UC San Diego
JACOBS SCHOOL OF ENGINEERING

Convex and Nonconvex Optimization for Control

Yang Zheng

Department of Electrical and Computer Engineering,
University of California San Diego

May 29, 2026

MAY 26-29 2026 ACC AMERICAN CONTROL CONFERENCE 2026 New Orleans

Control as the hidden technology[1,2]

A talk given by Prof. Karl J. Åström in 2006 [2]

- [1]Åström, Karl J. "Automatic control - the hidden technology." *Advances in control: highlights of ECC'99*. Springer London, 1999. 1-28.
- [2]https://archive.control.lth.se/media/Staff/KarlJohanAstrom/Lectures/HiddenTechnologyMIT2006.pdf

From the talk[2]:

- Widely used; Very successful; Seldom talked about
- Except when disaster strikes

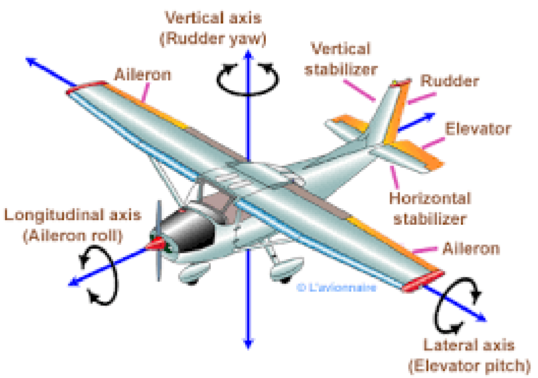

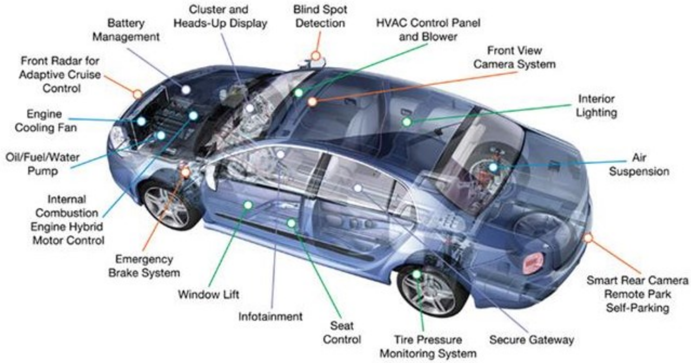

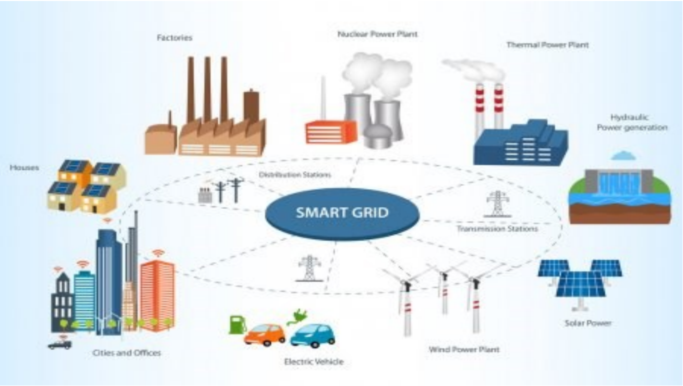

Control as the hidden technology[1,2]

Classical control combines **modeling**, **feedback**, and **robustness** with rigorous **guarantees**.

Enabled by Riccati theory, LMIs, and **convex analysis/optimization**, etc.

From the talk[2]:

- Widely used; Very successful; Seldom talked about
- Except when disaster strikes

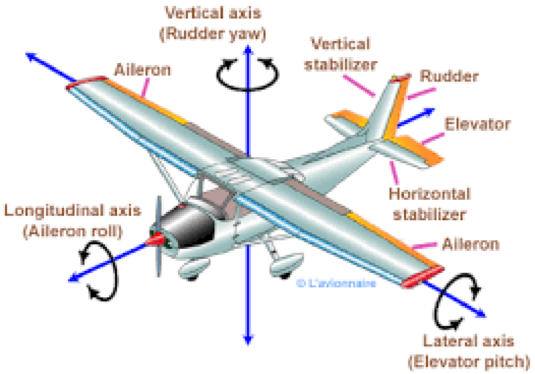

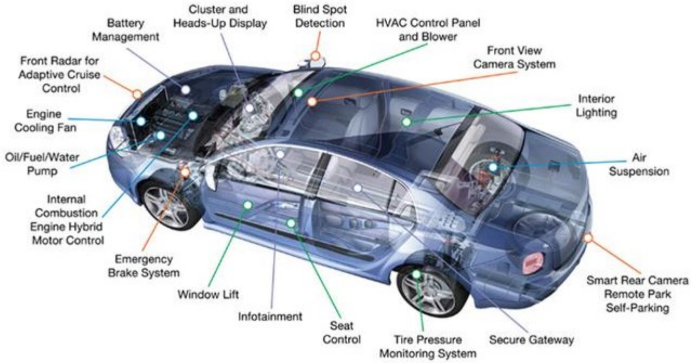

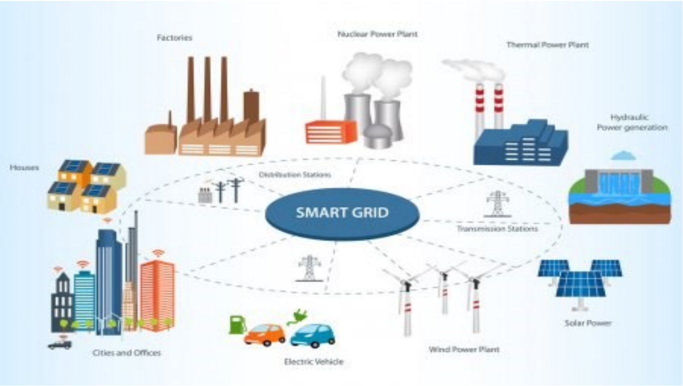

A new success story: learning-based decision making

- Data-driven decision making and Reinforcement learning have achieved striking performance in games, robotics, sequential decision problems, and language models.
- These successes often rely on massive **data**, **computation**, and **flexible policy parameterizations**.
- They shift the design paradigm from "**derive a controller**" to "**optimize a policy**".

