Extended Convex Lifting for Policy Optimization in Control

Yang Zheng

Assistant Professor,
ECE Department, UC San Diego

Department Colloquium, Department of Mathematics and Statistics,
Queen's University

March 27, 2026

UC San Diego
JACOBS SCHOOL OF ENGINEERING
Electrical and Computer Engineering

Scalable Optimization and
Control (SOC) Lab
https://zhengy09.github.io/soclab.html

Collaborators

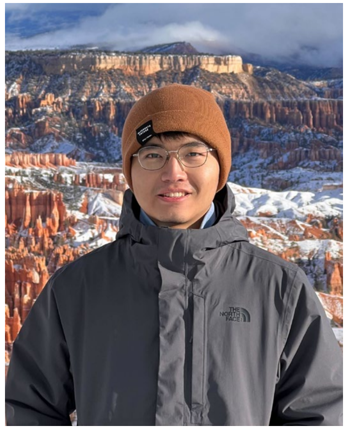

Chih-Fan (Rich) Pai
UC San Diego

Yuto Watanabe
UC San Diego

Feng-Yi Liao
UC San Diego

Yujie Tang
Peking University

Na Li
Harvard University

Supported by

NSF ECCS-2154650,
CMMI-2320697,
CAREER-2340713

Introduction

Control theory is the study of dynamical systems with inputs

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

$x_t \in \mathbb{R}^n$ is called the system state

$u_t \in \mathbb{R}^m$ is called the control input

$y_t \in \mathbb{R}^p$ is called the system output

Optimal control

$$\min_{u_t} \quad \sum_{t=0}^{T} c_t(x_t, u_t)$$

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

Control as **Optimization** over time constrained by dynamics

Introduction: a simple control example

Done flying: position control

Target state

Newton's Laws

$$p_{t+1} = p_t + v_t \Delta_t$$

$$v_{t+1} = v_t + a_t \Delta_t$$

$$m a_t = u_t$$

State $x_t = \begin{bmatrix} p_t \\ v_t \end{bmatrix}$

Optimal control

$$\min_{u_t} \quad \sum_{t=0}^{T} \left(\|x_t\|^2 + R \cdot u_t^2 \right)$$

$$\text{subject to} \quad x_{t+1} = \begin{bmatrix} 1 & \Delta_t \\ 0 & 1 \end{bmatrix} x_t + \begin{bmatrix} 0 \\ \frac{\Delta_t}{m} \end{bmatrix} u_t$$

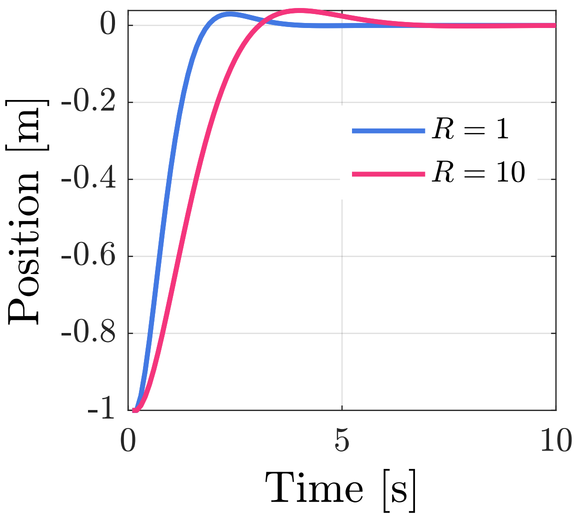

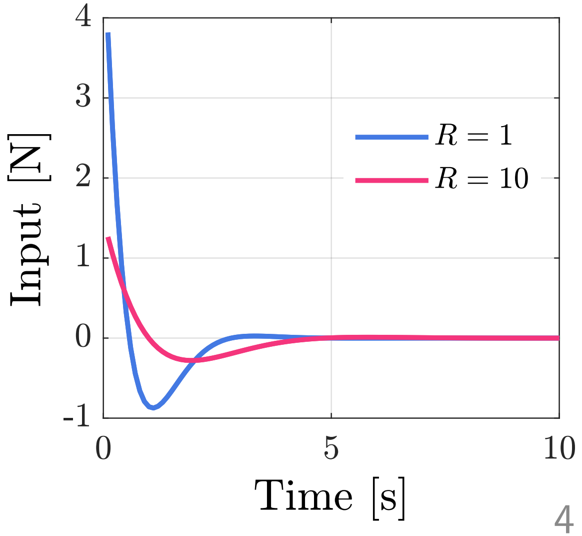

Many control applications

Temperature/HVAC control

Robotics

Drone formations

Transportation

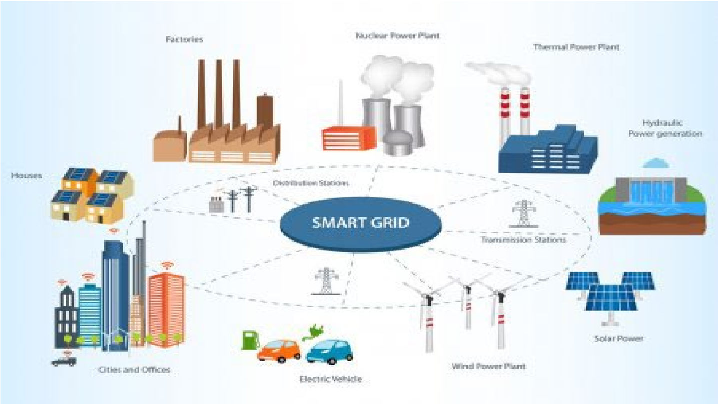

Smart grid

Control theory offers the principled use of **feedback loops** and **algorithms** to drive a dynamic system to its goal.

Today's topic: Policy optimization

Feedback control

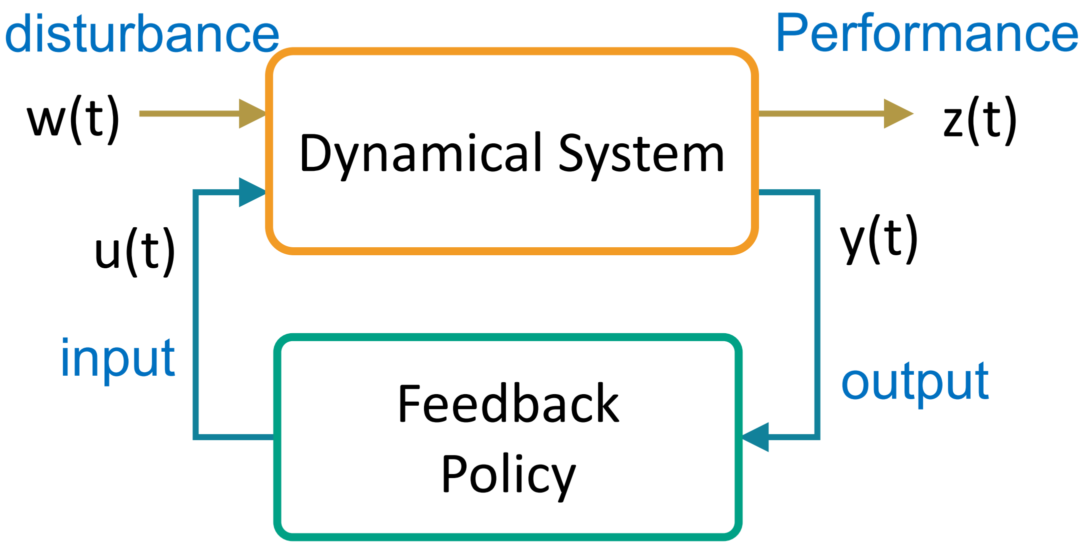

Policy parametrization

Policy Optimization

$$\min_{\mathsf{K}} \quad J(\mathsf{K})$$
$$\text{s.t.} \quad \mathsf{K} \in \mathcal{C}$$

- Traditional methods are often **model-based;**
- Well-established theory on **Optimal and Robust control** etc.
- **Tools:** Riccati equations, linear matrix inequalities, Dynamic programming, etc.

- Use direct policy updates (**model-free**),
$$K_{t+1} = K_t - \alpha_t \nabla J(K_t)$$
- Become very popular in both academia and practice, from game playing, **robotics** to **drone flying**, etc.

Today's topic: Policy optimization

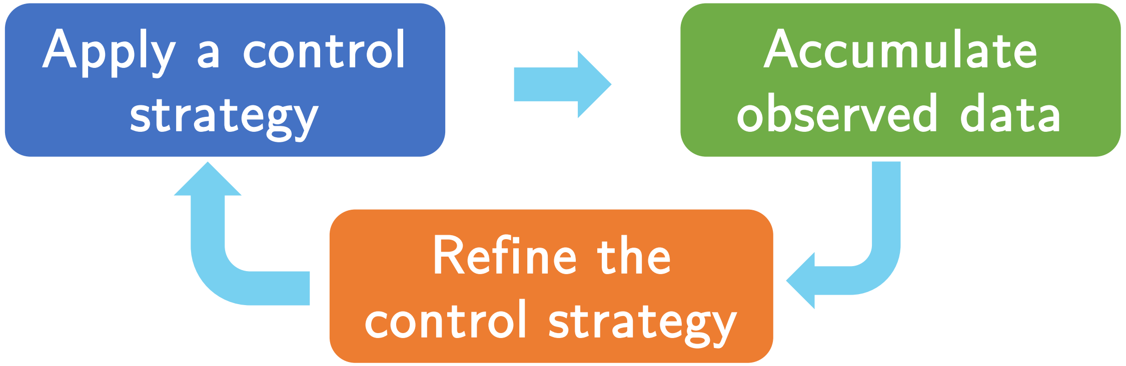

$$K_{t+1} = K_t - \alpha_t \nabla J(K_t)$$

$$\min_{\mathsf{K} \in \mathcal{C}} \quad J(\mathsf{K})$$

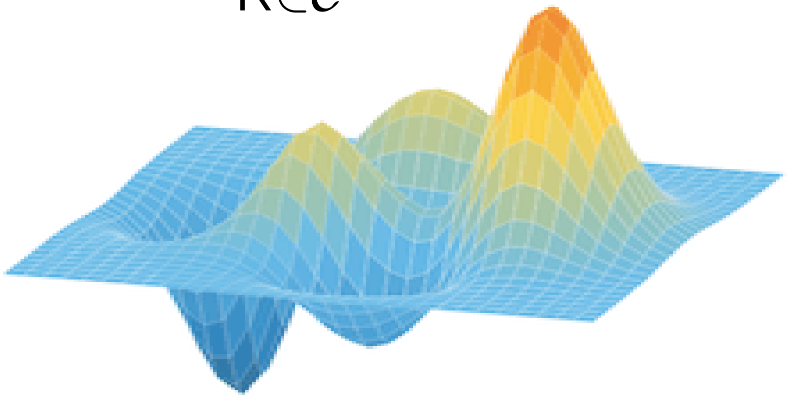

Opportunities

- Easy-to-implement
- Scalable to high-dimensional problems
- Enable model-free search with rich observations

Challenges

- Nonconvex optimization
- Lack of principled algorithms for optimality (e.g., avoiding saddles/local minimizers)
- Hard to obtain theoretical guarantees (e.g., robustness/stability, sample efficiency)

Key Messages

Message 1: Nonconvex Landscapes in Benchmark Control Problems are Benign

- Informally, any (non-degenerate) stationary point is globally optimal!

[Fazel et al., 2018], [Mohammad et al., 2022] [Fatkhullin & Polyak, 2021], [Guo & Hu, 2022], [Moghaddam, Olshevsky & Gharesifard, 2025], [Tang, **Zheng**, & Li, 2023], [**Zheng**, Pai & Tang, 2023] etc.

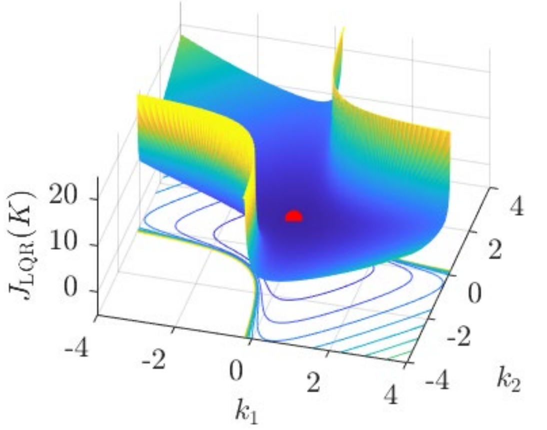

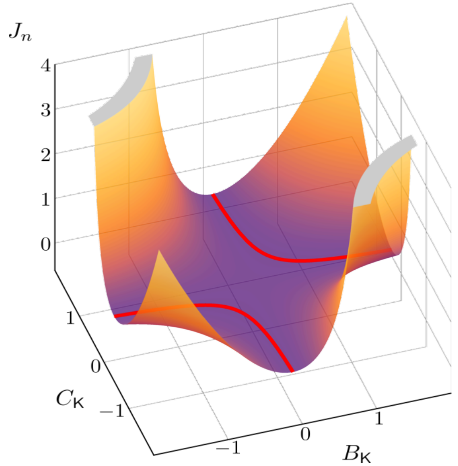

Message 2: A Unified Extended Convex Lifting Framework

- Bridge non-convex policy optimization with convex reformulation!

