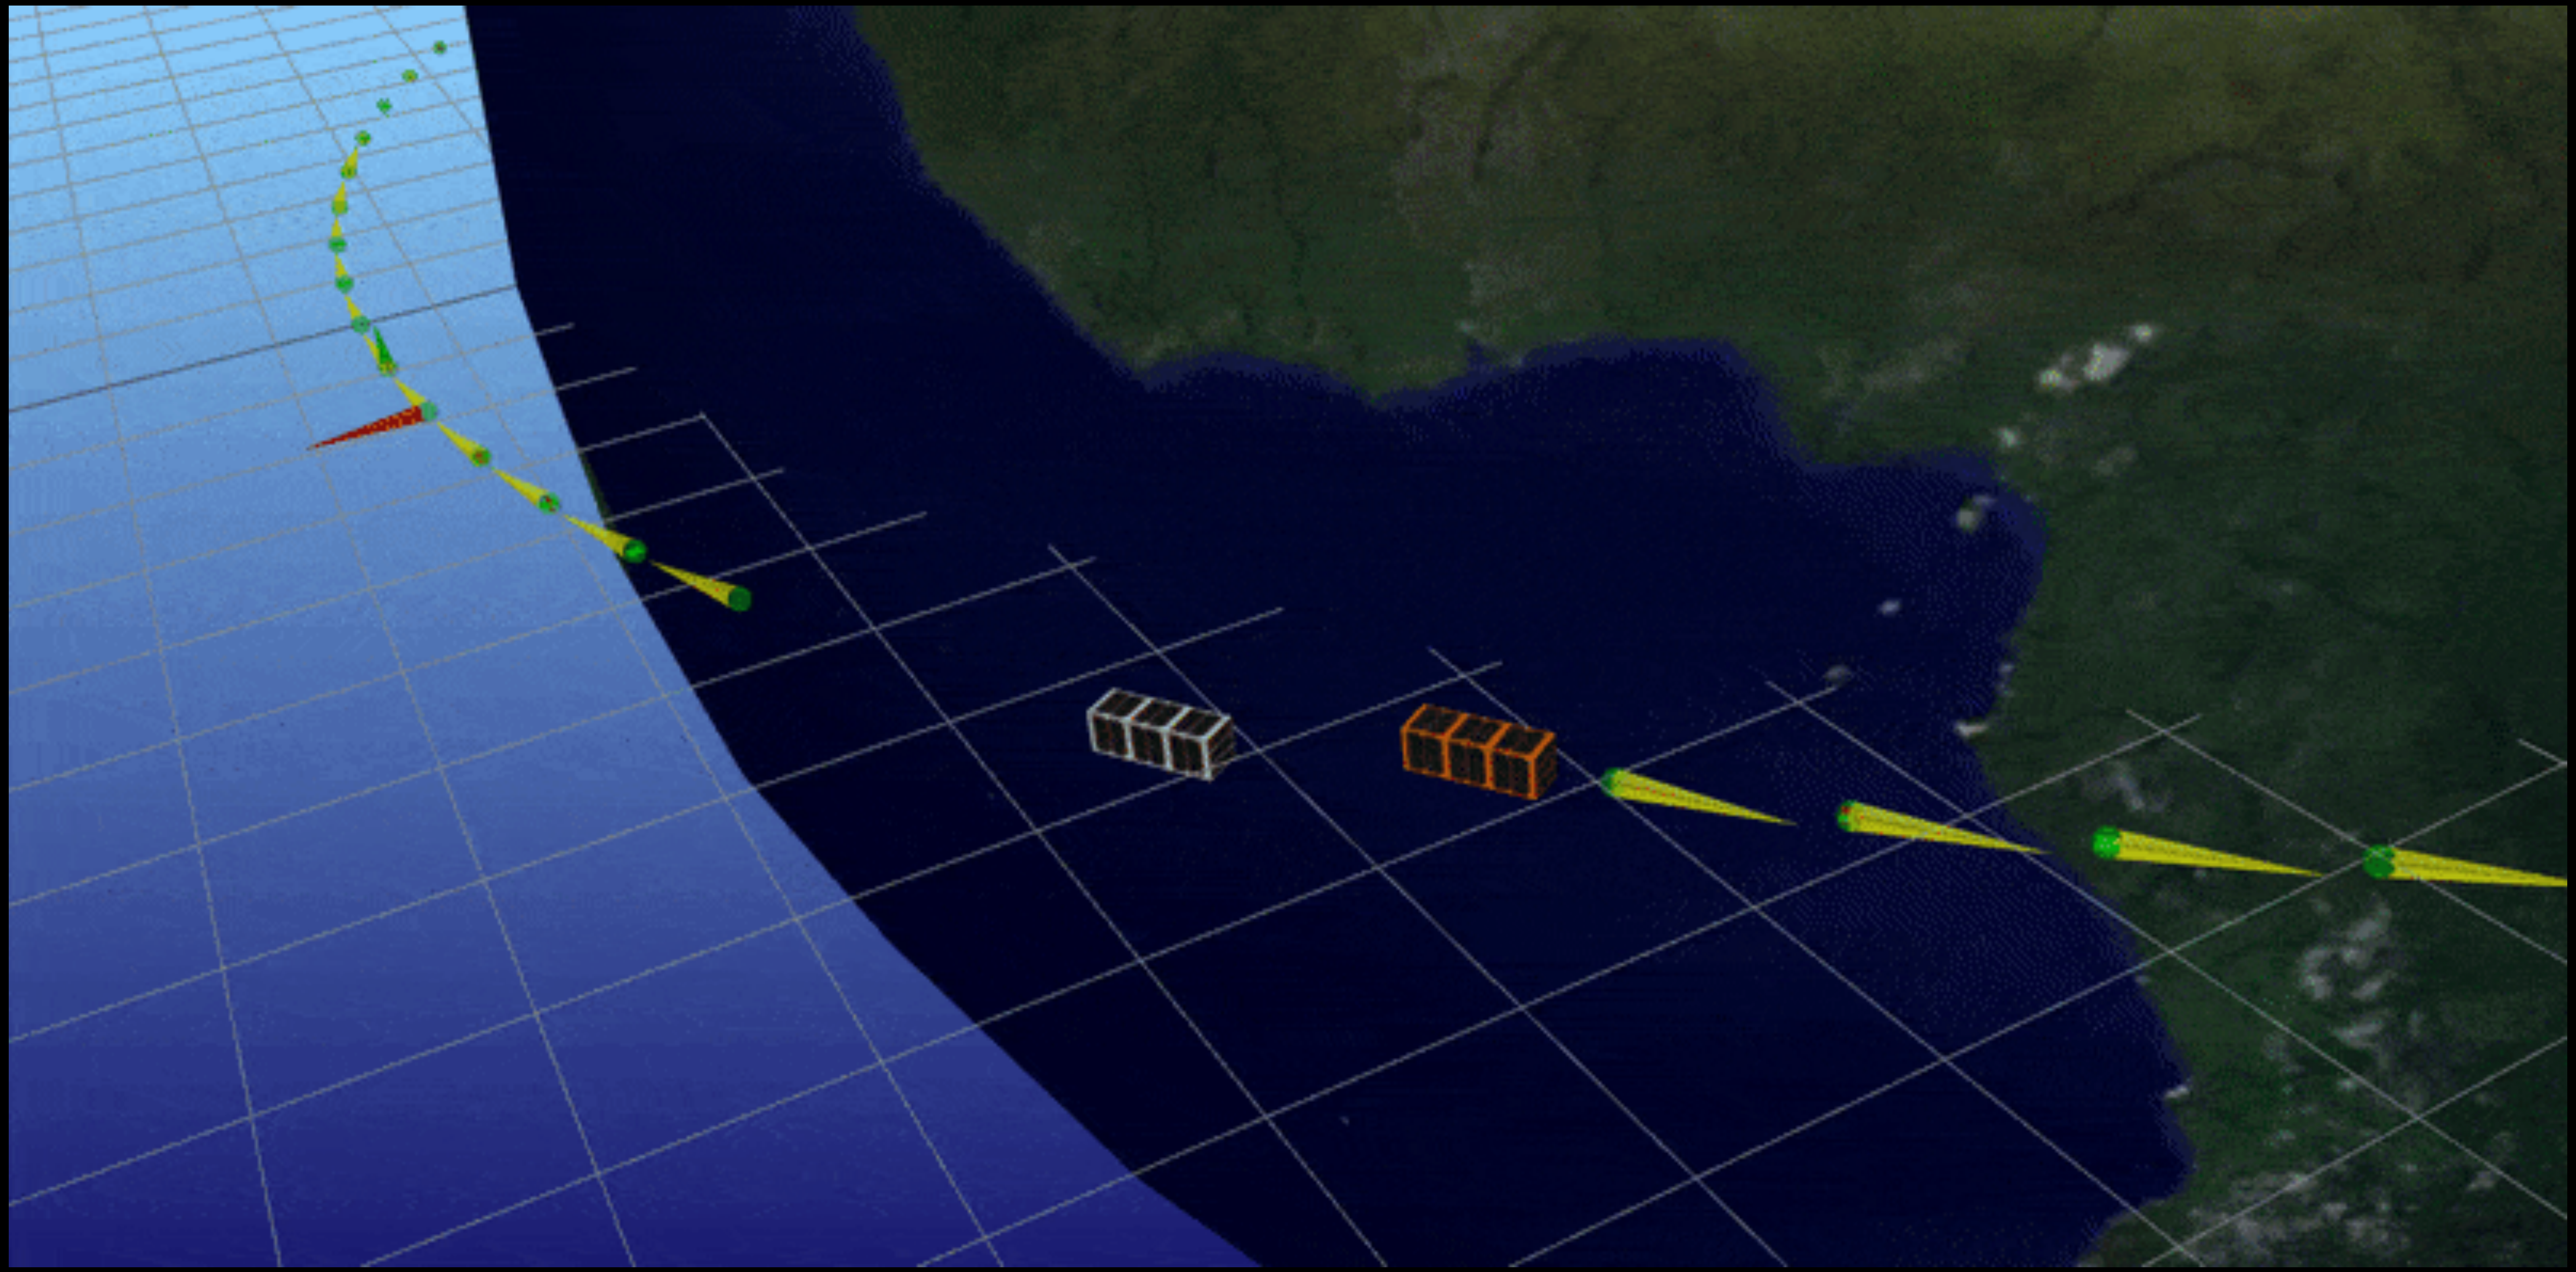

Video from [Le Cleac'h and Manchester 2019]

Structured Non-Smooth Cost Functions

$$\ell_t(x_t, u_t) = f_t(x_t, u_t) + \sum_{i=1}^{M} \max\{g_t^i(x_t, u_t), \bar{g}_t^i(x_t, u_t)\}$$

- Functions $\{f_t, g_t, \bar{g}_t\}_{t=1}^{T}$ are all twice-differentiable and convex
- $\max$ operator makes the resulting cost function non-differentiable
- All previous examples conform to this structure
- Non-smooth cost function with smooth components

Reduction to Simple Formulation

$$\min_{x_{0:T}, u_{0:T-1}} \quad \ell_T(x_T) + \sum_{t=0}^{T-1} \ell_t(x_t, u_t)$$

subject to $\quad x_{t+1} = \kappa(x_t, u_t)$

$$\ell_t(x_t, u_t) = f_t(x_t, u_t) + \sum_{i=1}^{M} \max\{g_t^i(x_t, u_t), \bar{g}_t^i(x_t, u_t)\}$$

Functions f, g_1, g_2 are convex and smooth

$$\min_{y \in Y} f(y) + \max\{g_1(y), g_2(y)\}$$

Equivalent Problem