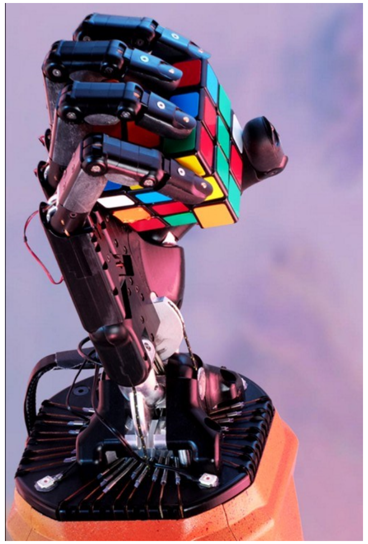

ChatGPT

Policy optimization

Feedback control

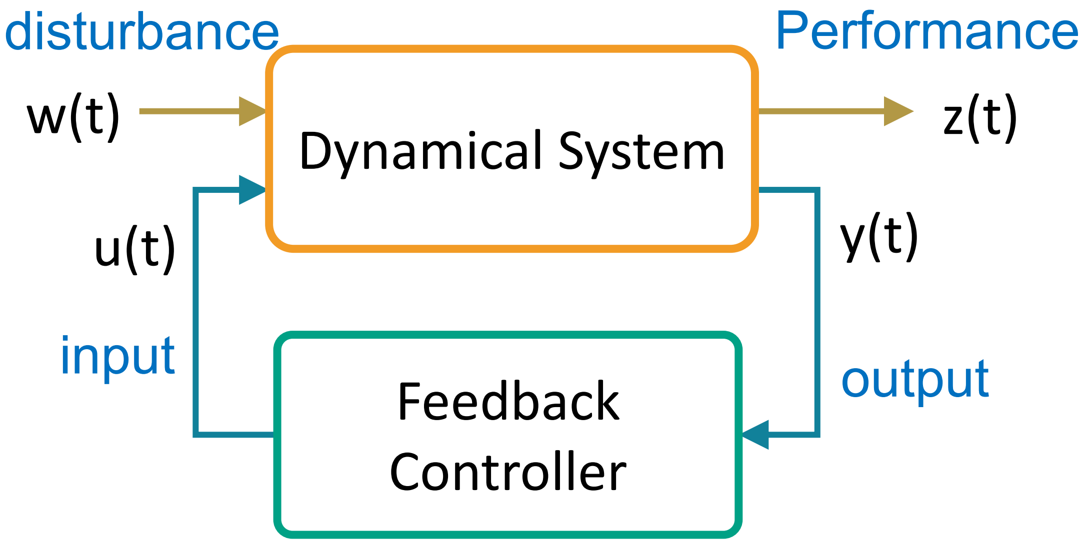

Policy parametrization

Policy Optimization

$$\min_{\mathsf{K}} \; J(\mathsf{K})$$

$$\text{s.t.} \;\; \mathsf{K} \in \mathcal{C}$$

Use direct policy updates (**model-free**)

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

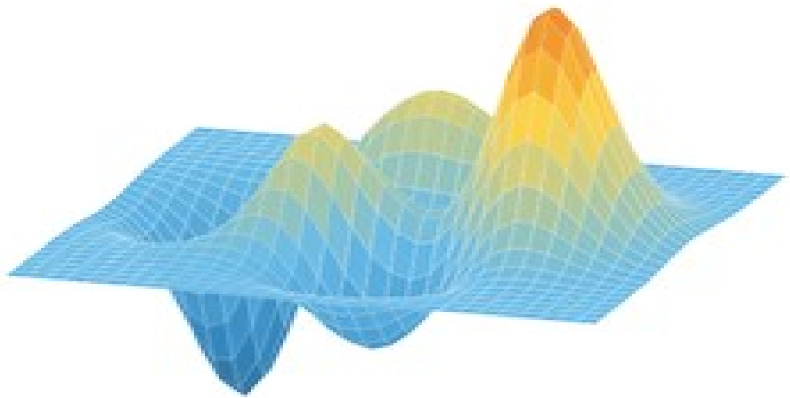

- **Nonconvex optimization**
- Lack of principled algorithms for **optimality**
- Hard to derive **theoretical guarantees**

Two optimization frameworks for control

❑ Convex optimization

- LMIs, Semidefinite programs, etc.
- Have become popular since 1980s
  - ✓ [Boyd et al., 1994] [Scherer & Weiland, 2000] [Dullerud & Paganini, 2013], etc.
- Enjoy global guarantees and reliable interior-point solvers
- Model-based re-parameterizations (do not optimize over controller space directly)

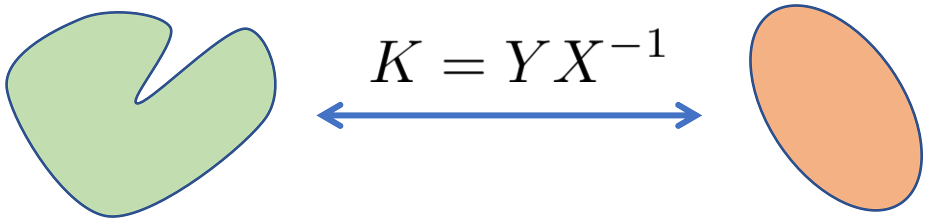

❑ Nonconvex policy optimization

- Policy gradient, direct search, etc.
- Have a long history in control theory
  - ✓ [Levine and Athans, 1970] [Apkarian & Noll, 2006] [Saeki, 2006] [Apkarian et al., 2008] [Arzelier et al., 2011], etc.
  - ✓ Packages: HIFOO, hinfstruct
- Model-free with great empirical performance (scalability, flexibility)
- Weak guarantees in general; not popular for control theorists in the past

Two optimization frameworks for control

❑ Renewed interest and recent progress on nonconvex policy optimization

- Favorable properties have been revealed for a range of benchmark problems:

✓ LQR [Fazel et al., 2019; Mohammadi et al., 2022; Fatkhullin & Polyak, 2021; Zhang, Hu&Basar, 2020; Cui, Jiang & Sontag, 2024; Talebi & Mesbahi, 2023; Jansch-Porto, Hu & Dullerud, 2022; WZ, SICON 2026] etc.

✓ LQG [Kraisler & Mesbahi, 2024; Sadamoto & Tanaka, 2026; TZL, Math Prog. 2023; ZPT, TAC 2026] etc.

✓ $\mathcal{H}_\infty$ robust control, [Guo & Hu, 2022; HZ, L-CSS 2023; ZPT, TAC 2026] etc.

Nonconvex policy space

Convex space

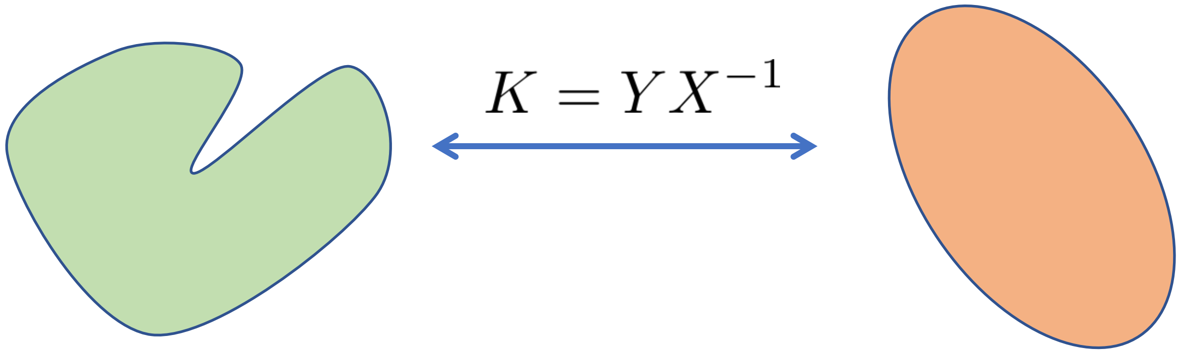

Are these two optimization frameworks separate for control?