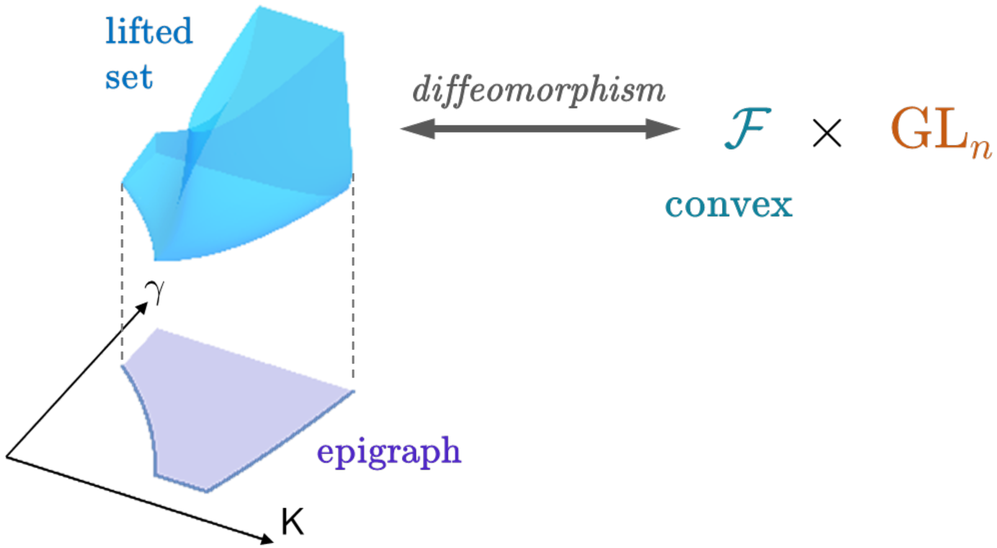

Message 3: Efficient first-order algorithm for (non)-convex optimization

Outline

Geometry:
Benign Nonconvex Landscape

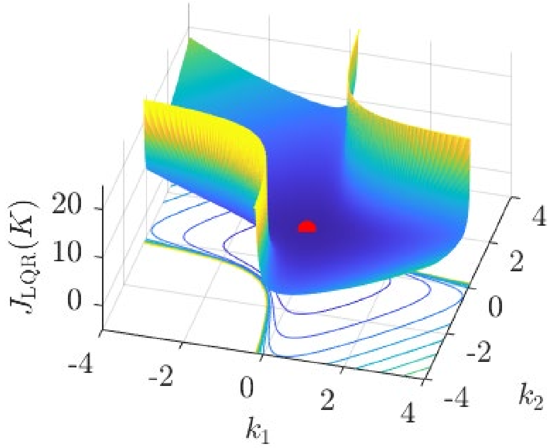

Framework:
Extended Convex Lifting

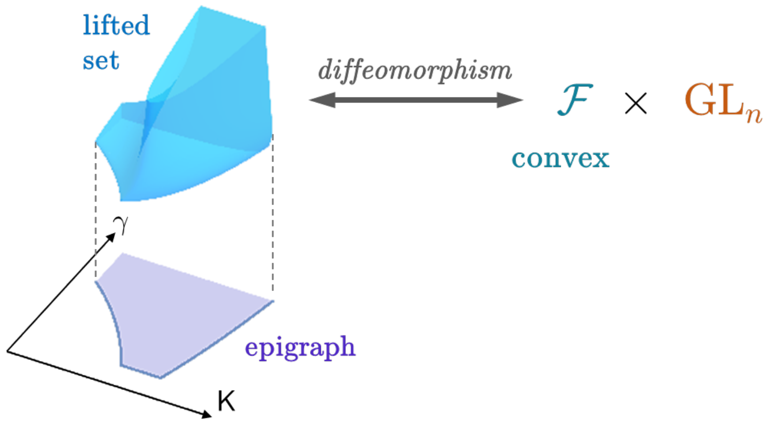

Algorithm:
Proximal Descent Method

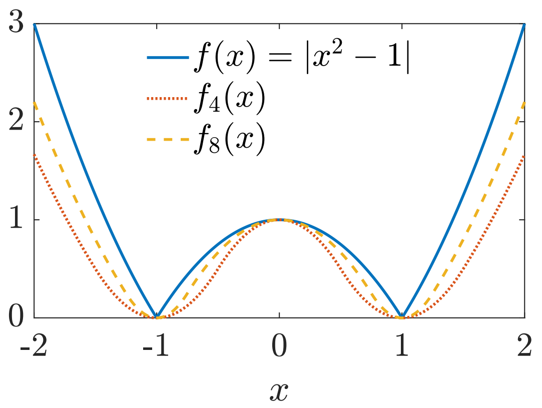

Non-convexity in control

Policy optimization in control is generally **nonconvex**!

$$\min_{\mathsf{K}} \; J(\mathsf{K}) \qquad \text{s.t.} \; \mathsf{K} \in \mathcal{C}$$

$$\min_{K \in \mathbb{R}^{m \times n}} \quad J(K) := \sum_{t=0}^{\infty} (x_t^{\mathsf{T}} Q x_t + u_t^{\mathsf{T}} R u_t) \longrightarrow \text{Quadratic costs}$$

$$\text{subjct to} \quad x_{t+1} = A x_t + B u_t, \longrightarrow \text{Linear dynamics}$$

$$u_t = K x_t \longrightarrow \text{Linear static policy}$$

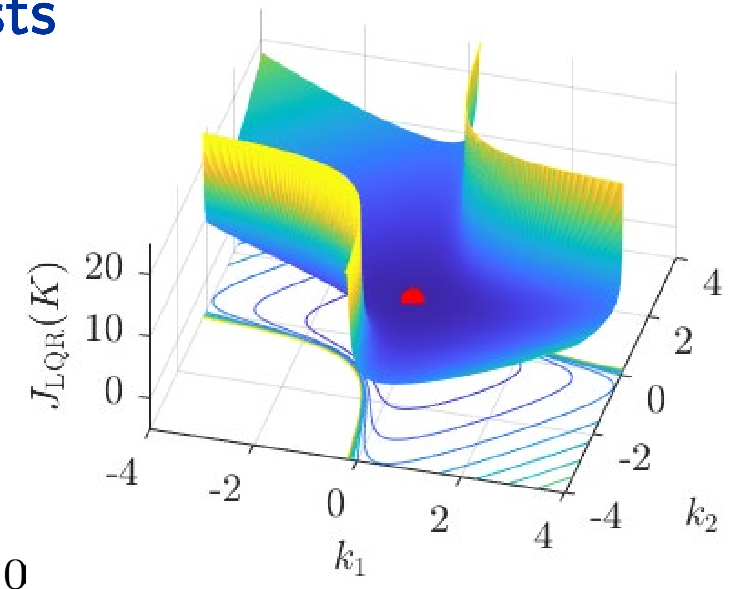

- State trajectories are nonlinear in K

$$x_1 = (A + BK)x_0,$$

$$x_2 = (A + BK)x_1 = (A + BK)^2 x_0$$

$$x_t = (A + BK)^t x_0$$

- The cost function J is nonconvex in K

Non-convexity in control

Policy optimization in control is generally **nonconvex**!

$$\min_{\mathsf{K}} \quad J(\mathsf{K})$$
$$\text{s.t.} \quad \boxed{\mathsf{K} \in \mathcal{C}}$$

- **Stabilization** is already a **non-convex** constraint!
  - ✓ A policy K is said to **stabilize** the system, if the system trajectory converges to the origin.

$$\lim_{t\to\infty} x_t = \lim_{t\to\infty} (A + BK)^t x_0 = 0$$

  - ✓ It requires the spectral radius of the closed-loop matrix less than 1.

$$\rho(A + BK) < 1$$

  - ✓ A simple example: $A = 0, \quad B = I_2$

$$x_{t+1} = u_t \qquad \mathcal{C} = \{K \in \mathbb{R}^{2\times 2} \mid A + BK \text{ is stable}\}$$

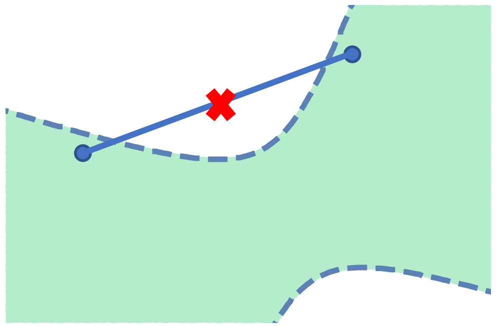

$$K_1 = \begin{bmatrix} 0.5 & 2 \\ 0 & 0.5 \end{bmatrix} \in \mathcal{C}, \quad K_2 = \begin{bmatrix} 0.5 & 0 \\ 2 & 0.5 \end{bmatrix} \in \mathcal{C},$$

$$\frac{1}{2}(K_1 + K_2) = \begin{bmatrix} 0.5 & 1 \\ 1 & 0.5 \end{bmatrix} \notin \mathcal{C}$$

Non-convexity in control

Policy optimization in control is generally **nonconvex**!

$$\min_{\mathsf{K}} \quad J(\mathsf{K})$$
$$\text{s.t.} \quad \boxed{\mathsf{K} \in \mathcal{C}}$$

- The set of dynamic stabilizing policies is **nonconvex** and may even be **disconnected**. [Tang, **Zheng**, Li, *Math. Prog.* 2023]

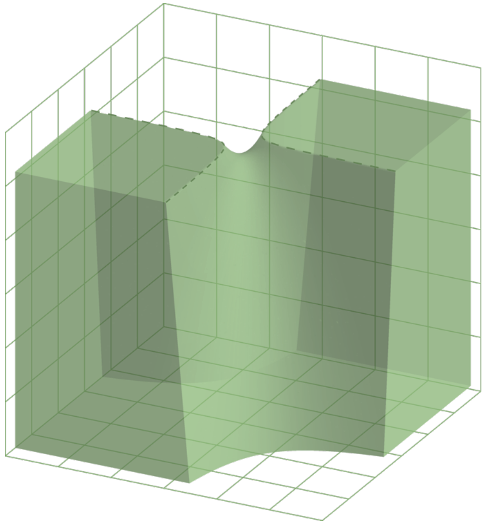

Example 1: Nonconvex but **connected**

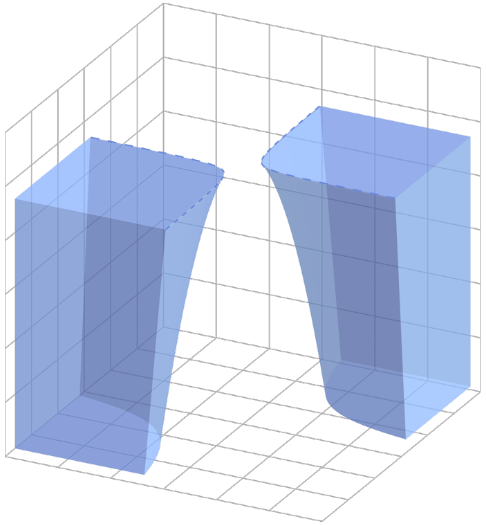

Example 2: Nonconvex and **dis-connected**

Benign non-convexity in control

Policy optimization in **linear** control is generally **nonconvex but benign**!

$$\min_{\mathsf{K}} \quad J(\mathsf{K})$$
$$\text{s.t.} \quad \mathsf{K} \in \mathcal{C}$$

- Including all iconic benchmark control problems:
  - ✓ LQR, LQG, H2/Hinf optimal control, etc.

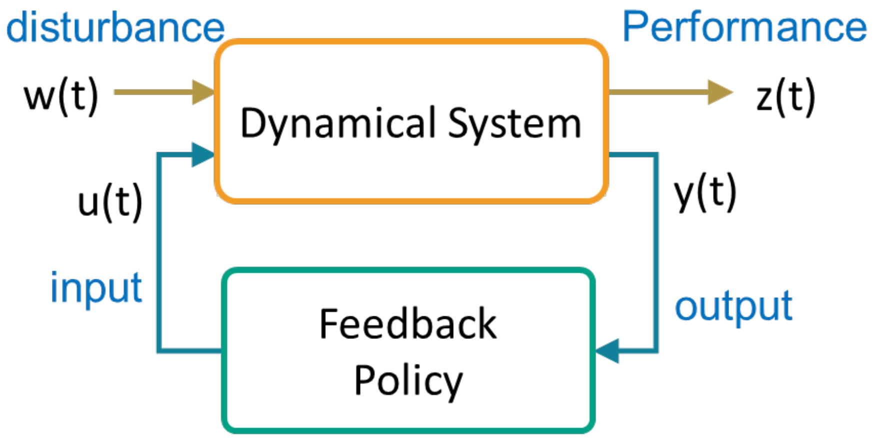

❑ Linear Quadratic Regulator (LQR)

Dynamics: $\dot{x}(t) = Ax(t) + Bu(t) + B_w w(t),$

Static policies: $u(t) = Kx(t)$

Stability: $\mathcal{C} = \{K \in \mathbb{R}^{m\times n} \mid A + BK \text{ is stable}\}$

Performance: $$J_{\mathrm{LQR}}(K) := \lim_{T\to\infty} \mathbb{E}\left[\frac{1}{T}\int_0^T x^{\mathsf{T}}(t)Qx(t) + u^{\mathsf{T}}(t)Ru(t)\,dt\right]$$

Benign non-convexity in LQR

- **Linear Quadratic Regulator (LQR)**

$$\min_{\mathsf{K}} \quad J(\mathsf{K})$$
$$\text{s.t.} \quad \mathsf{K} \in \mathcal{C}$$

Stability: $\mathcal{C} = \{K \in \mathbb{R}^{m\times n} \mid A + BK \text{ is stable}\}$

Fact 1: the set of stabilizing gains is always **path-connected**

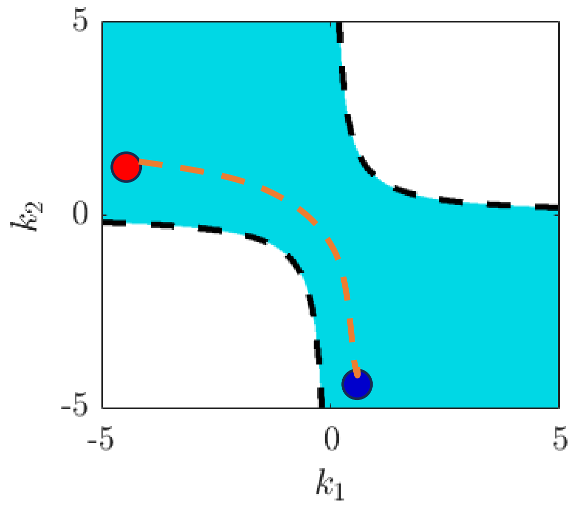

Proof: **convexification** via a **nonlinear change of variables**

$$
\begin{aligned}
&\{K \in \mathbb{R}^{m\times n} \mid A + BK \text{ is stable}\}\\
\iff &\{K \in \mathbb{R}^{m\times n} \mid \exists P \succ 0, P(A+BK)^{\mathsf{T}} + (A+BK)P \prec 0\}\\
\iff &\{K \in \mathbb{R}^{m\times n} \mid \exists P \succ 0, PA^{\mathsf{T}} + L^{\mathsf{T}}B^{\mathsf{T}} + AP + BL \prec 0, L = KP\}\\
\iff &\{K = LP^{-1} \in \mathbb{R}^{m\times n} \mid \exists P \succ 0, PA^{\mathsf{T}} + L^{\mathsf{T}}B^{\mathsf{T}} + AP + BL \prec 0\}.
\end{aligned}
$$

Any convex set is path-connected, and thus so is its continuous image.

Benign non-convexity in LQR

- **Linear Quadratic Regulator (LQR)**

$$J_{\mathrm{LQR}}(K) := \lim_{T\to\infty} \mathbb{E}\left[\frac{1}{T}\int_0^T x^{\mathsf{T}}(t)Qx(t) + u^{\mathsf{T}}(t)Ru(t)\,dt\right]$$

$$\min_{\mathsf{K}} \quad J(\mathsf{K}) \qquad \text{s.t.} \quad \mathsf{K} \in \mathcal{C}$$

Fact 2: the LQR cost is real analytical and thus **infinitely differentiable** over C.