$$\min_{y \in Y} f(y) + \max\{g_1(y), g_2(y)\}$$

$$\min_{y \in Y} f(y) + \max_{\theta \in \Delta_2}(\theta_1 g_1(y) + \theta_2 g_2(y))$$

Δ_2 is the 2D simplex and $\theta = [\theta_1, \theta_2]^T \in \Delta_2$ implies $\theta_1 + \theta_2 = 1$ and $\theta_1, \theta_2 \geq 0$

Equivalence not useful as the term $\max_{\theta \in \Delta_2} \theta_1 g_1(y) + \theta_2 g_2(y)$ is non-smooth in y

Regularized Objective

$$\min_{y \in Y} f(y) + \max_{\theta \in \Delta_2} (\theta_1 g_1(y) + \theta_2 g_2(y))$$

$$\min_{y \in Y} f(y) + \max_{\theta \in \Delta_2} (\theta_1 g_1(y) + \theta_2 g_2(y) - \eta^k \mathtt{KL}(\theta || \theta^{k-1}))$$

At iteration k, add regularization term penalizing deviation from previous solution θ^{k-1}

$$\mathtt{KL}(\theta || \theta^{k-1}) = \theta_1 \log \frac{\theta_1}{\theta_1^{k-1}} + \theta_2 \log \frac{\theta_2}{\theta_2^{k-1}}$$