Of course, Not

We highlight a clean and precise connection, for many benchmark control problems!

Today's talk: A unifying perspective

❑ Geometry: Nonconvex Landscapes in Benchmark Control Problems are Benign

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

  [Fazel et al., 2018], [Mohammadi et al., 2022] [Fatkhullin & Polyak, 2021], [Guo & Hu, 2022], [Sontag, 2025], [TZL, Math Prog. 2023], [HZ, L-CSS 2023]; [ZPT, TAC 2026], [WZ, SICON 2026] etc.

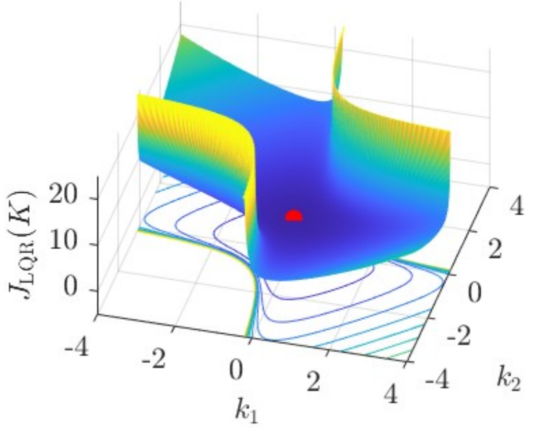

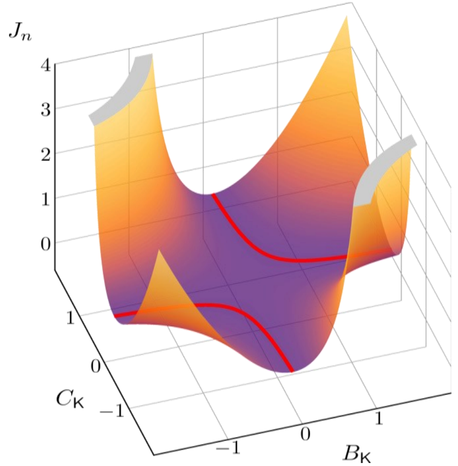

❑ Framework: Extended Convex Lifting

- **Bridge non-convex policy optimization with convex reformulation!**

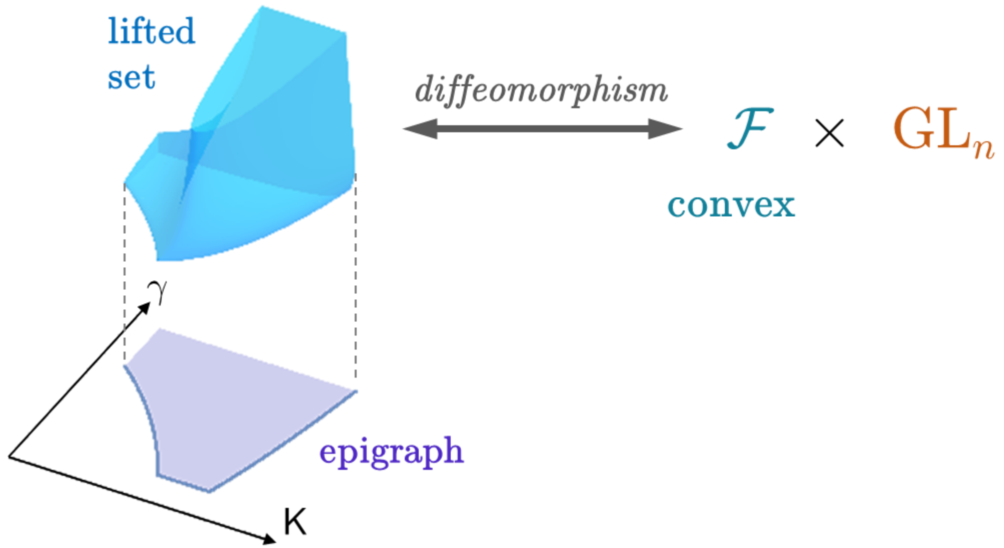

❑ Algorithm: Decomposition for scalability & proximal descent methods

Outline

Geometry: Benign Nonconvex Landscape

Framework: Extended Convex Lifting

Algorithm: Decomposition & Proximal Descent

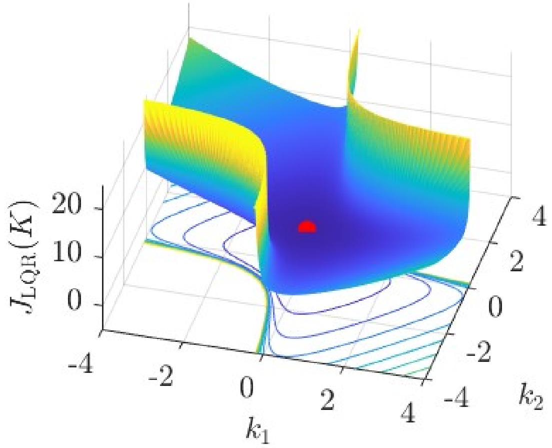

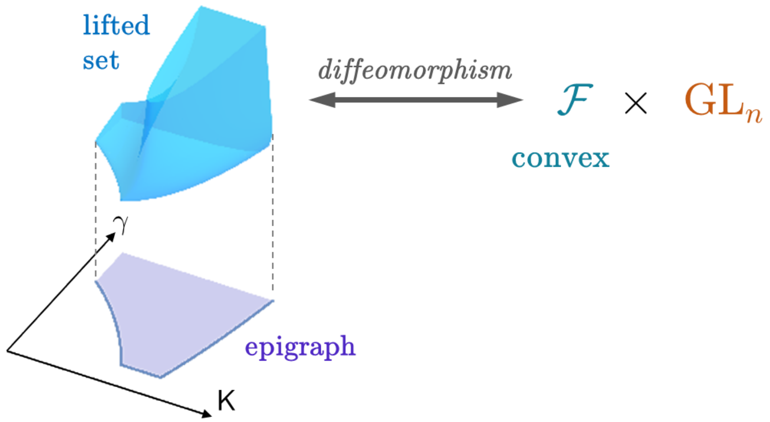

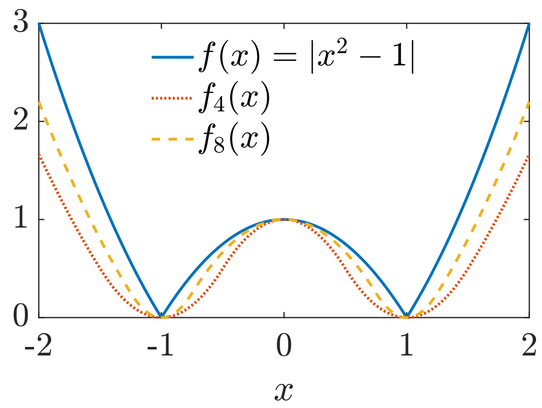

Non-convexity in control

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

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

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

$$\text{subjct to} \quad x_{t+1} = Ax_t + Bu_t, \longrightarrow \text{Linear dynamics}$$