✓ The LQR cost is actually a **rational function** in K

✓ Another form $\quad J_{\mathrm{LQR}}(K) = \mathrm{Tr}\left[(Q + K^{\mathsf{T}}RK)X\right]$

Lyapunov equation $\quad (A + BK)X + X(A + BK)^{\mathsf{T}} + W = 0$

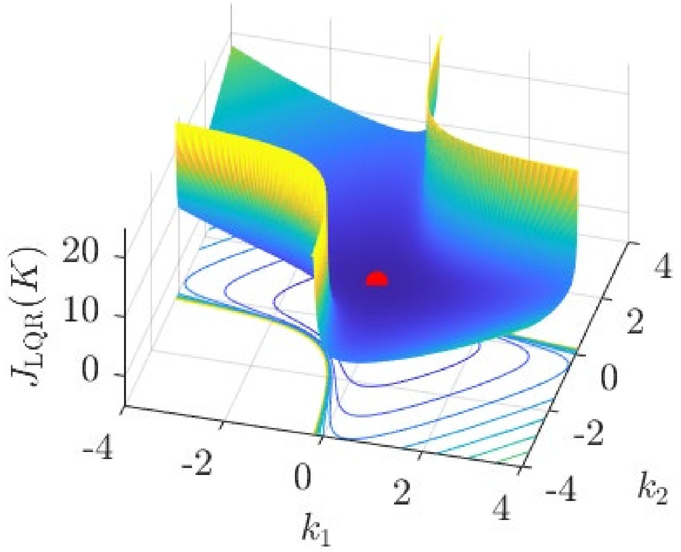

Fact 3: **Any stationary point $\nabla J(K) = 0$ is globally optimal**

Local Stationarity → **Global Optimality**

Benign non-convexity in LQR

- Linear Quadratic Regulator (LQR)

$$\min_{\mathsf{K}} \quad J(\mathsf{K})$$
$$\text{s.t.} \quad \mathsf{K} \in \mathcal{C}$$

Fact 4: the LQR cost satisfies the **gradient dominance** property over any sublevel set (Fazel et al., 2019; Mohammadi et al., 2022).

$$J_{\text{LQR}}(K) - J^\star \le \mu \|\nabla J_{\text{LQR}}(K)\|_F^2, \quad \forall K \in [J_{\text{LQR}} \le \nu]$$

- LQR almost behaves like a **strongly convex** problem under mild assumptions (Watanabe and **Zheng, 2025**).
- The basic gradient descent algorithm achieves linear convergence (Fazel et al., 2019; Mohammad et al., 2022; Fatkhullin & Polyak, 2021)

$$K_{t+1} = K_t - \alpha \nabla J(K_t)$$

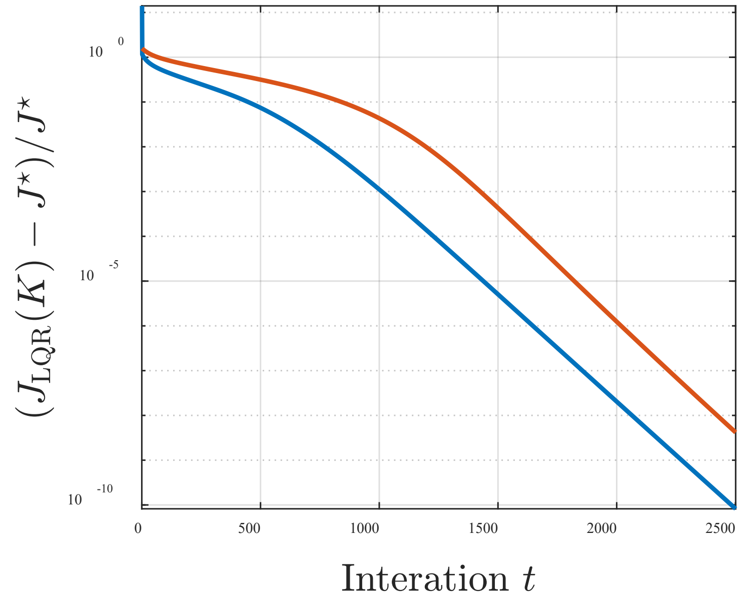

Benign non-convexity in robust control

❑ State-feedback Hinf robust control

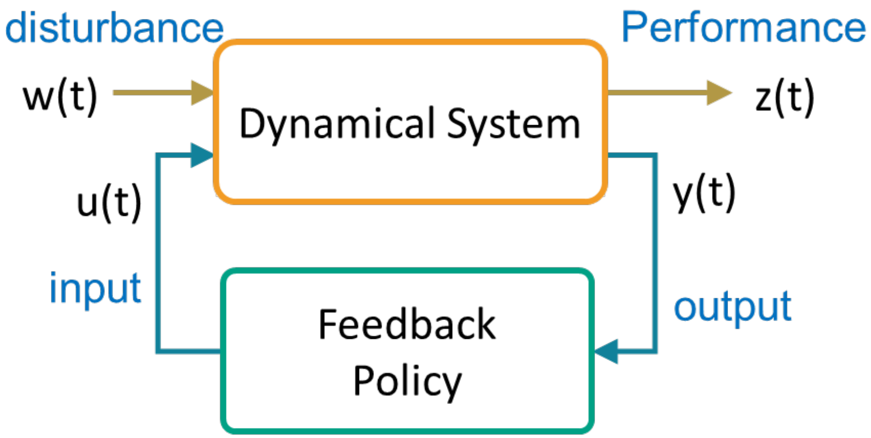

Dynamics: $\dot{x}(t) = Ax(t) + Bu(t) + B_w w(t),$

Static policies: $u(t) = Kx(t)$

Stability: $\mathcal{C} = \{K \in \mathbb{R}^{m\times n} \mid A + BK \text{ is stable}\}$

Unlike LQR with stochastic noises, we here assume **adversarial noises** with bounded energy and consider the **worst-case performance**

Performance: $$J_\infty(K) := \sup_{\|w(t)\|_2 \leq 1} \int_0^\infty x^{\mathsf{T}}(t)Qx(t) + u^{\mathsf{T}}(t)Ru(t)\, dt$$

Benign non-convexity in robust control

❑ State-feedback Hinf robust control

$$\min_{\mathsf{K}} \quad J(\mathsf{K})$$
$$\text{s.t.} \quad \mathsf{K} \in \mathcal{C}$$

Performance: $$J_\infty(K) := \sup_{\|w(t)\|_2 \le 1} \int_0^\infty x^{\mathsf{T}}(t) Q x(t) + u^{\mathsf{T}}(t) R u(t)\, dt$$

Fact 1: The Hinf cost is generally **non-convex**, and **non-smooth**, but **locally weakly convex**.

Fact 2: **Any sublevel set is path-connected**, and **any stationary point is globally optimal** (Guo and Hu, 2022; **Zheng** et al., 2023)

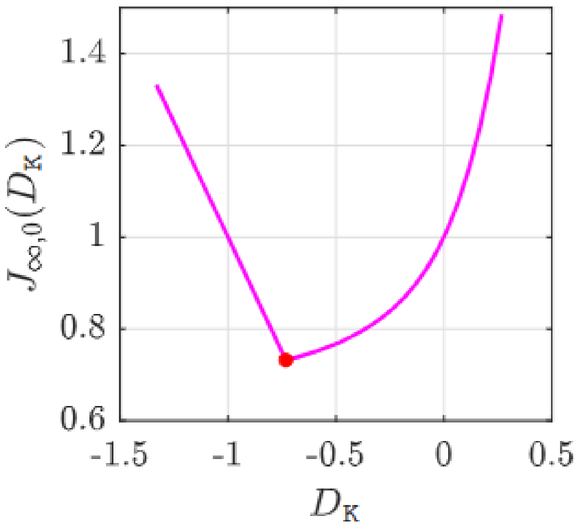

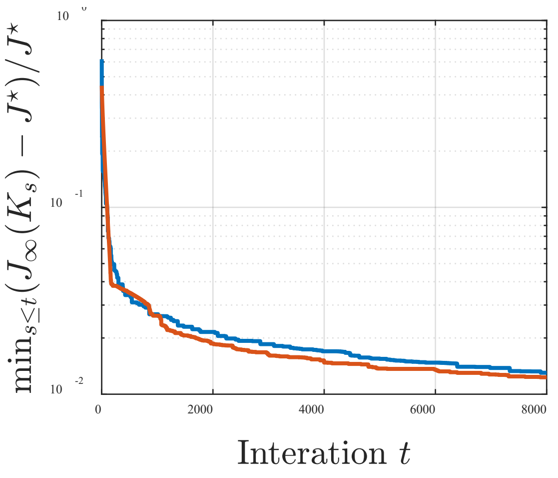

- We may use the subgradient method to find a global minimizer despite the non-convexity (Watanabe, Liao, & **Zheng**, 2025)

$$K_{t+1} = K_t - \alpha_t g_t,$$
$$g_t \in \partial J_\infty(K_t)$$

Benign non-convexity in LQG

❑ Linear Quadratic Gaussian (LQG) Control

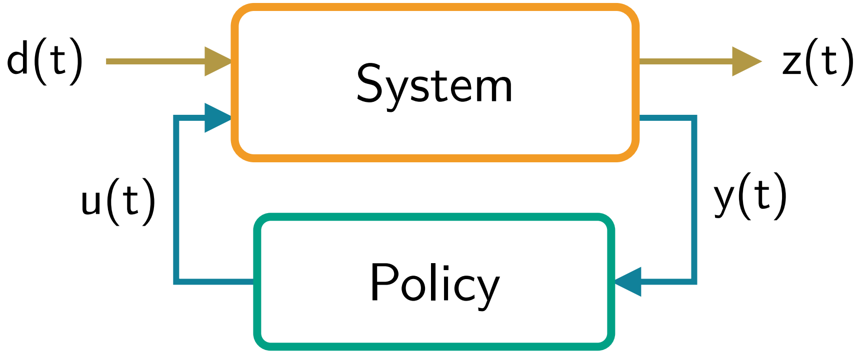

Dynamics: $\dot{x}(t) = Ax(t) + Bu(t) + B_w w(t)$

$$\boxed{y(t) = Cx(t) + D_v v(t)}$$

Performance:

$$J_{\mathrm{LQG}} := \lim_{T\to\infty} \mathbb{E}\left[\frac{1}{T}\int_0^T x^{\mathsf{T}}(t)Qx(t) + u^{\mathsf{T}}(t)Ru(t)\,dt\right]$$

➢ We consider dynamic policies:

$$\dot{\xi}(t) = A_{\mathsf{K}}\xi(t) + B_{\mathsf{K}}y(t)$$
$$u(t) = C_{\mathsf{K}}\xi(t)$$

$$\mathsf{K} = (A_{\mathsf{K}}, B_{\mathsf{K}}, C_{\mathsf{K}}) \in \mathbb{R}^{n\times n} \times \mathbb{R}^{n\times p} \times \mathbb{R}^{m\times n}$$

Policy optimization for LQG

$$\min_{\mathsf{K}} \quad J(\mathsf{K})$$
$$\text{s.t.} \quad \mathsf{K} = (A_{\mathsf{K}}, B_{\mathsf{K}}, C_{\mathsf{K}}) \in \mathcal{C}_{\text{full}}$$

Benign nonconvexity in LQG

❑ Linear Quadratic Gaussian (LQG) Control

$$\mathcal{C}_{\text{full}} = \left\{ \mathsf{K} \;\middle|\; \begin{bmatrix} A & BC_{\mathsf{K}} \\ B_{\mathsf{K}}C & A_{\mathsf{K}} \end{bmatrix} \text{ is Hurwitz stable} \right\}$$

Fact 1: The set $\mathcal{C}_{\text{full}}$ can be **disconnected** but has at most **2 connected** components [Tang, **Zheng**, Li, 2023].

Proof: **convexification** via a **nonlinear change of variables**

$$\boxed{\mathcal{F} \times \mathrm{GL}} \xrightarrow{\Phi} \boxed{\mathcal{C}_{\text{full}}}$$

Continuous surjective map

at most 2 connected components

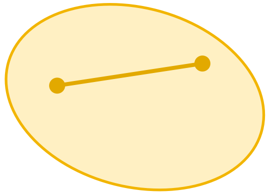

Convex thus connected

General linear group: the set of invertible matrices (similarity transformation)

Policy optimization for LQG

$$\begin{aligned} \min_{\mathsf{K}} \quad & J(\mathsf{K}) \\ \text{s.t.} \quad & \mathsf{K} = (A_{\mathsf{K}}, B_{\mathsf{K}}, C_{\mathsf{K}}) \in \mathcal{C}_{\text{full}} \end{aligned}$$

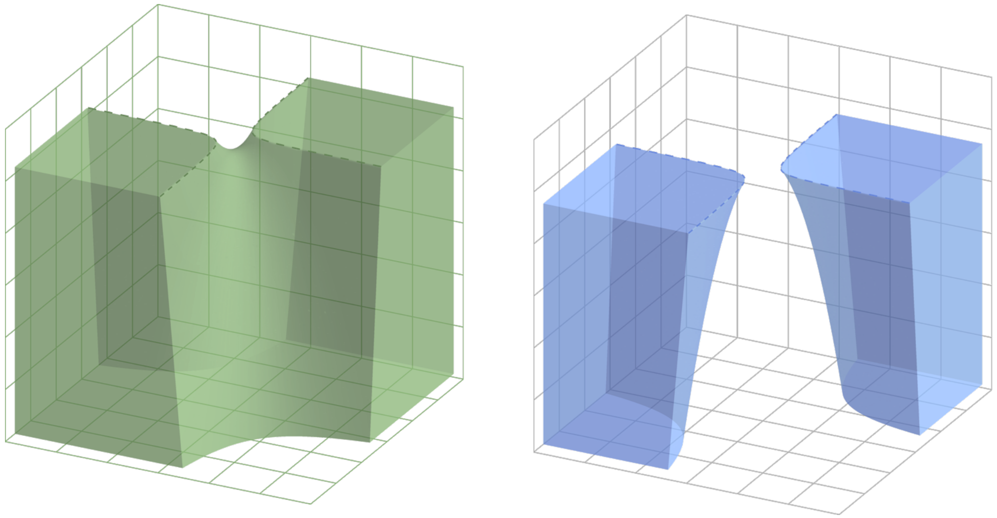

Two connected components

$$\mathrm{GL}_n^+ = \{\Pi \in \mathbb{R}^{n\times n} \mid \det \Pi > 0\},$$
$$\mathrm{GL}_n^- = \{\Pi \in \mathbb{R}^{n\times n} \mid \det \Pi < 0\}.$$

Benign nonconvexity in LQG

❑ Linear Quadratic Gaussian (LQG) Control

Fact 2: For a controllable and observable policy, if it is a **stationary** point $\nabla J(\mathsf{K}) = 0$, then it is a **global minimizer** [Tang, **Zheng**, Li, 2023].

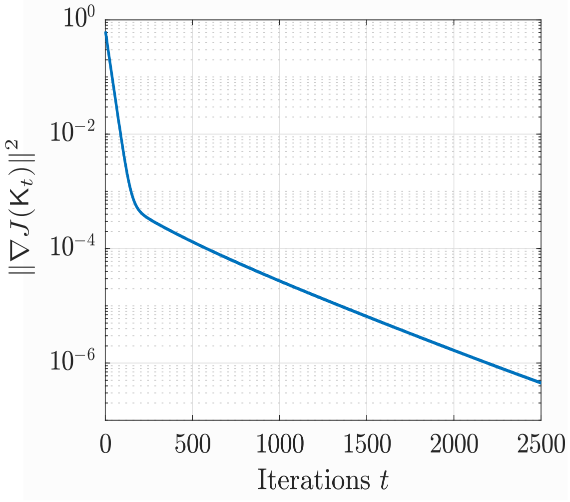

Local Zero Gradient ✚ Structural Information ➡ Global Optimality Certificate

Fact 3: For policies that are uncontrollable or unobservable, there exists suboptimal stationary points (e.g., saddle points) [Tang, **Zheng**, Li, 2023].