A Two-Player Min-Max Game

$$\min_{y \in Y} f(y) + \max_{\theta \in \Delta_2} (\theta_1 g_1(y) + \theta_2 g_2(y) - \eta^k \mathrm{KL}(\theta || \theta^{k-1}))$$

Take a small gradient step to play

$y \in Y$ that decreases objective

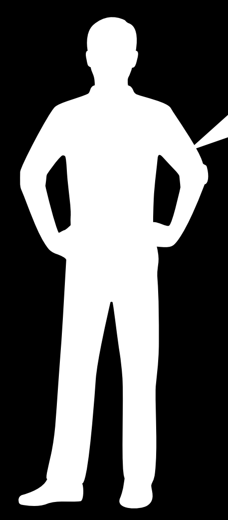

Player

Play $\theta \in \Delta_2$

that maximizes objective

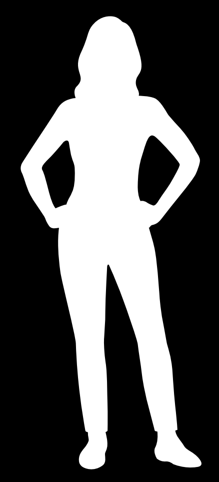

Adversary

Player and Adversary Updates

$$\min_{y \in Y} f(y) + \max_{\theta \in \Delta_2}(\theta_1 g_1(y) + \theta_2 g_2(y) - \eta^k \texttt{KL}(\theta || \theta^{k-1}))$$

- The inner maximization can be solved in closed form!

$$\theta^k = \frac{\theta^{k-1} \exp\left(\frac{g(y)}{\eta^k}\right)}{\sum_{i=1}^{2} \theta_i^{k-1} \exp\left(\frac{g_i(y)}{\eta^k}\right)}$$

- Simplifying, we get

$$\min_{y \in Y} f(y) + \eta^k \log(\theta_1^{k-1} \exp(\frac{g_1(y)}{\eta^k}) + \theta_2^{k-1} \exp(\frac{g_2(y)}{\eta^k}))$$

Smooth and twice-differentiable in y

TRON: Application to Trajectory Optimization

$$\min_{x_{0:T}, u_{0:T}} \quad \sum_{t=0}^{T} f_t(x_t, u_t) + \eta^k \log(\theta_t^{k-1} \exp(\frac{g_t(x_t, u_t)}{\eta^k}) + \bar{\theta}_t^{k-1} \exp(\frac{\bar{g}_t(x_t, u_t)}{\eta^k}))$$

subject to $\quad x_{t+1} = \kappa(x_t, u_t)$

- Solve the above trajectory optimization problem with smooth, twice-differentiable cost function using iLQR
- After obtaining $x_{0:T}, u_{0:T-1}$ from iLQR, update θ
- Repeat until a fixed number of iterations

TRON: Theoretical Guarantees

$$\min_{x_{0:T}, u_{0:T-1}} \quad \ell_T(x_T) + \sum_{t=0}^{T-1} \ell_t(x_t, u_t)$$

$$\text{subject to} \quad x_{t+1} = \kappa(x_t, u_t)$$

- When dynamics $x_{t+1} = \kappa(x_t, u_t)$ are linear, converges to global minimum
- With non-linear dynamics, convergence to stationary point
- Requires $\eta^k \to 0$ as $k \to \infty$