$$u_t = Kx_t \longrightarrow \text{Linear static policy}$$

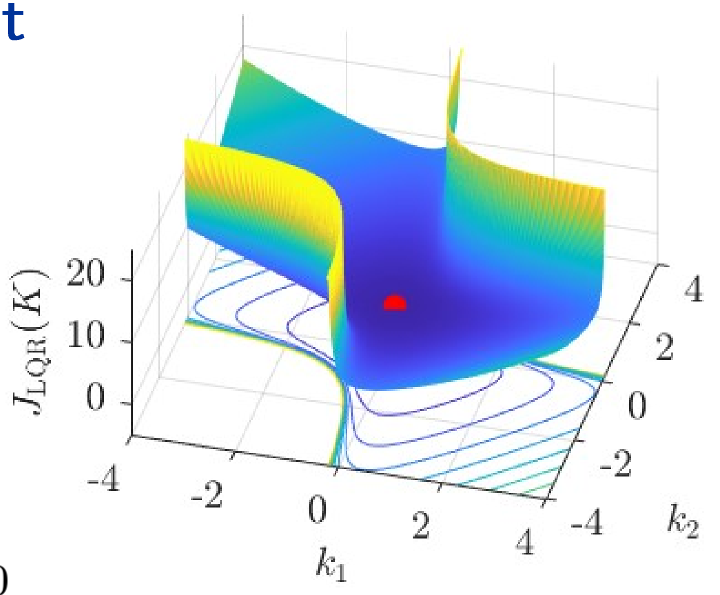

- Closed-loop state trajectories are nonlinear in K

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

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

- The cost function 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 closed-loop trajectory converges to the origin.

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

  - ✓ The spectral radius of the closed-loop matrix should be less than 1.

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

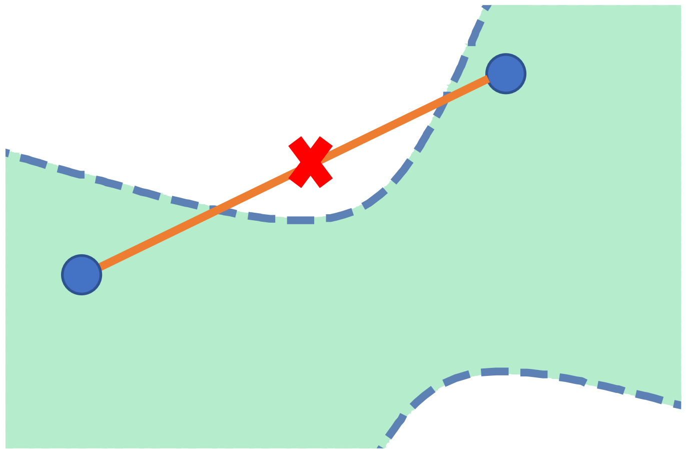

  - ✓ A simple example: $A = 0, \quad B = I_2$

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

Non-convexity in control

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

Fact: The set of dynamic stabilizing policies is **nonconvex** and may even be **disconnected**. [TZL, *Math. Prog.* 2023].

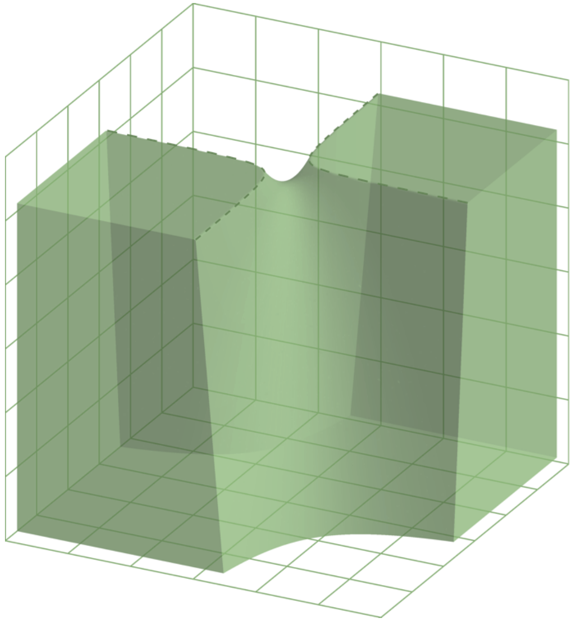

Example 1: Nonconvex but still **connected**

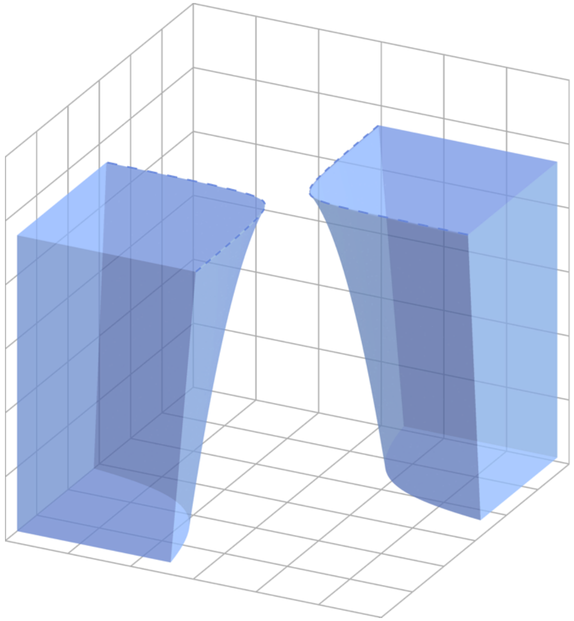

Example 2: Nonconvex and **disconnected**

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, $\mathcal{H}_2$ / $\mathcal{H}_\infty$ optimal control, etc.