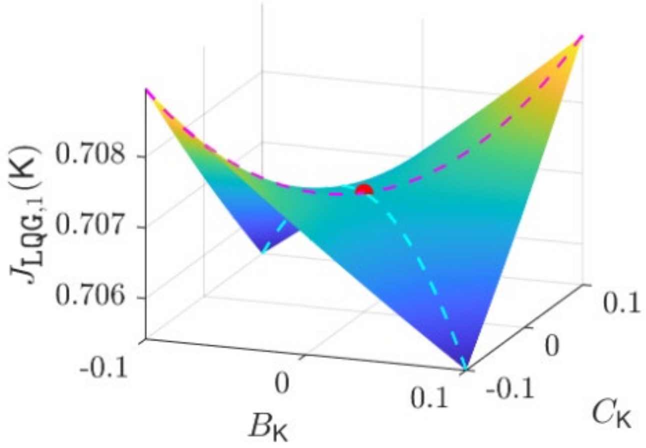

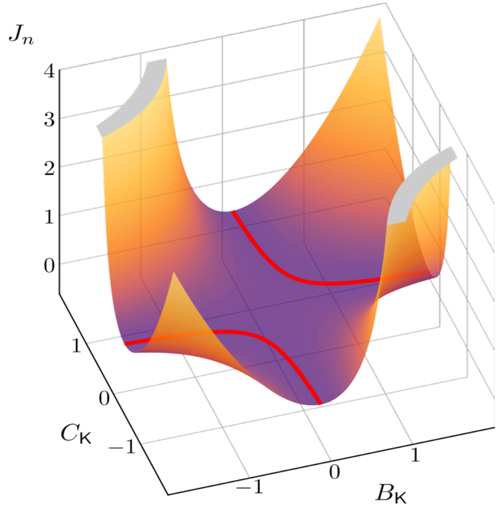

Benign Nonconvex Policy Optimization

❑ Recent progress on non-convex policy optimization

- Favorable properties have been revealed for policy optimization in many benchmark control problems:
  - ✓ LQR [Fazel et al., 2018] [Mohammad et al., 2022] [Moghaddam, Olshevsky & Gharesifard, 2025], etc.
  - ✓ LQG [Tang, **Zheng,** & Li, 2021], [**Zheng** et al., 2022], [Ren et al., 2023] [Kraisler & Mesbahi, 2024], etc.
  - ✓ $\mathcal{H}_\infty$ state-feedback/output-feedback, [Guo & Hu, 2022] [Hu & **Zheng**, 2022] [**Zheng** et al, 2023 & 2024]

ANNUAL REVIEW OF CONTROL, ROBOTICS, AND AUTONOMOUS SYSTEMS Volume 6, 2023

Review Article | **Open Access**

Toward a Theoretical Foundat
Learning Control Policies

Bin Hu[1], Kaiqing Zhang[2,3], Na Li[4], Mehran Mesbahi[5], Maryam F

Policy Optimization in Control: Geometry and Algorithmic Implications

Shahriar Talebi[a], Yang Zheng[b], Spencer Kraisler[c], Na Li[a], Mehran Mesbahi[c]

[a] *Harvard University, School of Engineering and Applied Sciences, 150 Western Ave, Boston, 02134, MA, US*
[b] *University of California San Diego, Department of Electrical and Computer Engineering, 9500 Gilman Drive, La Jolla, 92093, CA, US*
[c] *University of Washington, Department of Aeronautics and Astronautics, 3940 Benton Ln NE, Seattle, 98195, WA, US*

Outline

Geometry:
Benign Nonconvex Landscape

Framework:
Extended Convex Lifting

Algorithm:
Proximal Descent Method

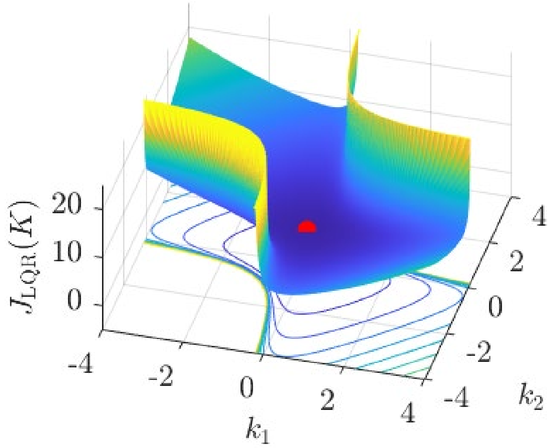

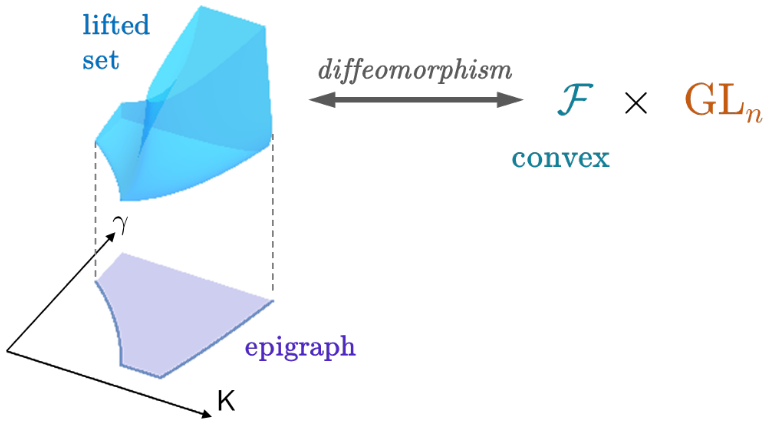

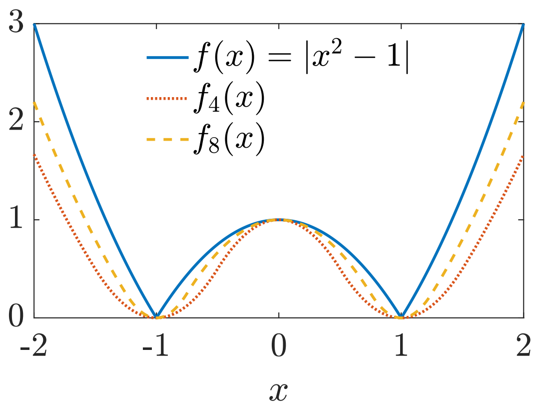

Yang Zheng, Chih-Fan Pai, and Yujie Tang. "**Benign Nonconvex Landscapes in Optimal and Robust Control, Part II: Extended Convex Lifting.**" IEEE Transactions on Automatic Control (2026): https://arxiv.org/abs/2406.04001

Key idea

Benign Non-convexity in Control via Extended Convex Lifting (ECL)

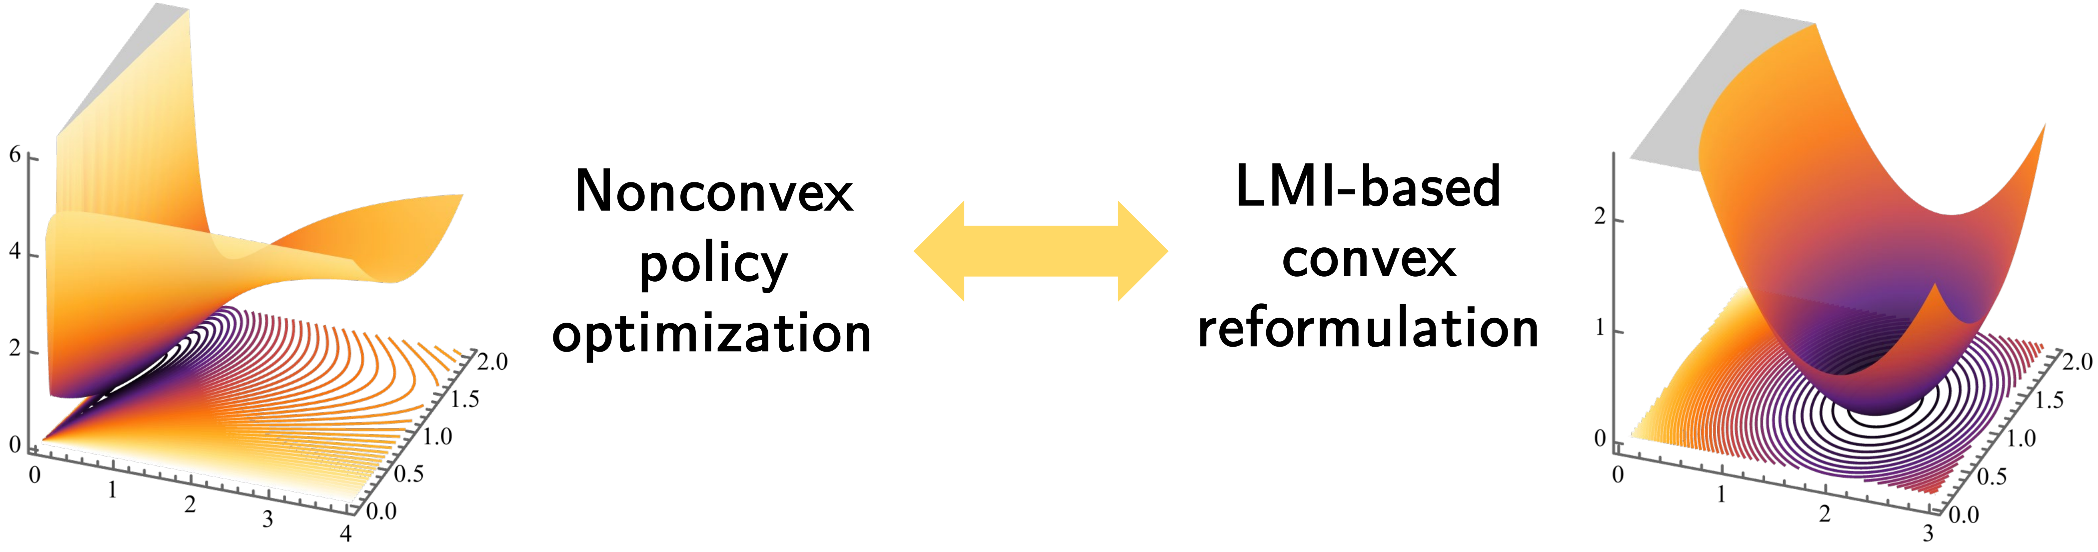

- Our idea: Exploit **LMI-based convex reformulations** of control problems
- They reveal the hidden convexity of policy optimization landscapes.
- For non-degenerate policies, all Clarke stationary points are globally optimal and there is no spurious local minimum in policy optimization.

Example 1

❑ Nonconvex and smooth function

$$f_1(x_1, x_2) = \left(\underbrace{\boxed{\frac{x_2}{x_1}}}_{y_1} - 2\right)^2 + (\underbrace{\boxed{x_2}}_{y_2} - 1)^2, \qquad \mathrm{dom}(f_1) = \{x \in \mathbb{R}^2 \mid x_1 > 0, x_2 > 0\}.$$

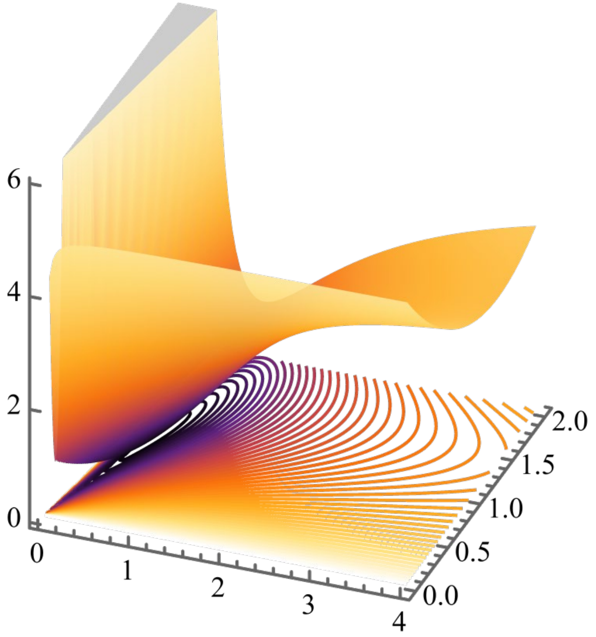

Its global **minimizer** is

$$x^* = (0.5, 1)$$

Define an **invertible map**

$$g(x) := (x_2/x_1, x_2),$$
$$\forall x_1 > 0, x_2 > 0,$$

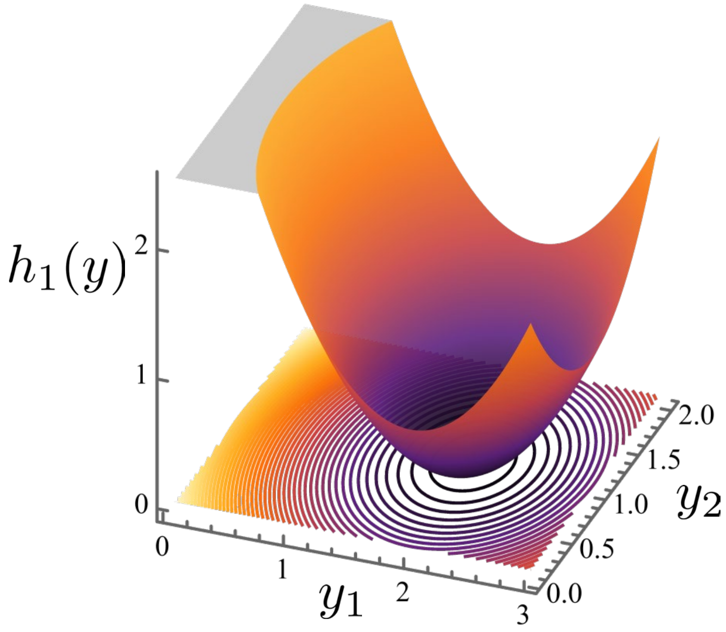

$$h_1(y) := f_1(g^{-1}(y)) = (y_1 - 2)^2 + (y_2 - 1)^2, \qquad \forall y_1 > 0, y_2 > 0.$$

Example 2

❑ Nonconvex and non-smooth function

$$f_2(x_1, x_2) = \left|\underbrace{\boxed{\frac{x_2}{x_1}}}_{y_1} - 2\right| + |\underbrace{\boxed{x_2}}_{y_2} - 1|, \qquad \mathrm{dom}(f_2) = \{x \in \mathbb{R}^2 \mid x_1 > 0, x_2 > 0\}.$$