Control for a Differential Drive Robot

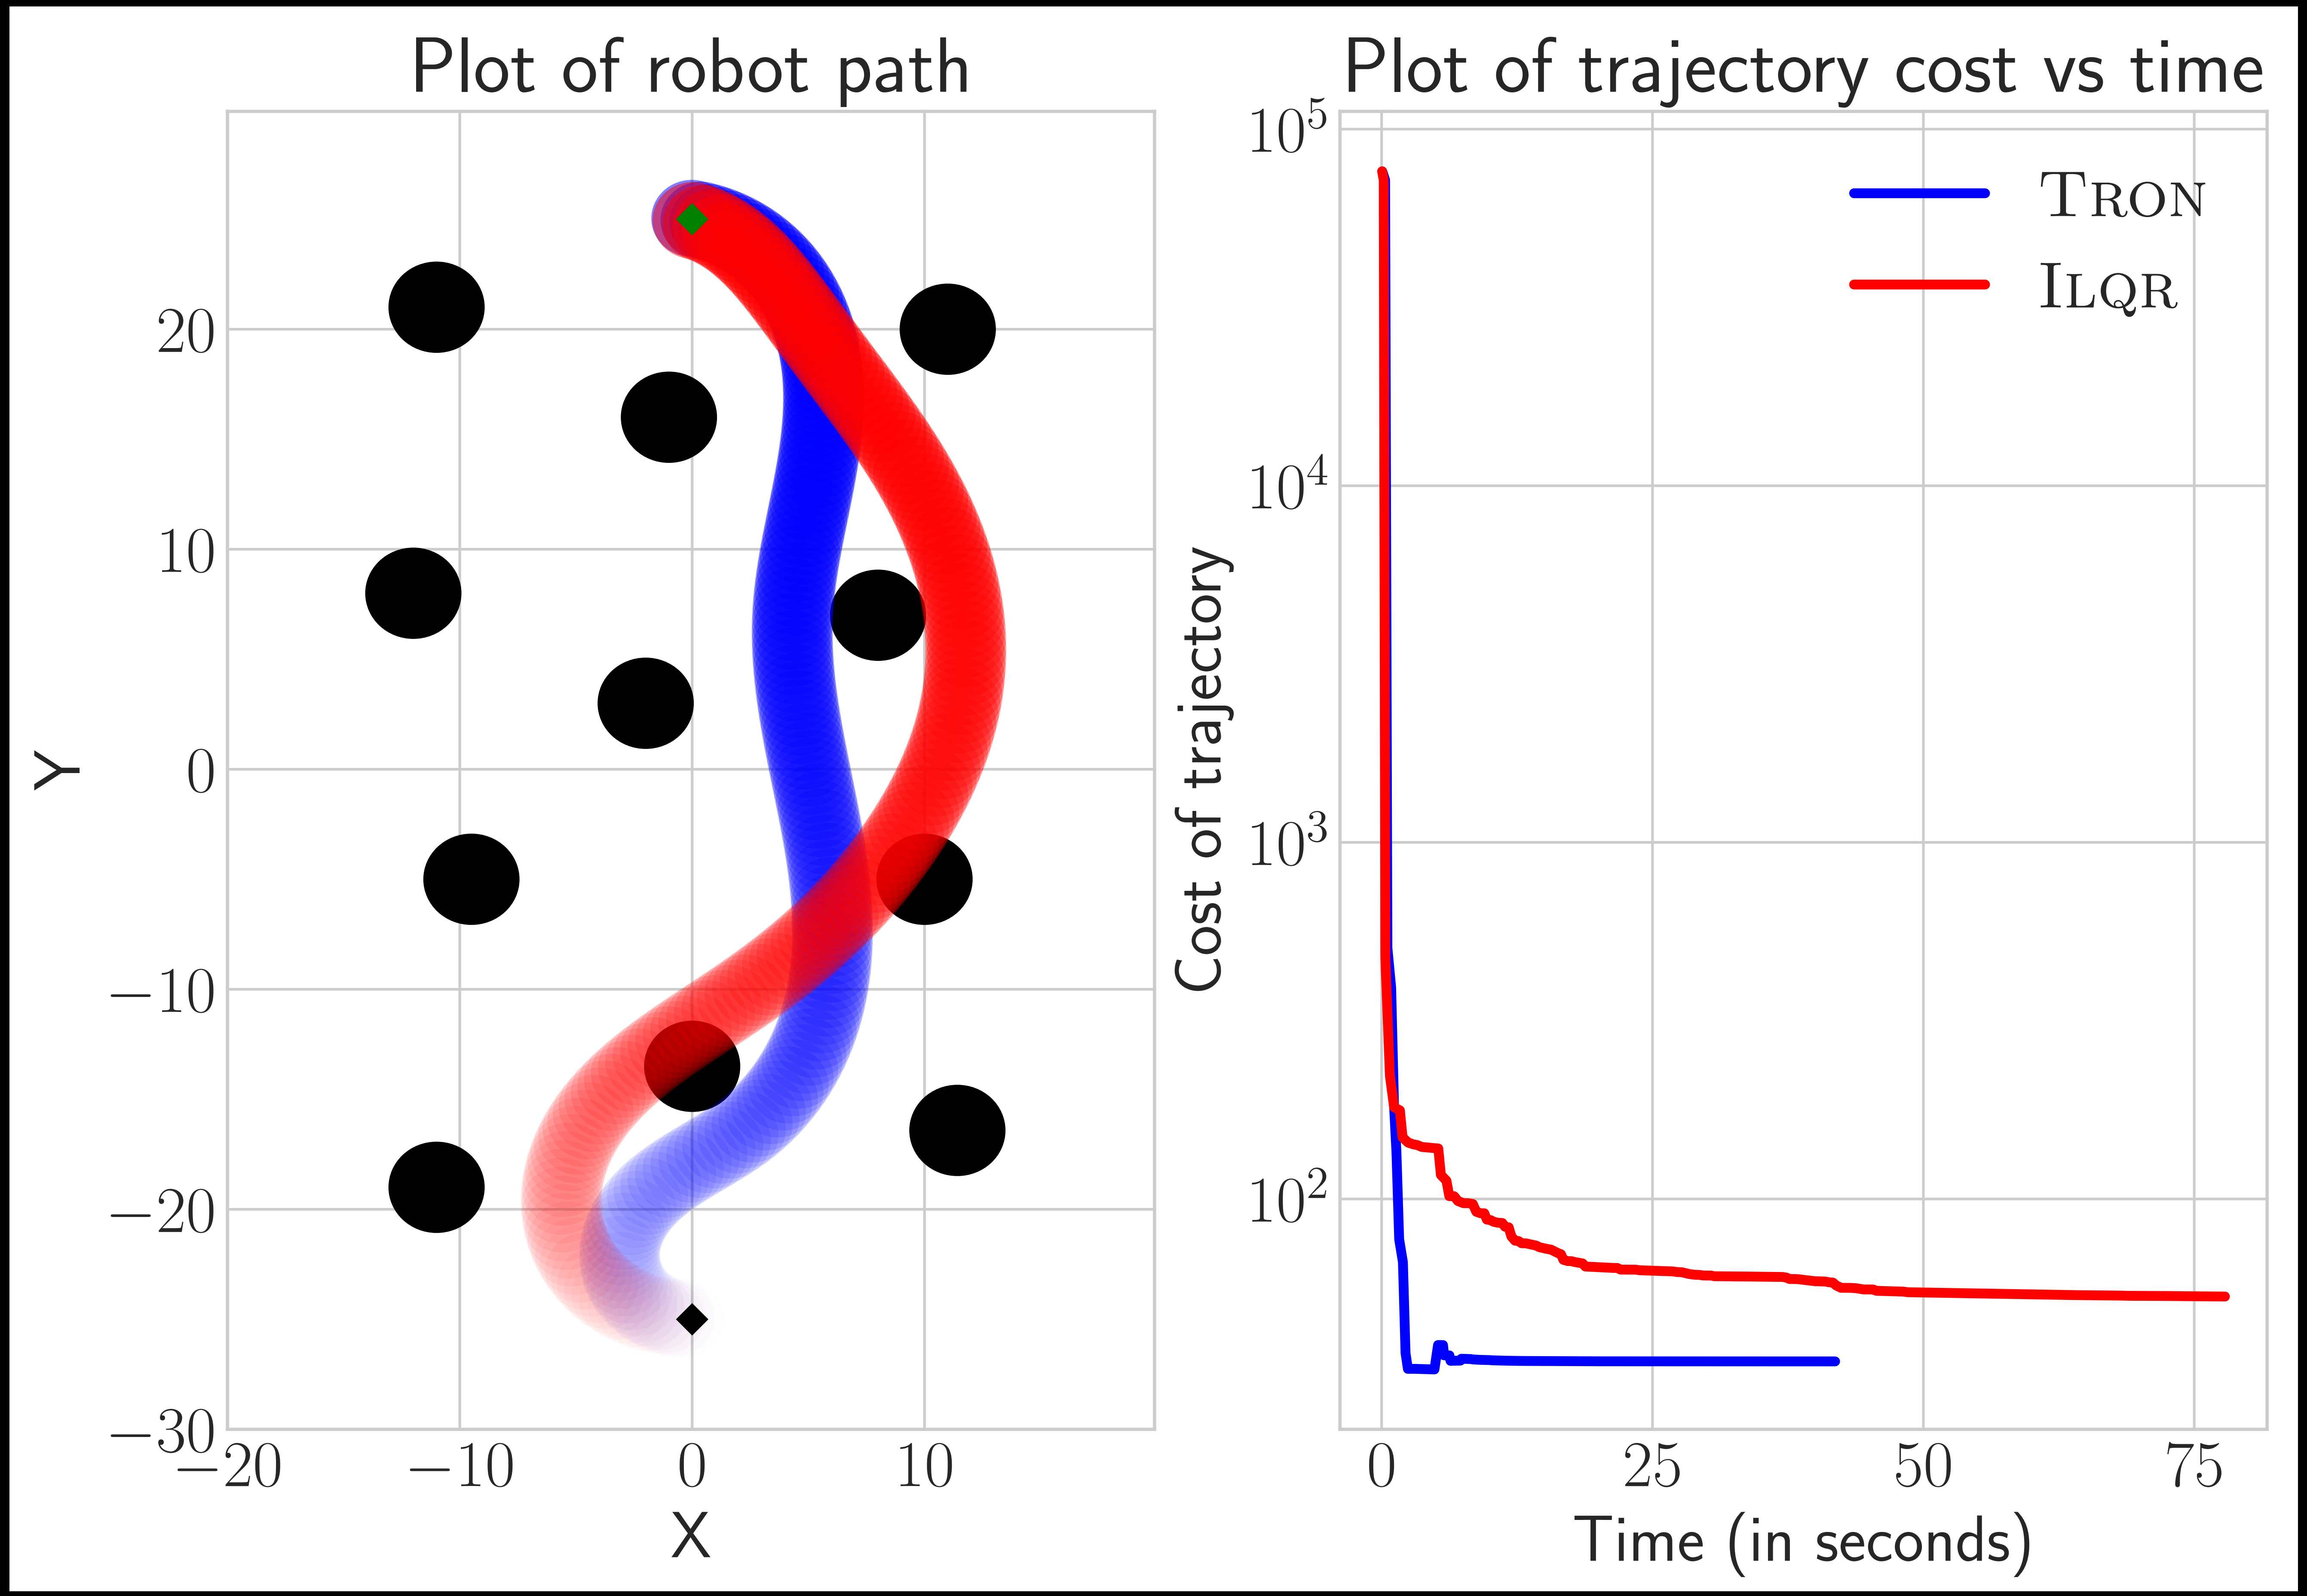

$$\ell_t(x_t, u_t) = \max(0, -d(x_t))$$