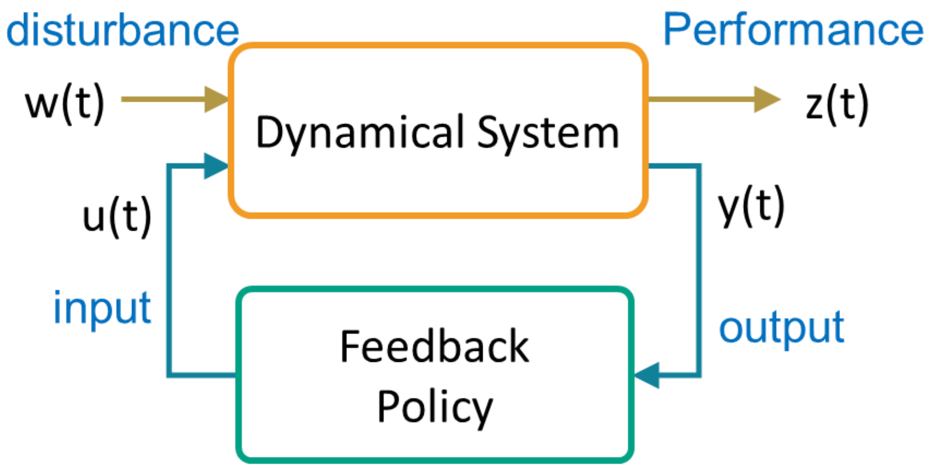

❑ Linear Quadratic Regulator (LQR)

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

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

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

Linear Matrix inequality

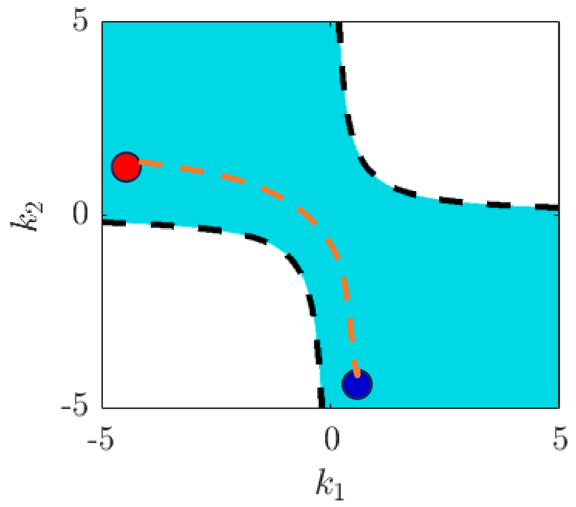

Any convex set is path-connected, and thus so is its continuous image.

Benign non-convexity in LQR

- Linear Quadratic Regulator (LQR)

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

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

Fact: **Any stationary point $\nabla J(K) = 0$ is globally optimal,** under a mild assumption (Fazel et al., 2019; Mohammadi et al., 2022).

Local Stationarity → Global Optimality

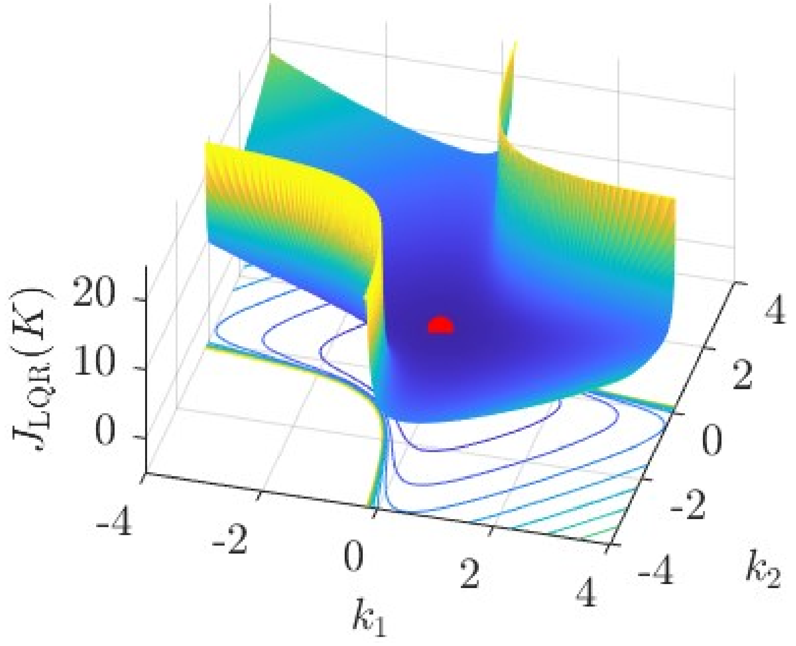

Benign non-convexity in LQR

- ❑ Linear Quadratic Regulator (LQR)

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

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

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

- LQR almost behaves like a strongly convex problem under mild assumptions (WZ, SICON 2026).
- The basic gradient descent algorithm achieves linear convergence (Fazel et al., 2019; Mohammadi et al., 2022; Fatkhullin & Polyak, 2021)

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

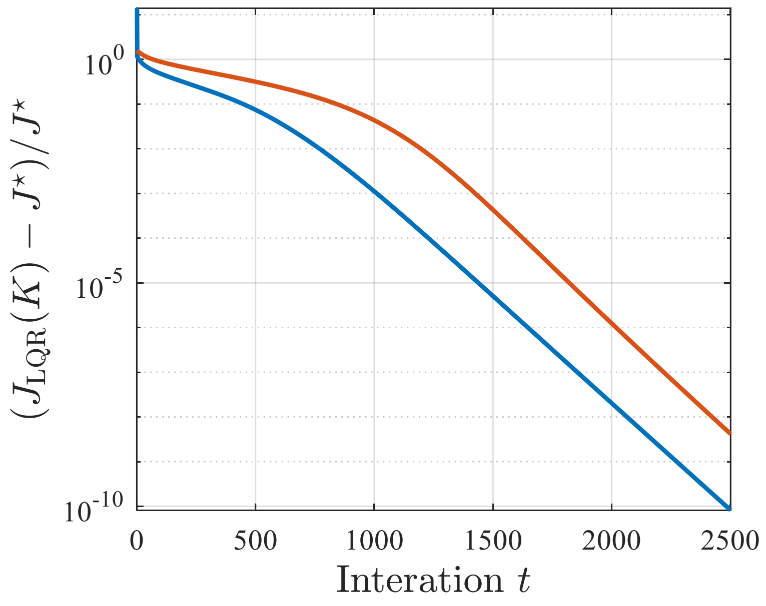

Benign non-convexity in robust control

- State-feedback $\mathcal{H}_\infty$ robust control

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

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

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

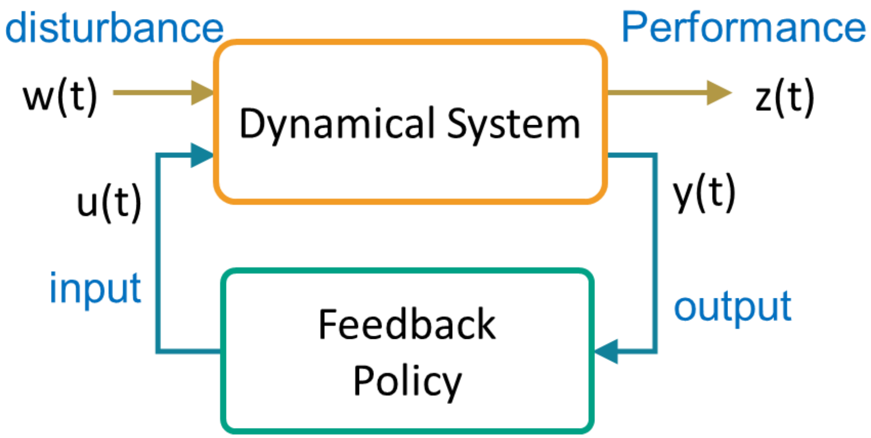

Unlike LQR with stochastic noise, we here assume **adversarial disturbances** with bounded energy and consider the **worst-case performance**

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 $\mathcal{H}_\infty$ 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: The $\mathcal{H}_\infty$ cost is generally **non-convex**, and **non-smooth**, but **locally weakly convex** (WLZ, ACC 2026).