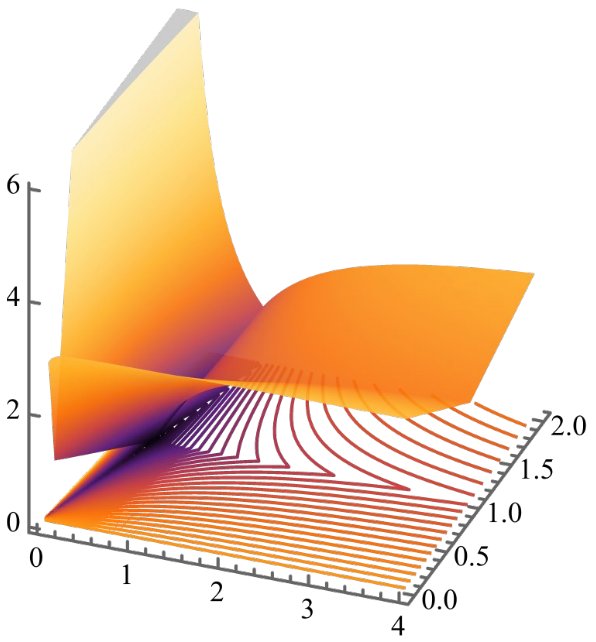

Its global **minimizer** is

$$x^* = (0.5, 1)$$

Define an **invertible map**

$$g(x) := (x_2/x_1, x_2),$$
$$\forall x_1 > 0, x_2 > 0,$$

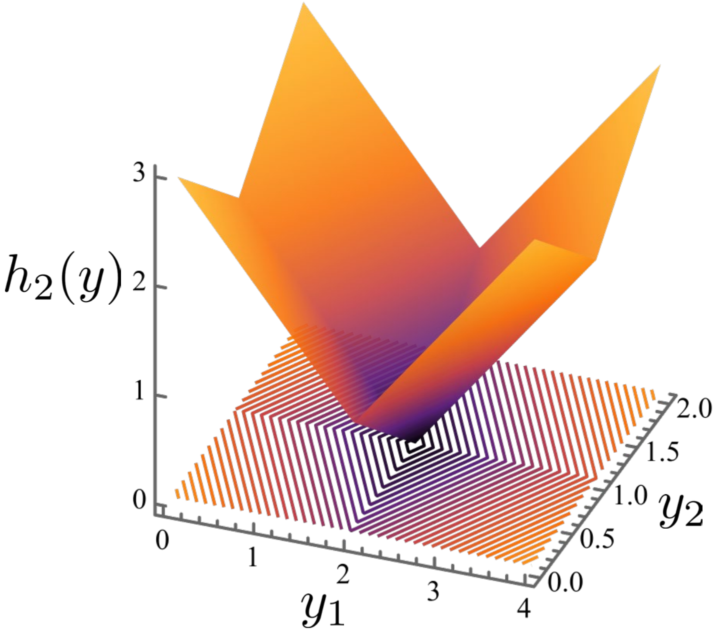

$$h_2(y) := f_2(g^{-1}(y)) = |y_1 - 2| + |y_2 - 1|, \qquad \forall y_1 > 0, y_2 > 0,$$

Direct Convex Reformulation

- Direct convex reformulation (no lifting)
  - Consider a continuous function $J(\mathsf{K}) : \mathcal{D} \to \mathbb{R}$.

    Denote its epigraph as $\operatorname{epi}_{\geq}(f) := \{(\mathsf{K}, \gamma) \in \mathcal{D} \times \mathbb{R} \mid \gamma \geq J(\mathsf{K})\}$.
  - Suppose there exists a **smooth and invertible map** Φ between $\operatorname{epi}_{\geq}(J)$ and a **convex set** $\mathcal{F}_{\mathrm{cvx}}$
  - and we further have $(y, \gamma) = \Phi(\mathsf{K}, \gamma), \ \forall (\mathsf{K}, \gamma) \in \operatorname{epi}_{\geq}(J)$

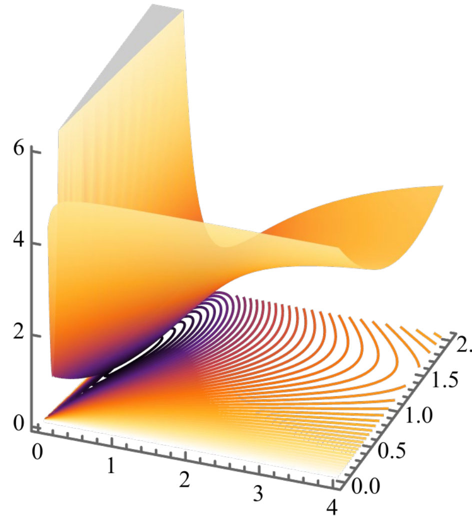

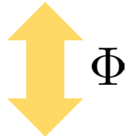

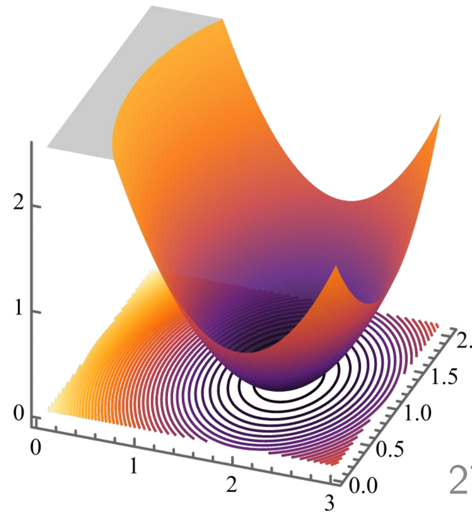

Guarantee 1: Minimizing J is equivalent to a convex problem

$$\inf_{\mathsf{K} \in \mathcal{D}} J(\mathsf{K}) = \inf_{(y,\gamma) \in \mathcal{F}_{\mathrm{cvx}}} \gamma.$$

Guarantee 2: Any stationary point to J is globally optimal; in other words, $0 \in \partial J(\mathsf{K}^*)$ implies globally optimality

Extended Convex Lifting (ECL)

A schematic illustration of ECL:

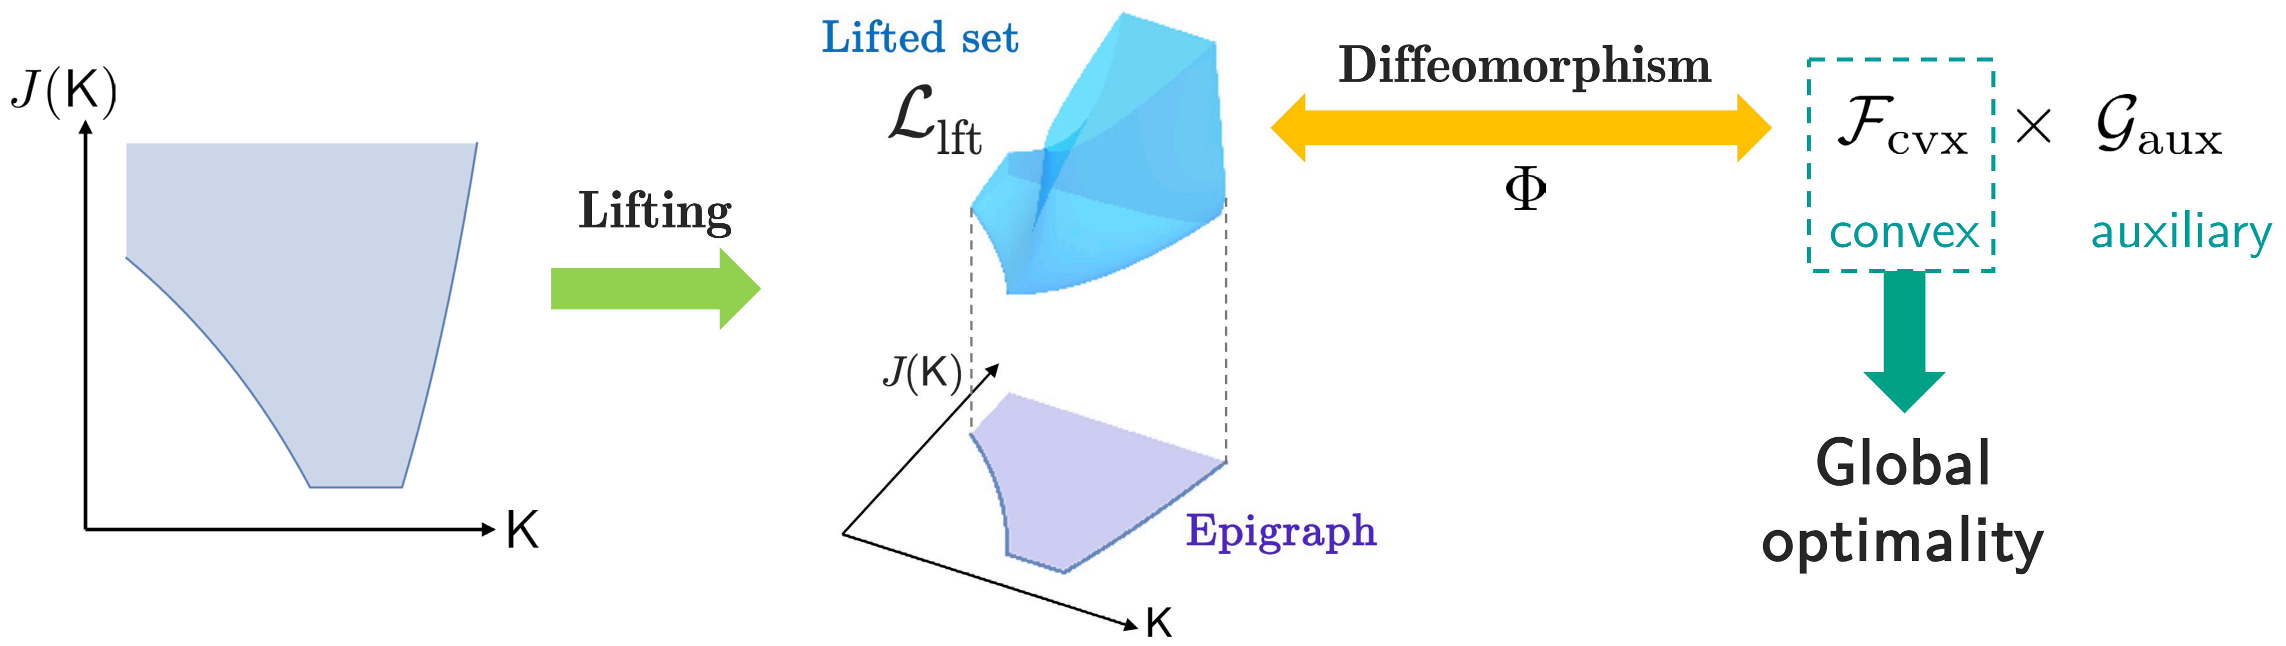

Two key features

Feature 1: a lifting procedure

Feature 2: an auxiliary set

Extended Convex Lifting (ECL)

A schematic illustration of ECL:

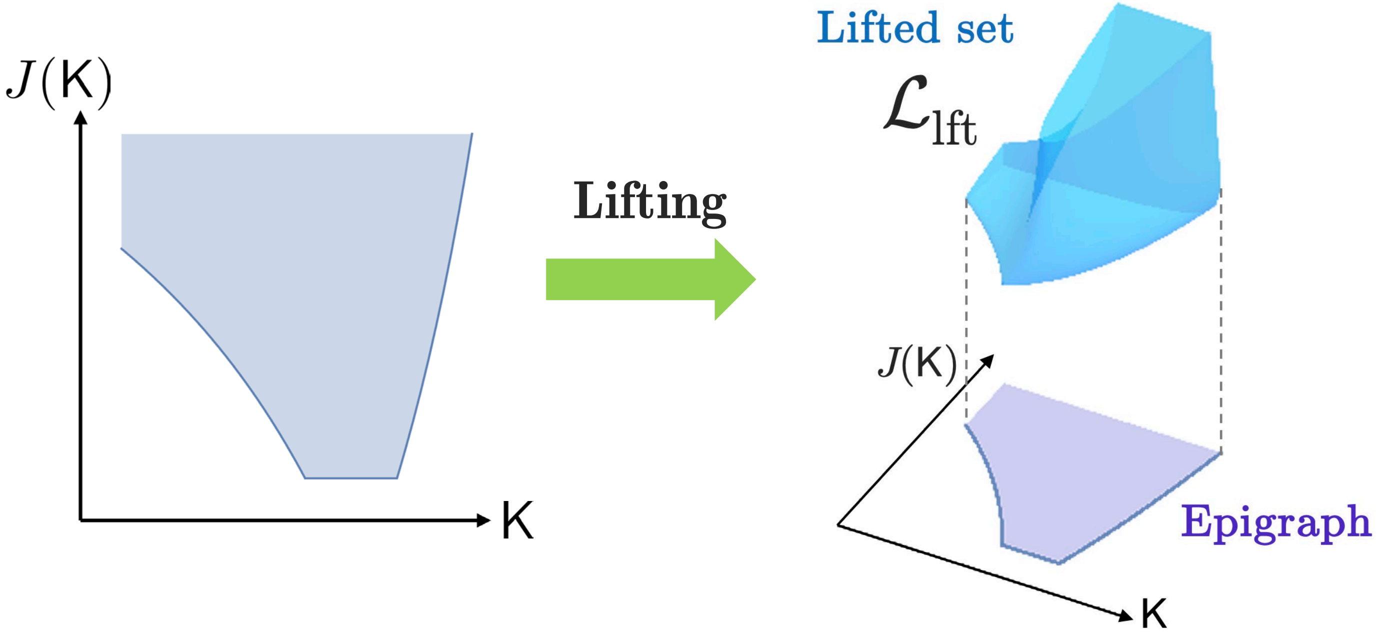

$$\boxed{\begin{matrix}\mathrm{epi}_{>}(J) \\ (\mathsf{K}, \gamma)\end{matrix}} \longrightarrow \boxed{\begin{matrix}\mathcal{L}_{\text{lft}} \\ (\mathsf{K}, \gamma, \xi)\end{matrix}}$$

Why lifting?

- For many control problems, a **direct convexification is not possible**
- **A lifting procedure** corresponding to **Lyapunov variables** is necessary.

Extended Convex Lifting (ECL)

A schematic illustration of ECL:

Why auxiliary set?