Sparse Control for a Surgical Needle

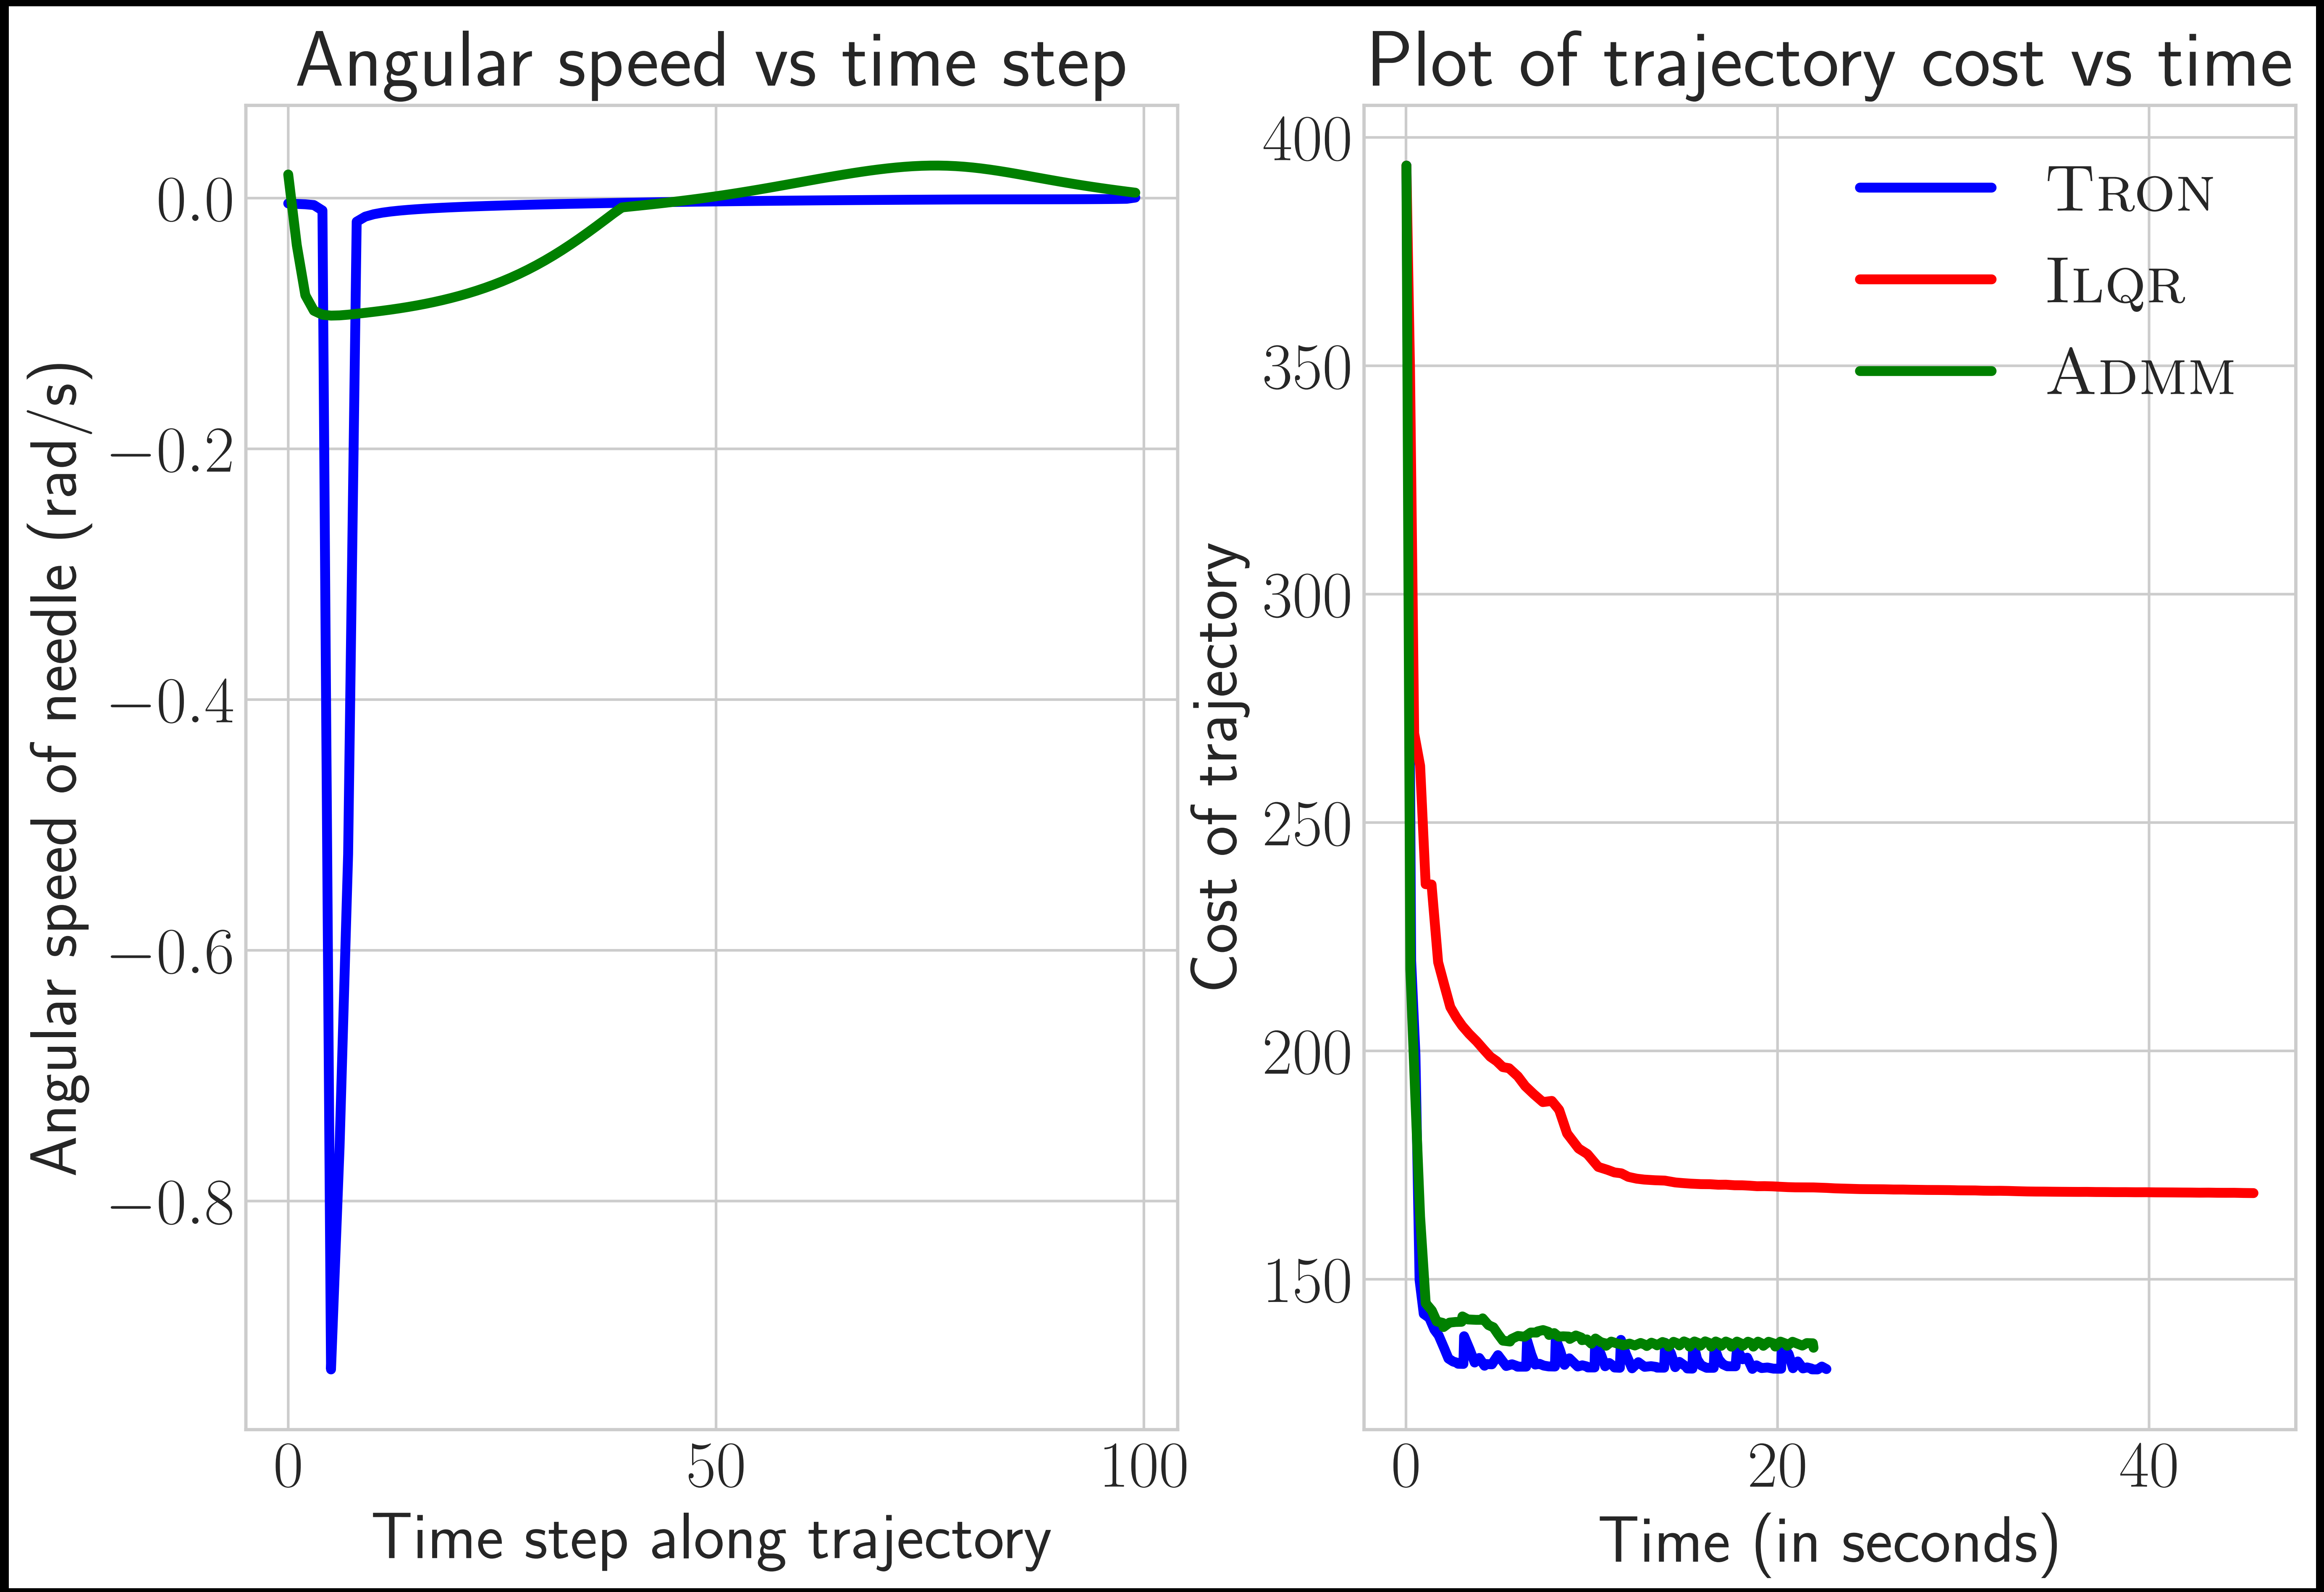

$$\ell_t(x_t, u_t) = \|u_t\|_1$$