Theorem: Any sublevel set is path-connected, and **any stationary point is globally optimal** (Guo and Hu, 2022; ZPT, TAC 2026)

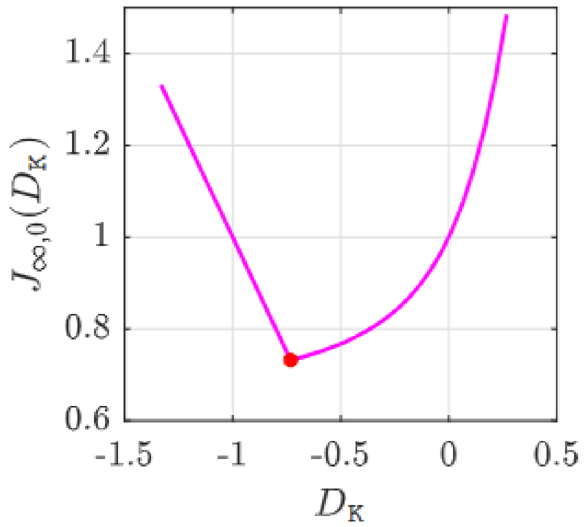

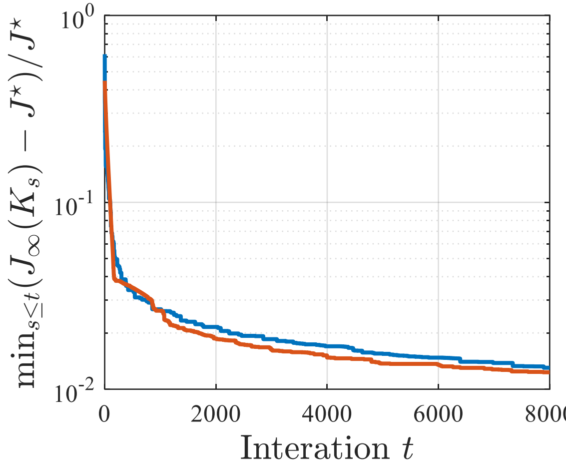

- We may use the subgradient method to find a global minimizer despite the non-convexity (WLZ, ACC 2026)

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

Benign non-convexity in LQG

❑ Linear Quadratic Gaussian (LQG) Control

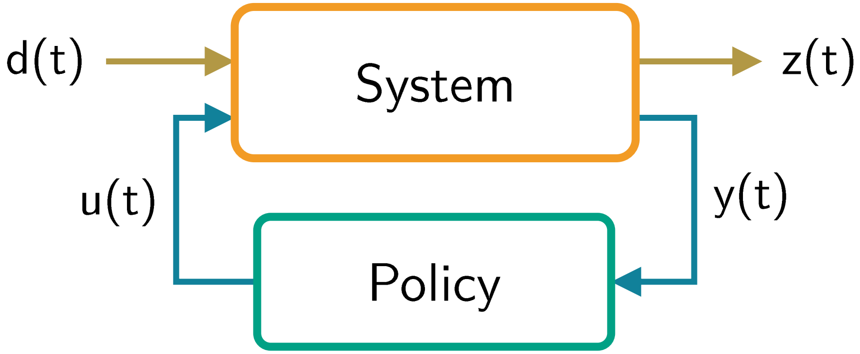

Dynamics: $$\dot{x}(t) = Ax(t) + Bu(t) + B_w w(t)$$
$$\boxed{y(t) = Cx(t) + D_v v(t)}$$

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 non-convexity in LQG

❑ Linear Quadratic Gaussian (LQG) Control

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

Policy optimization for LQG

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

Theorem: The set $\mathcal{C}_{\text{full}}$ can be **disconnected** but has at most **2 connected** components that are algebraically identical [TZL, *Math. Prog.* 2023]

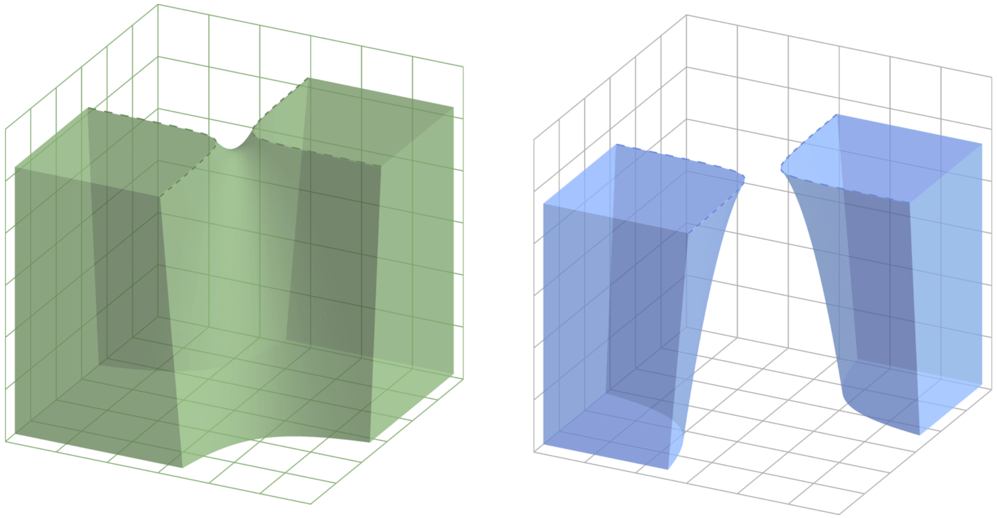

Proof: convexification via a nonlinear change of variables

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

Continuous surjective map

at most 2 connected components

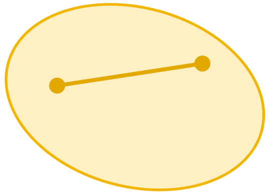

Convex thus connected

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

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 non-convexity in LQG

❑ Linear Quadratic Gaussian (LQG) Control

Theorem: For a controllable and observable policy, if it is a **stationary** point $\nabla J(\mathsf{K}) = 0$, then it is a **global minimizer** [TZL, *Math. Prog.* 2023; Hyland and Bernstein, TAC 1984].