- Allows us to isolate the **redundancy** or **symmetry** in the original nonconvex domain
- Related to **similarity transformations** of dynamic policies in control
- Needed for **output-feedback control** problems

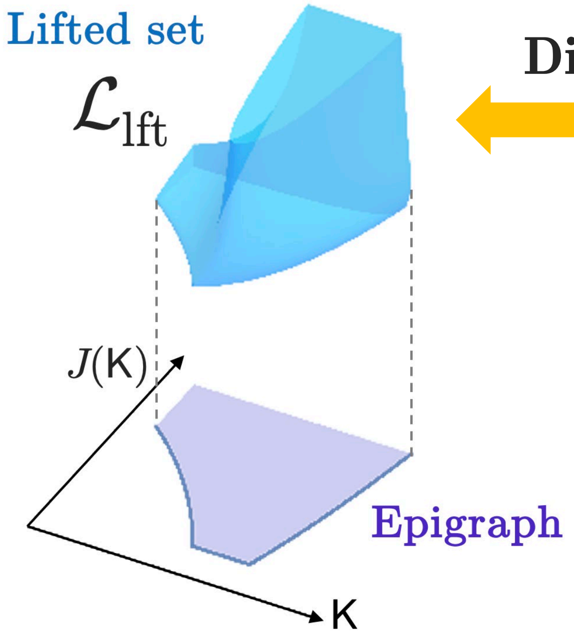

$$\mathcal{L}_{\mathrm{lft}} \xleftrightarrow[\Phi]{\textbf{Diffeomorphism}} \mathcal{F}_{\mathrm{cvx}} \times \mathcal{G}_{\mathrm{aux}}$$

convex auxiliary

$$\boxed{\begin{matrix}\mathcal{L}_{\mathrm{lft}} \\ (\mathsf{K}, \gamma, \xi)\end{matrix}} \longleftrightarrow \boxed{\begin{matrix}\mathcal{F}_{\mathrm{cvx}} \times \mathcal{G}_{\mathrm{aux}} \\ (\gamma, \zeta_1, \zeta_2)\end{matrix}}$$

Formal ECL Definition

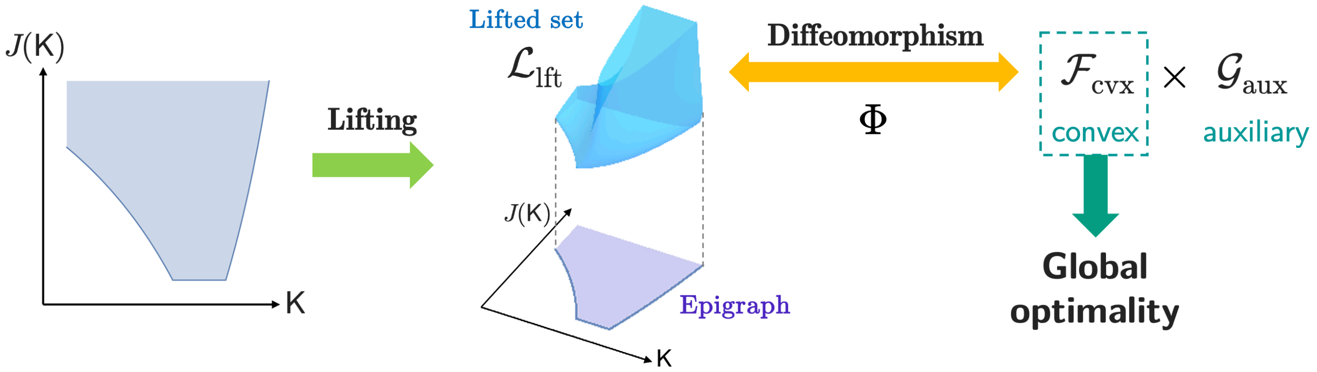

$$\begin{array}{c}\mathrm{epi}_>(J)\\ (\mathsf{K},\gamma)\end{array} \longrightarrow \begin{array}{c}\mathcal{L}_{\mathrm{lft}}\\ (\mathsf{K},\gamma,\xi)\end{array} \longleftrightarrow \begin{array}{c}\mathcal{F}_{\mathrm{cvx}} \times \mathcal{G}_{\mathrm{aux}}\\ (\gamma,\zeta_1,\zeta_2)\end{array}$$

- Consider a continuous function $J(\mathsf{K}) : \mathcal{D} \to \mathbb{R}$ where $\mathcal{D} \subseteq \mathbb{R}^d$.
- Denote its strict and non-strict epigraph as

$$\mathrm{epi}_>(J) := \{(\mathsf{K},\gamma) \in \mathcal{D} \times \mathbb{R} \mid \gamma > J(\mathsf{K})\},$$
$$\mathrm{epi}_\geq(J) := \{(\mathsf{K},\gamma) \in \mathcal{D} \times \mathbb{R} \mid \gamma \geq J(\mathsf{K})\}.$$

Extended Convex Lifting (ECL)

We say a tuple $(\mathcal{L}_{\mathrm{lft}}, \mathcal{F}_{\mathrm{cvx}}, \mathcal{G}_{\mathrm{aux}}, \Phi)$ is an ECL of $J(\mathsf{K}) : \mathcal{D} \to \mathbb{R}$ if

- A lifted set $\mathcal{L}_{\mathrm{lft}}$ satisfying

$$\boxed{\mathrm{epi}_>(J) \subseteq \pi_{\mathsf{K},\gamma}(\mathcal{L}_{\mathrm{lft}}) \subseteq \mathrm{cl}\,\mathrm{epi}_\geq(J)}$$

- A diffeomorphism $\Phi : \mathcal{L}_{\mathrm{lft}} \to \mathcal{F}_{\mathrm{cvx}} \times \mathcal{G}_{\mathrm{aux}}$ such that

$$\Phi(\mathsf{K},\gamma,\xi) = (\gamma,\zeta_1,\zeta_2)$$

A special ECL

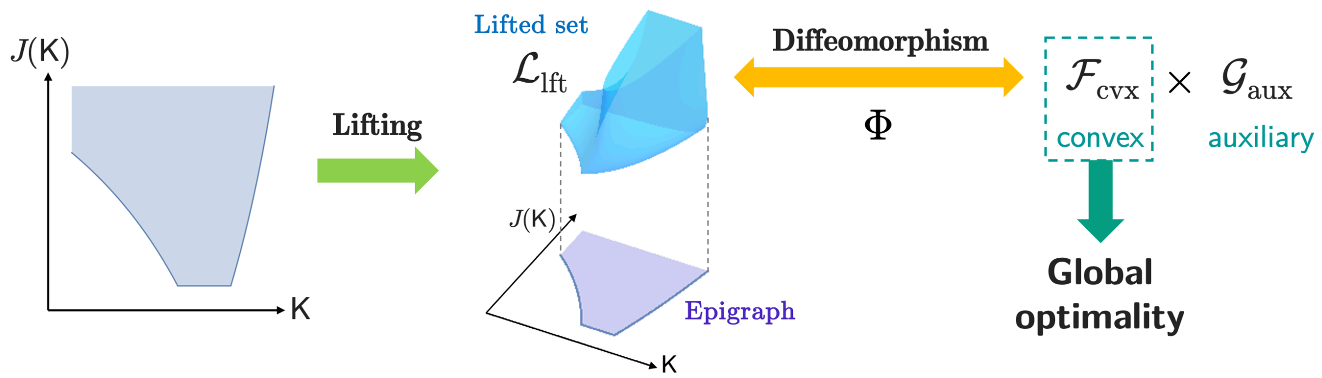

A more intuitive condition

- A lifted set $\mathcal{L}_{\mathrm{lft}}$ satisfying
  $\mathrm{epi}_{>}(J) \subseteq \pi_{\mathsf{K},\gamma}(\mathcal{L}_{\mathrm{lft}}) \subseteq \mathrm{cl}\,\mathrm{epi}_{\geq}(J)$

- A lifted set $\mathcal{L}_{\mathrm{lft}}$ satisfying
  $\pi_{\mathsf{K},\gamma}(\mathcal{L}_{\mathrm{lft}}) = \mathrm{epi}_{\geq}(f).$

Does this "simpler" lifting condition work?

- Apparently, the condition on the left is **less restrictive**, and works for more general situations
- The simpler condition on the right is indeed sufficient for **state-feedback control** problems
- However, it is too restrictive for **dynamic output-feedback control** problems (e.g., LQG)

Non-degenerate points

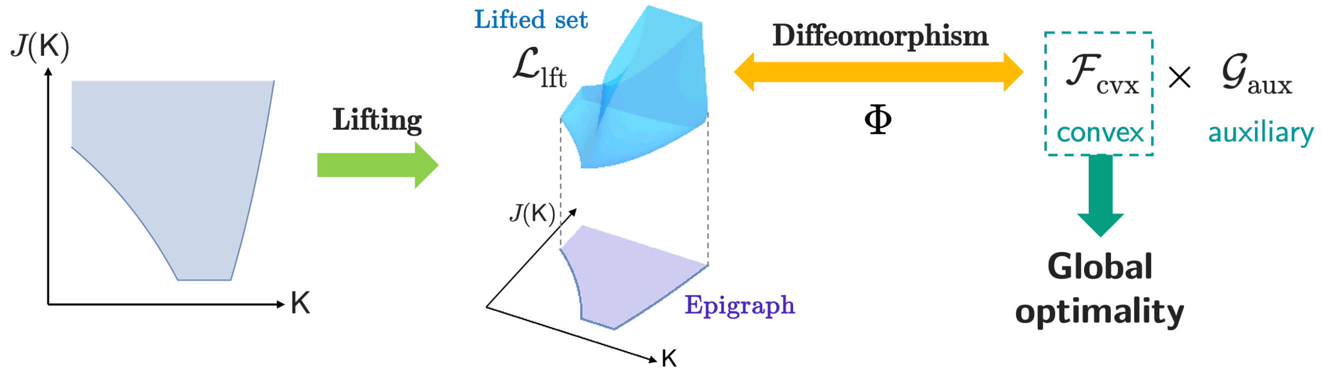

$\mathrm{epi}_>(J)$ (K, γ) → $\mathcal{L}_{\mathrm{lft}}$ $(\mathsf{K}, \gamma, \xi)$ ↔ $\mathcal{F}_{\mathrm{cvx}} \times \mathcal{G}_{\mathrm{aux}}$ $(\gamma, \zeta_1, \zeta_2)$

Extended Convex Lifting:

- A lifted set $\mathcal{L}_{\mathrm{lft}}$ satisfying
  $$\boxed{\mathrm{epi}_>(J) \subseteq \pi_{\mathsf{K},\gamma}(\mathcal{L}_{\mathrm{lft}}) \subseteq \mathrm{cl}\,\mathrm{epi}_\geq(J)}$$
- A diffeomorphism $\Phi : \mathcal{L}_{\mathrm{lft}} \to \mathcal{F}_{\mathrm{cvx}} \times \mathcal{G}_{\mathrm{aux}}$ such that
  $$\Phi(\mathsf{K}, \gamma, \xi) = (\gamma, \zeta_1, \zeta_2)$$

- By construction, some points in $\mathrm{epi}_\geq(J)$ may **not be covered** in the lifted set

→ Those points will be called **degenerate** → bad behavior may happen (e.g., **saddles**)

Definition. K is called **non-degenerate** if $(\mathsf{K}, J(\mathsf{K})) \in \pi_{\mathsf{K},\gamma}(\mathcal{L}_{\mathrm{lft}})$ → **well-behaved**

ECL Guarantees

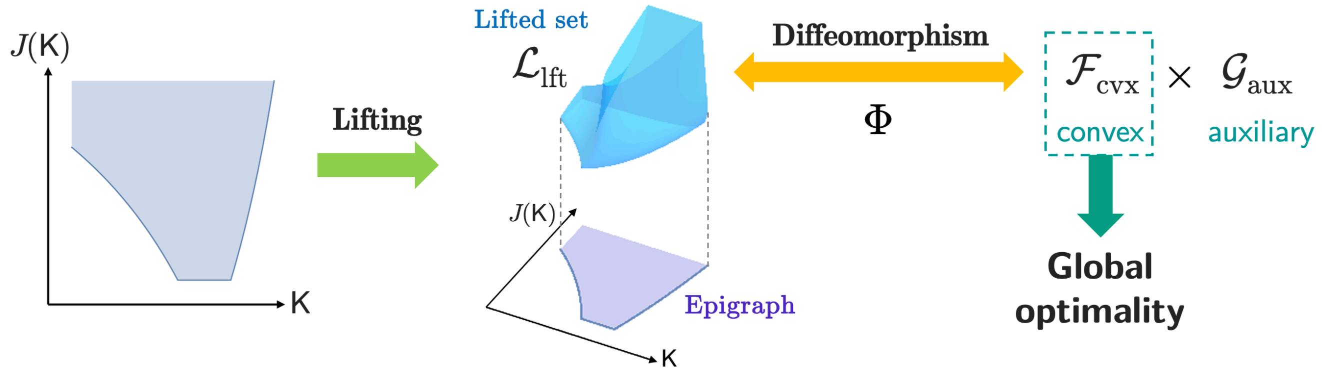

Given an Extended Convex Lifting

$$(\mathcal{L}_{\text{lft}}, \mathcal{F}_{\text{cvx}}, \mathcal{G}_{\text{aux}}, \Phi)$$

Guarantee 1: Convex Reformulation

Optimization $\min_{\mathsf{K}\in\mathcal{D}} J(\mathsf{K})$ is equivalent to a convex problem

$$\inf_{\mathsf{K}\in\mathcal{D}} J(\mathsf{K}) = \inf_{(y,\gamma)\in\mathcal{F}_{\text{cvx}}} \gamma.$$

Guarantee 2: Global Optimality

All non-degenerate Clarke stationary points are globally optimal (for subdifferentially regular functions)

- **Clarke stationary points**: Generalization of stationary points to nonsmooth functions, based on the notion of Clarke subdifferential

Framework:
Extended Convex Lifting

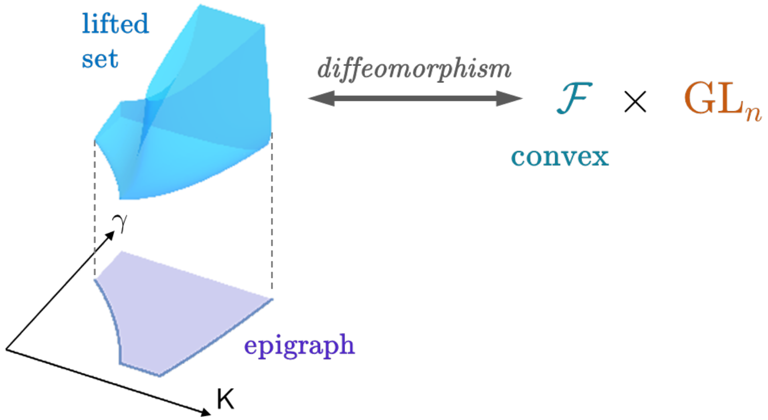

1. The ECL framework

2. **Applications in control**

Linear Quadratic Regulator (LQR)

❑ Problem setup

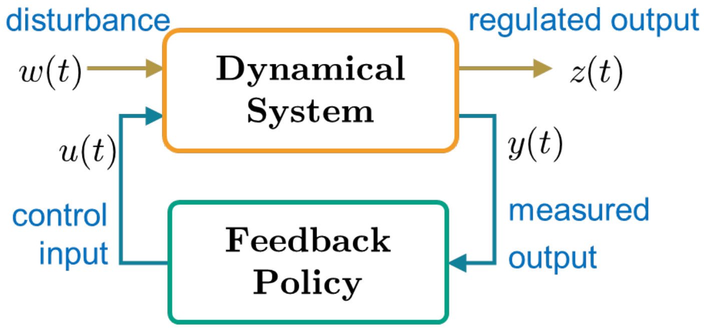

Dynamics: $$\dot{x}(t) = Ax(t) + Bu(t) + B_w w(t),$$

Static policies: $$u(t) = Kx(t)$$

Stability: $$\mathcal{C} = \{K \in \mathbb{R}^{m\times n} \mid A + BK \text{ is stable}\}$$

Performance: $$J_{\mathrm{LQR}}(K) := \lim_{T\to\infty} \mathbb{E}\left[\frac{1}{T}\int_0^T x^{\mathsf{T}}(t)Qx(t) + u^{\mathsf{T}}(t)Ru(t)\,dt\right]$$