Bang-Off-Bang Control for Satellite Rendezvous

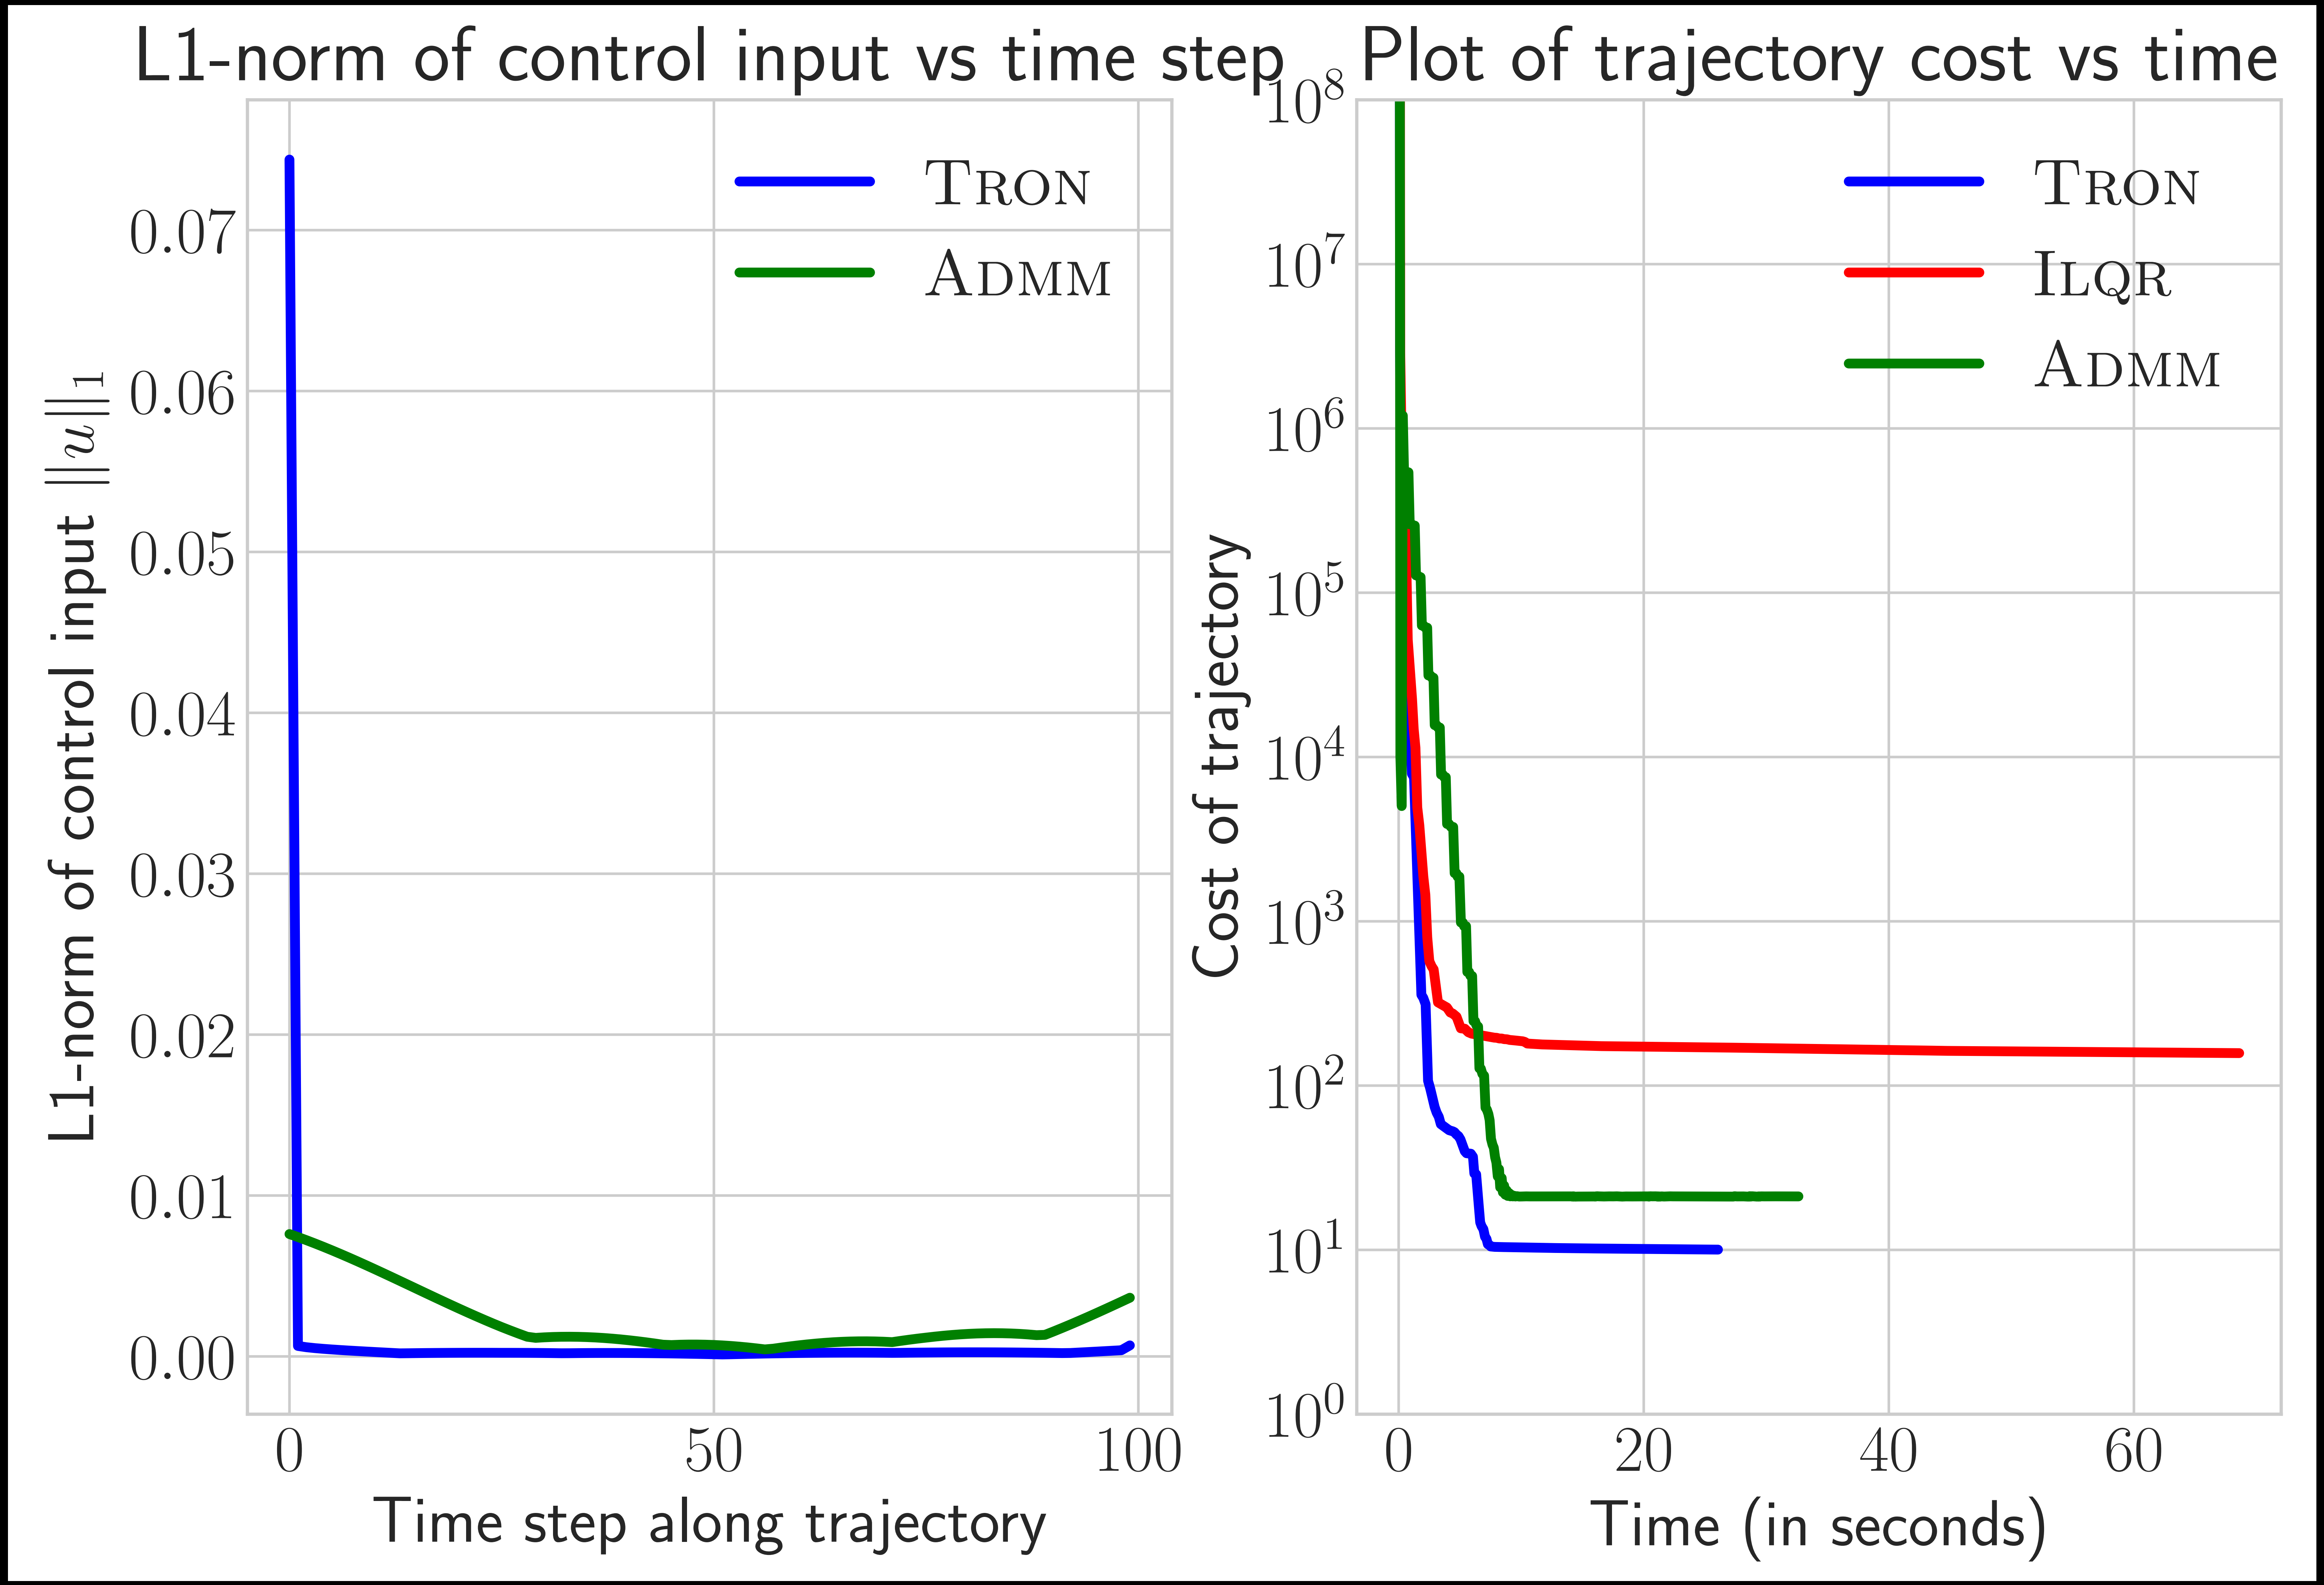

TRON works really well for trajectory optimization problems with structured non-smooth cost functions

Thank you.

The code is available at https://github.com/vvanirudh/TRON
Full paper is available at https://arxiv.org/abs/2003.14393