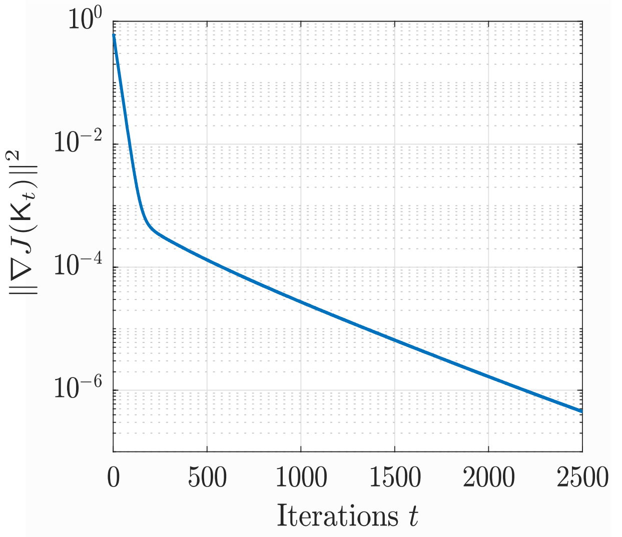

Local Zero Gradient + Structural Information → Global Optimality Certificate

Fact: For policies that are uncontrollable or unobservable, there exist suboptimal stationary points (e.g., **saddle points**) [TZL, *Math. Prog.* 2023; ZPT, TAC 2026].

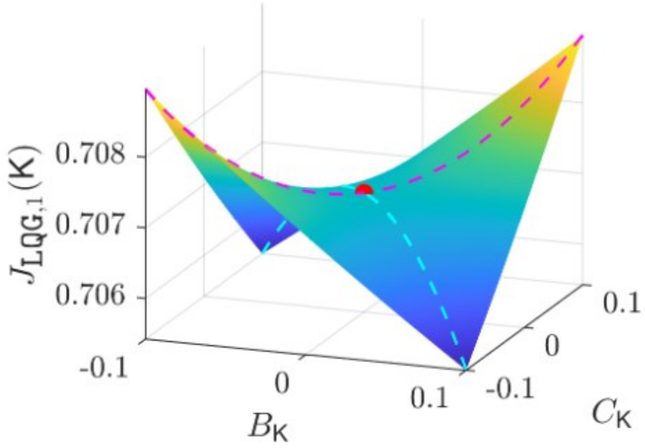

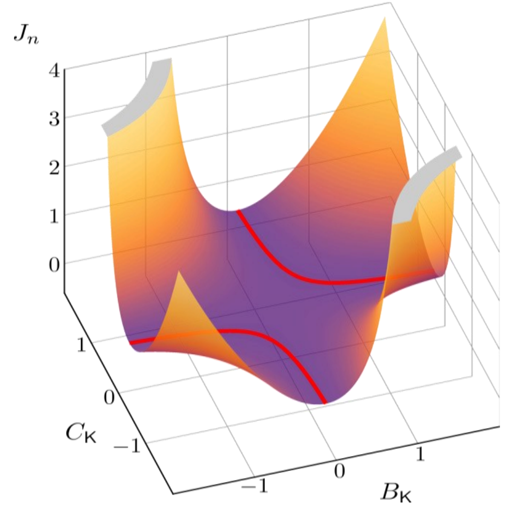

Other benign non-convexity

Zheng, C. Pai, and Y. Tang. "Benign nonconvex landscapes in optimal and robust control, Part I: Global optimality." *IEEE Transactions on Automatic Control* (2026), pp. 1-16. https://arxiv.org/abs/2312.15332

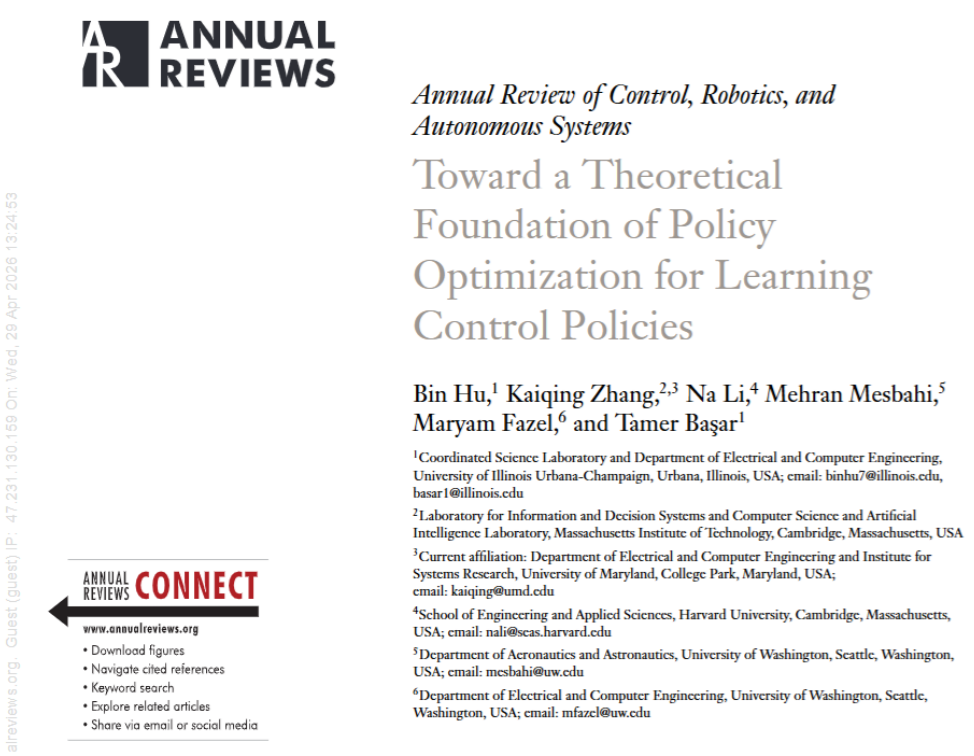

ANNUAL REVIEWS

Annual Review of Control, Robotics, and Autonomous Systems

Toward a Theoretical Foundation of Policy Optimization for Learning Control Policies

Bin Hu,[1] Kaiqing Zhang,[2,3] Na Li,[4] Mehran Mesbahi,[5] Maryam Fazel,[6] and Tamer Başar[1]

[1]Coordinated Science Laboratory and Department of Electrical and Computer Engineering, University of Illinois Urbana-Champaign, Urbana, Illinois, USA; email: binhu7@illinois.edu, basar1@illinois.edu

[2]Laboratory for Information and Decision Systems and Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

[3]Current affiliation: Department of Electrical and Computer Engineering and Institute for Systems Research, University of Maryland, College Park, Maryland, USA; email: kaiqing@umd.edu

[4]School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts, USA; email: nali@seas.harvard.edu

[5]Department of Aeronautics and Astronautics, University of Washington, Seattle, Washington, USA; email: mesbahi@uw.edu

[6]Department of Electrical and Computer Engineering, University of Washington, Seattle, Washington, USA; email: mfazel@uw.edu

Hu, B., Zhang, K., Li, N., Mesbahi, M., Fazel, M., & Başar, T. (2023). Toward a theoretical foundation of policy optimization for learning control policies. *Annual Review of Control, Robotics, and Autonomous Systems*, *6*(1), 123 158.