❑ Nonconvex and smooth landscape

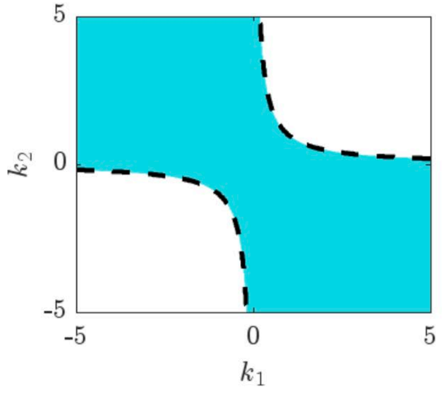

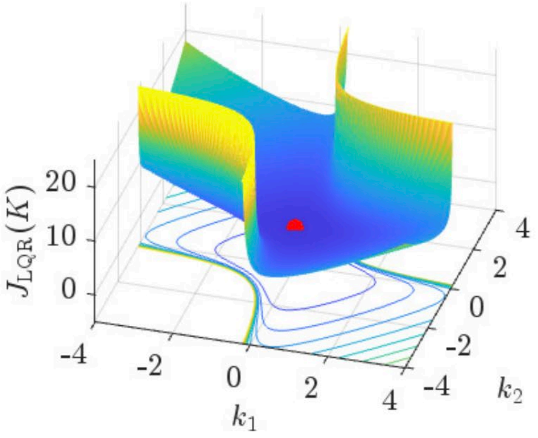

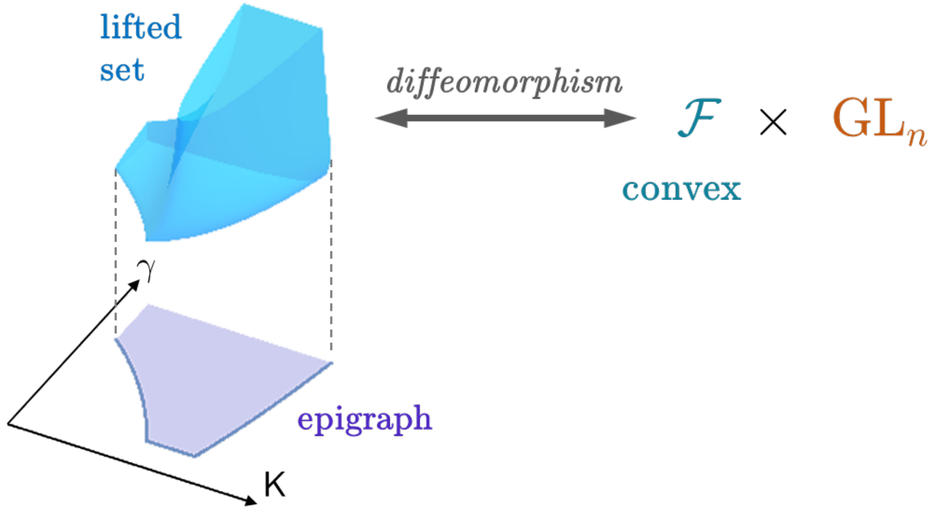

Linear Quadratic Regulator (LQR)

❑ Construction of ECL

Step 1: Lifting

$$\mathcal{L}_{\mathrm{LQR}} := \left\{(K,\gamma,X) : X \succ 0, (A+BK)X + X(A+BK)^{\mathsf{T}} + W = 0, \gamma \geq \mathrm{Tr}\left[(Q+K^{\mathsf{T}}RK)X\right]\right\}.$$

Step 2: Convex set

$$\mathcal{F}_{\mathtt{LQR}} = \left\{(\gamma,Y,X) : X \succ 0, AX + BY + XA^{\mathsf{T}} + Y^{\mathsf{T}}B^{\mathsf{T}} + W = 0, \gamma \geq \mathrm{tr}\big(QX + X^{-1}Y^{\mathsf{T}}RY\big)\right\}$$

Step 3: Diffeomorphism $\Phi(K,\gamma,X) = (\gamma, KX, X), \quad \forall (K,\gamma,X) \in \mathcal{L}_{\mathtt{LQR}}$

- No auxiliary set
- Lifted set satisfies $\mathrm{epi}_{\geq}(J) = \pi_{K,\gamma}(\mathcal{L}_{\mathtt{LQR}})$

➡ **All policies are non-degenerate**

Theorem. Any stationary point of the LQR cost function is globally optimal.

State-feedback Robust Control

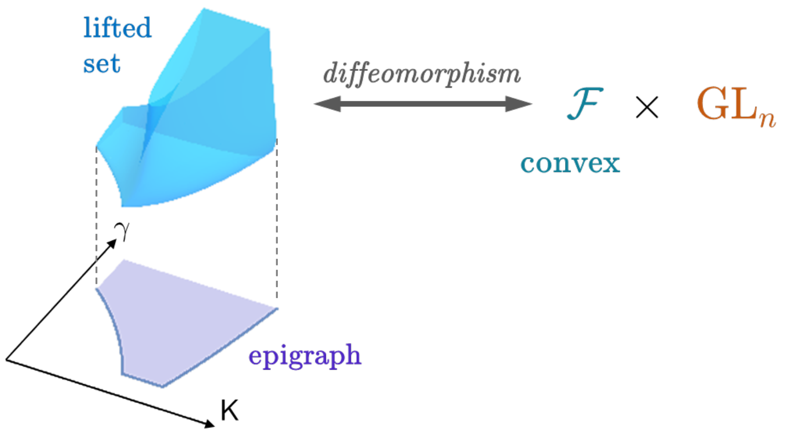

❑ Problem setup

Dynamics: $$\dot{x}(t) = Ax(t) + Bu(t) + B_w w(t),$$

Static policies: $$u(t) = Kx(t)$$

Stability: $$\mathcal{C} = \{K \in \mathbb{R}^{m\times n} \mid A + BK \text{ is stable}\}$$

Performance: $$J_\infty(K) := \sup_{\|w(t)\|_2 \le 1} \int_0^\infty x^\mathsf{T}(t)Qx(t) + u^\mathsf{T}(t)Ru(t)\,dt$$

❑ Building an ECL

Step 1: Lifting

$$\mathcal{L}_\infty := \left\{ (K, \gamma, P) : P \succ 0, \begin{bmatrix} (A+BK)^\mathsf{T}P + P(A+BK) & PB_w & C^\mathsf{T} \\ B_w^\mathsf{T}P & -\gamma I & 0 \\ C & 0 & -\gamma I \end{bmatrix} \preceq 0 \right\},$$

State-feedback Robust Control

❑ Building an ECL

Step 2: Convex set

$$\mathcal{F}_\infty = \left\{ (\gamma, Y, X) \,\middle|\, \begin{array}{r} X \succ 0, \\ Y \in \mathbb{R}^{m \times n}, \end{array} \begin{bmatrix} AX + XA^\mathsf{T} + BY + Y^\mathsf{T}B^\mathsf{T} & B_w & XQ^{1/2} & Y^\mathsf{T}R^{1/2} \\ B_w^\mathsf{T} & -\gamma I & 0 & 0 \\ Q^{1/2}X & 0 & -\gamma I & 0 \\ R^{1/2}Y & 0 & 0 & -\gamma I \end{bmatrix} \preceq 0 \right\},$$

Step 3: Diffeomorphism $\Phi(K, \gamma, P) = (\gamma, KP^{-1}, P^{-1}), \qquad \forall (K, \gamma, P) \in \mathcal{L}_\infty.$

- No auxiliary set
- Lifted set satisfies $\pi_{K,\gamma}(\mathcal{L}_\infty) = \mathrm{epi}_{\geq}(J_\infty)$

➡ **All policies are non-degenerate**

Theorem: Any Clarke stationary points are globally optimal.

Linear Quadratic Gaussian (LQG) Control

❑ **Construction of the ECL**: Based on the convexification proposed in [Scherer et al., 1997]

Theorem.
1. An ECL for LQG exists, of which the auxiliary set **is the set of invertible matrices.**
2. A policy K is non-degenerate if and only if it is **informative** in the sense that $\lim_{t\to\infty} \mathbb{E}\left[x(t)\xi(t)^{\mathsf{T}}\right]$ has full rank. So **any informative stationary point is globally optimal**.

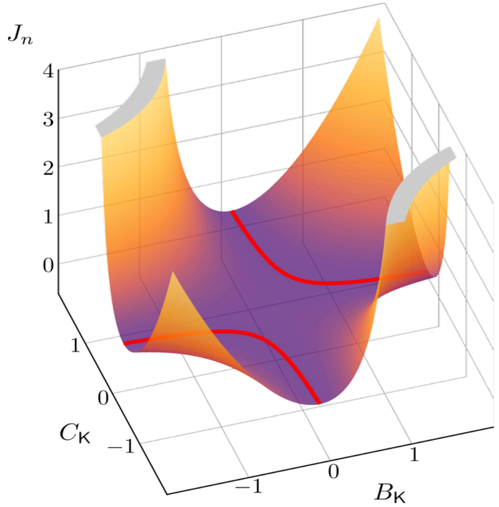

- We also show that **controllable and observable stationary policies are informative**, and are thus globally optimal [Tang, **Zheng**, Li, 2023].
- Similar ECL guarantees hold for output-feedback Hinf control and other control problems [**Zheng**, Pai & Tang, 2024].

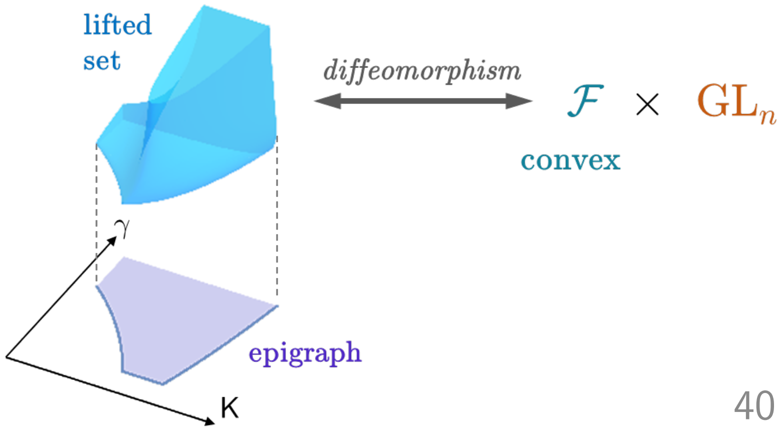

Outline

Geometry:
Benign Nonconvex Landscape

Framework:
Extended Convex Lifting

Algorithm:
Proximal Descent Method

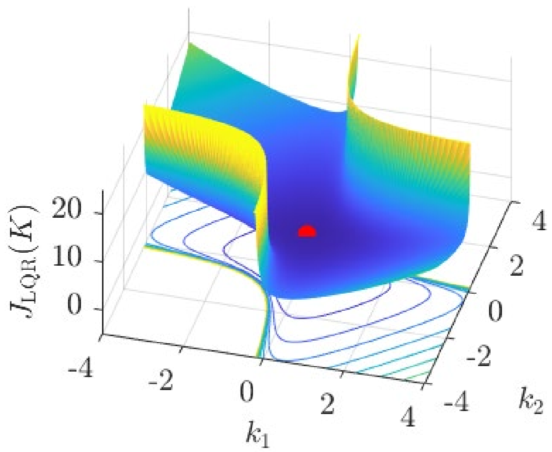

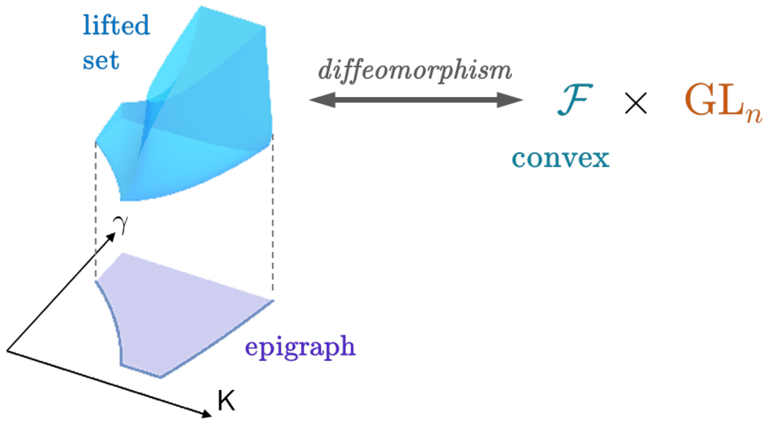

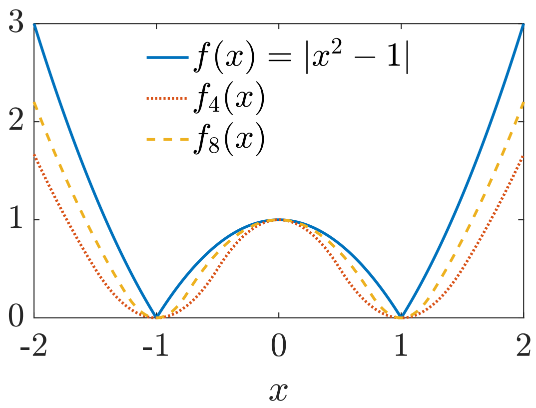

Algorithms for (non)-convex optimization

$$\min_x f(x)$$

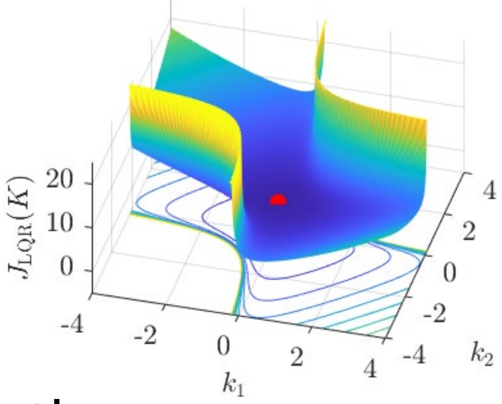

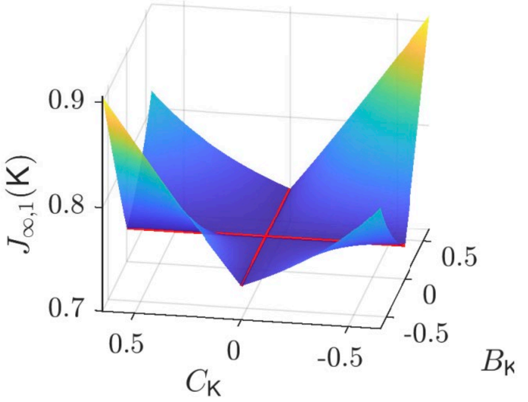

Standard algorithms:

- Gradient descent for L-smooth (non)-convex functions

$$x_{k+1} = x_k - \alpha_k \nabla f(x_k), \qquad 0 < \alpha_k < \frac{2}{L}$$

Very careful step size tuning

- Subgradient method for non-smooth convex functions

$$x_{k+1} = x_k - \alpha_k g_k, \qquad g_k \in \partial f(x_k)$$

Non-trivial to extend for nonsmooth/nonconvex case

❑ They achieve the best first-order complexity for convex functions.

Proximal Descent Method

❑ A conceptual descent algorithm

$$x_{k+1} = x_k - \alpha \times g_k$$

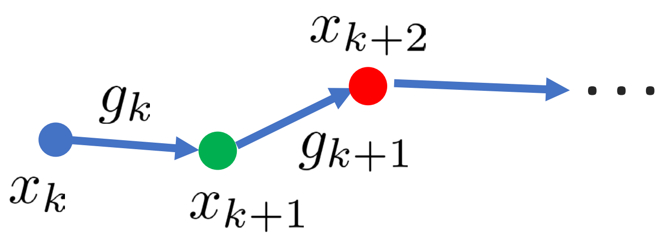

Suppose it satisfies a cost value drop

$$f(x_{k+1}) \leq f(x_k) - \frac{\alpha}{2} \|g_k\|^2$$

Complexity: the conceptual descent algorithm finds a point such that

$$\min_{k \leq T} \|g_k\|^2 \leq \frac{f(x_0) - f^\star}{T} \times \frac{2}{\alpha}$$

- Valid for any positive step size

When such a cost value drop is achievable? And How?

Weakly convex functions

Proximal bundle procedure

Feng-Yi Liao, and Yang Zheng. "**A proximal descent method for minimizing weakly convex optimization**." *arXiv preprint arXiv:2509.02804* (2025).

Proximal Descent Method

❑ Weakly convex function

$$f(x) + \frac{m}{2}\|x\|^2 \quad \text{is convex}$$

- Any **convex (nonsmooth)** functions
- Any **L-smooth (nonconvex)** functions
- **LQR** and **state-feedback Hinf control** under mild assumptions

---

Algorithm 1 Proximal descent method

Require: $x_1 \in \mathbb{R}^n$, $T > 0$,
  for $k = 1, 2, \ldots, T$ **do**
    $x_{k+1} = \texttt{ProxDescent}(x_k)$;
  end for

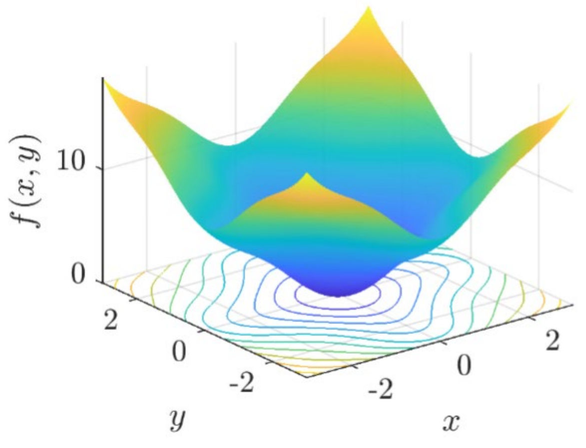

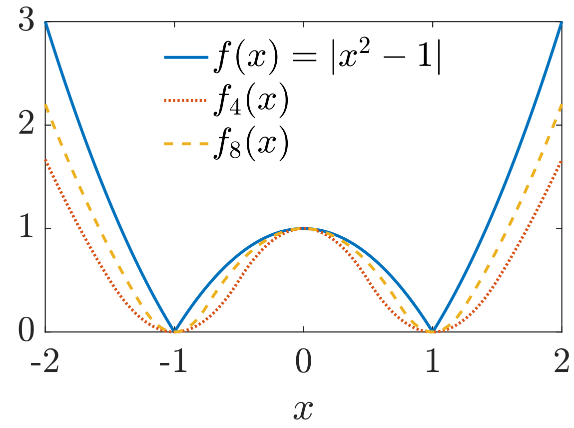

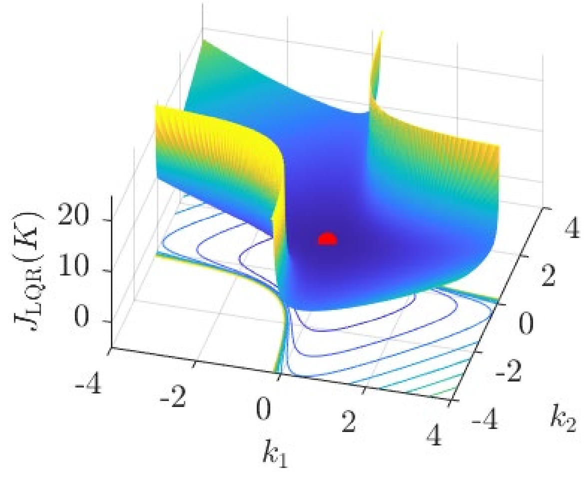

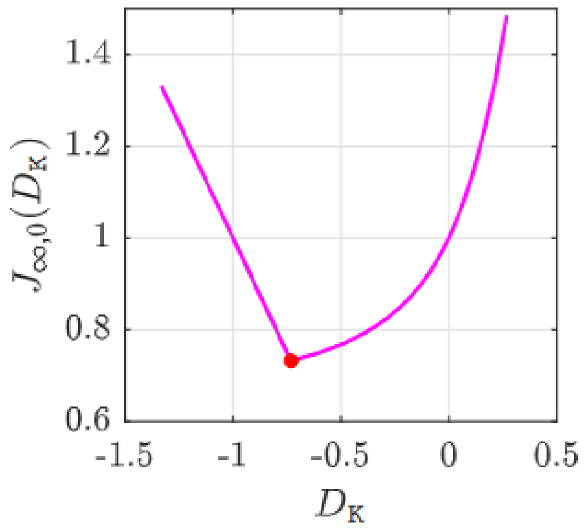

Convergence guarantees

$$\min_x f(x)$$

function	Growth	Measure	Complexity
L-Lipchitz	No growth	$\|\nabla f_\alpha(x_k)\| \leq \epsilon$	$\mathcal{O}\left(1/\epsilon^4\right)$
	μ_q-QG and $\mu_q > m$	$\|\nabla f_\alpha(x_k)\| \leq \epsilon$	$\mathcal{O}\left(1/\epsilon^2\right)$
M-Smooth	No growth	$\|\nabla f(x_k)\| \leq \epsilon$	$\mathcal{O}\left(1/\epsilon^2\right)$
	μ_q-QG and $\mu_q > m$	$\|\nabla f(x_k)\| \leq \epsilon$	$\mathcal{O}\left(\log\left(1/\epsilon\right)\right)$

- Holds for m-weakly convex functions, including both **convex** and **nonconvex** problems
- **Automatic adaptivity**: with no parameter tuning, **the algorithm accelerates** in the presence of smoothness and/or quadratic growth

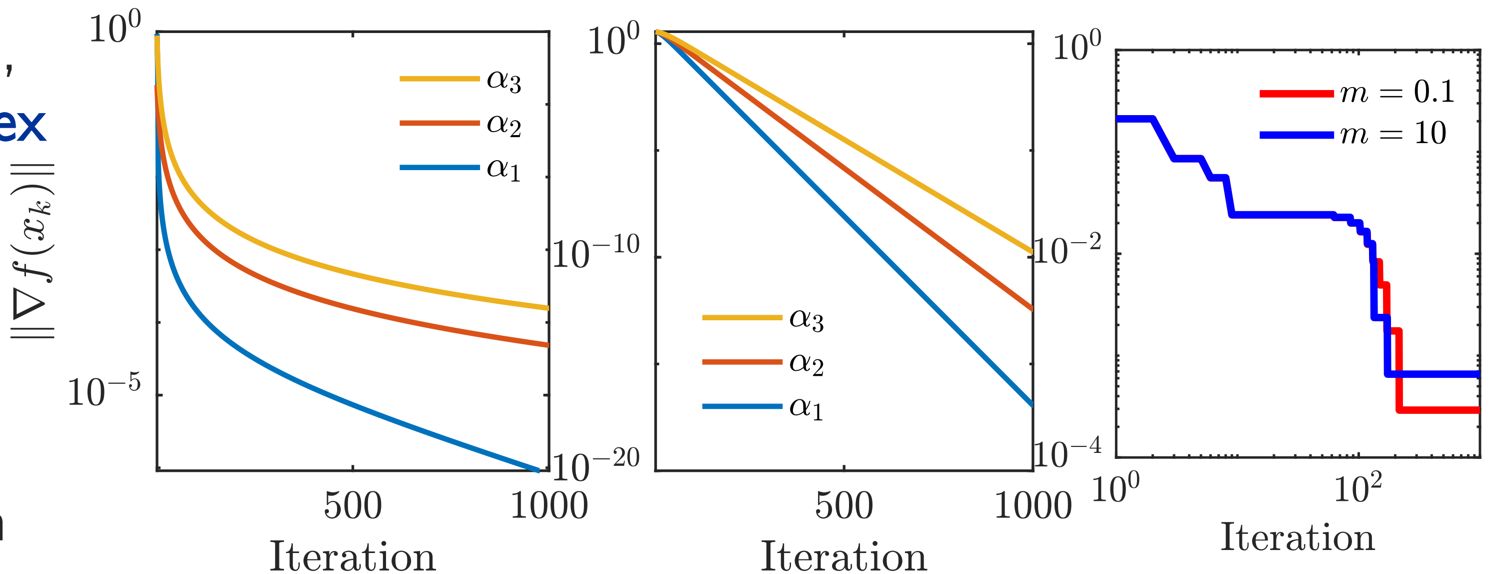

Smooth non-convex examples

Non-smooth example

Conclusion

Geometry:
Benign Nonconvex Landscape

Framework:
Extended Convex Lifting

Algorithm:
Proximal Descent Method

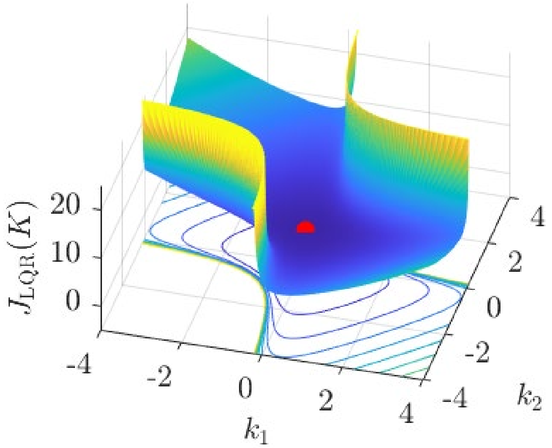

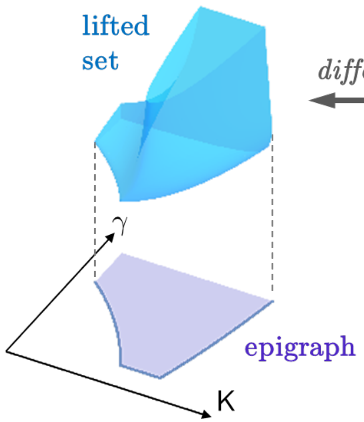

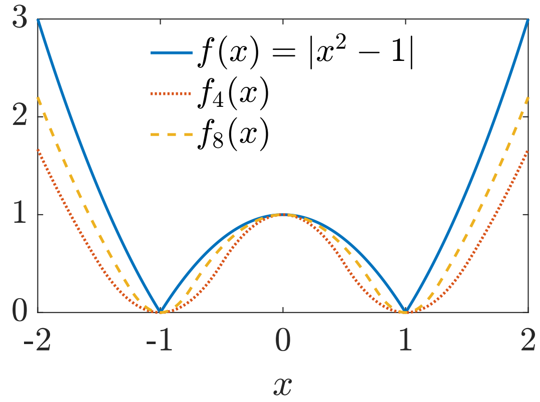

Local Stationarity + **Structural Information**

Global Optimality Certificate

Ongoing and Future work

- ❑ How to deal with degenerate points in local policy search? Avoiding **saddle points** with guarantees?
- ❑ Extend the analysis to **distributed control**?
- ❑ Establish guarantees **beyond linear control**?

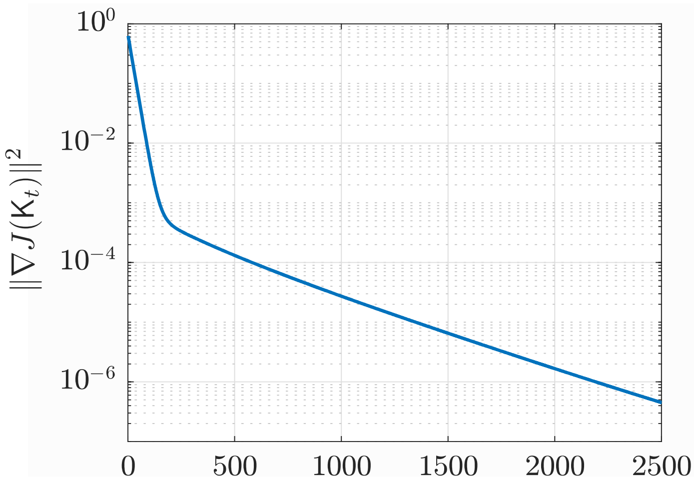

Control Theory ⟷ **(Non)convex Optimization**

Reinforcement Learning

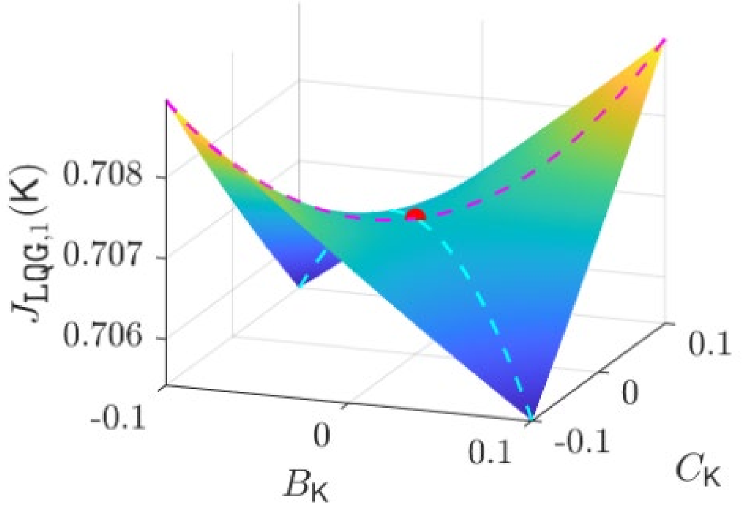

Selected references

- Yujie Tang, Yang Zheng, and Na Li. "**Analysis of the optimization landscape of Linear Quadratic Gaussian (LQG) control.**" *Mathematical Programming* 202.1 (2023): 399-444.
- Yang Zheng, Chih-Fan Pai, and Yujie Tang. "**Benign Nonconvex Landscapes in Optimal and Robust Control, Part I: Global Optimality.**" IEEE Transactions on Automatic Control (2026): https://arxiv.org/abs/2312.15332.
- Yang Zheng, Chih-Fan Pai, and Yujie Tang. "**Benign Nonconvex Landscapes in Optimal and Robust Control, Part II: Extended Convex Lifting.**" IEEE Transactions on Automatic Control (2026): https://arxiv.org/abs/2406.04001
- Yuto Watanabe, Feng-Yi Liao, and Yang Zheng. "**Policy Optimization in Robust Control: Weak Convexity and Subgradient Methods.**" *arXiv preprint arXiv:2509.25633* (2025). https://arxiv.org/abs/2509.25633
- Liao, Feng-Yi, and Yang Zheng. "**A proximal descent method for minimizing weakly convex optimization.**" *arXiv preprint arXiv:2509.02804* (2025).

Thank you for your attention!

Q & A