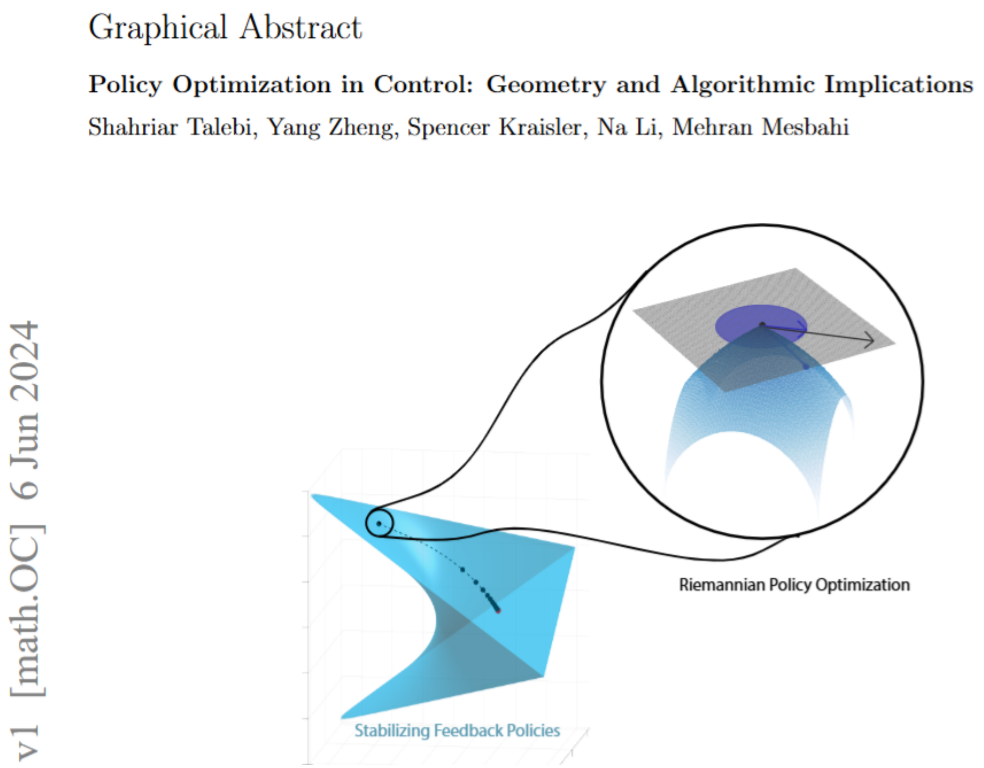

Graphical Abstract

Policy Optimization in Control: Geometry and Algorithmic Implications

Shahriar Talebi, Yang Zheng, Spencer Kraisler, Na Li, Mehran Mesbahi

Talebi, S., **Zheng**, Kraisler, S., Li, N., & Mesbahi, M. (2024). Policy Optimization in Control: Geometry and Algorithmic Implications. In *Encyclopedia of Systems and Control Theory*. Elsevier.

Outline

Geometry: Benign Nonconvex Landscape

Framework: Extended Convex Lifting

Algorithm: Decomposition & Proximal Descent

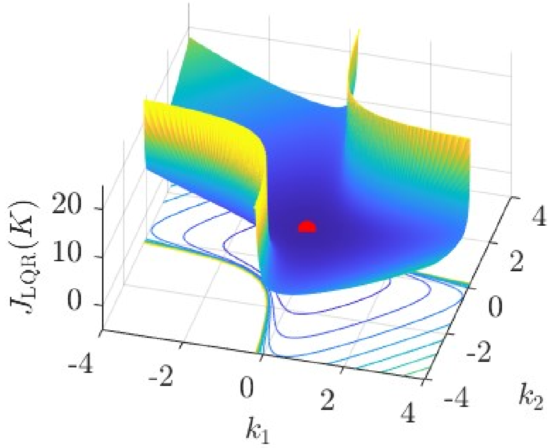

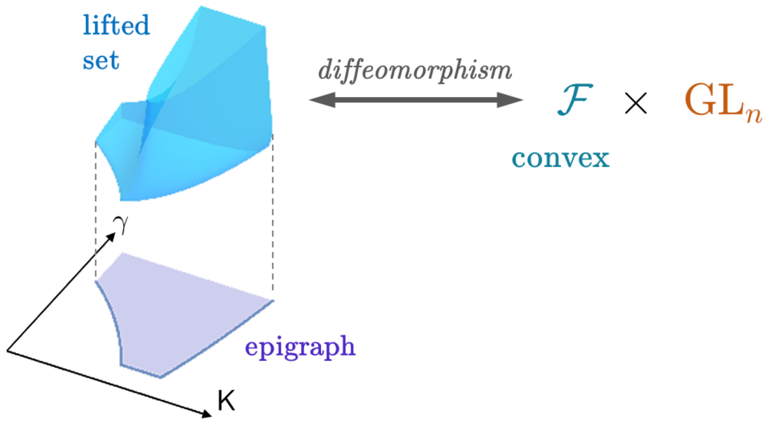

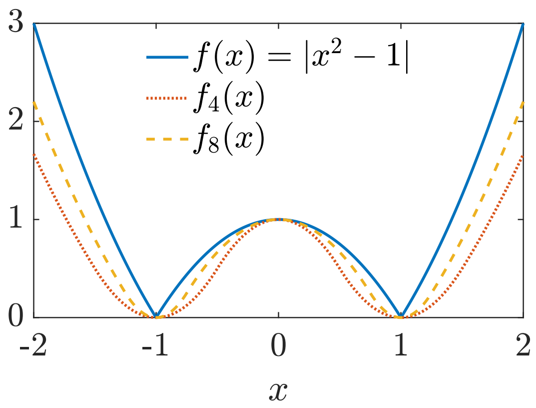

Zheng, C. Pai, and Y. Tang. "**Benign Nonconvex Landscapes in Optimal and Robust Control, Part II: Extended Convex Lifting**." IEEE Transactions on Automatic Control (2026), pp. 1-16: https://arxiv.org/abs/2406.04001

Key idea

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

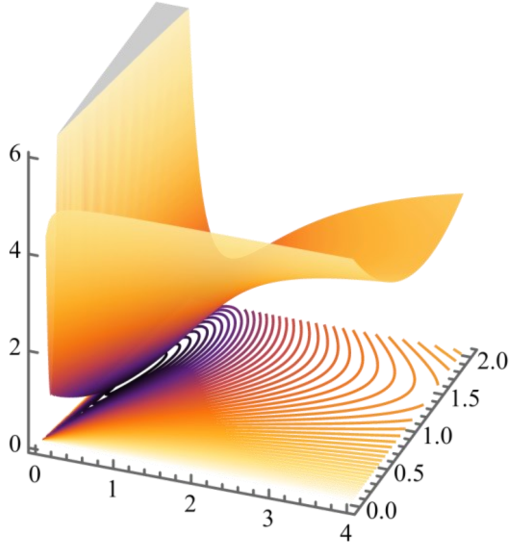

Nonconvex policy optimization ⟷ **LMI-based convex reformulation**

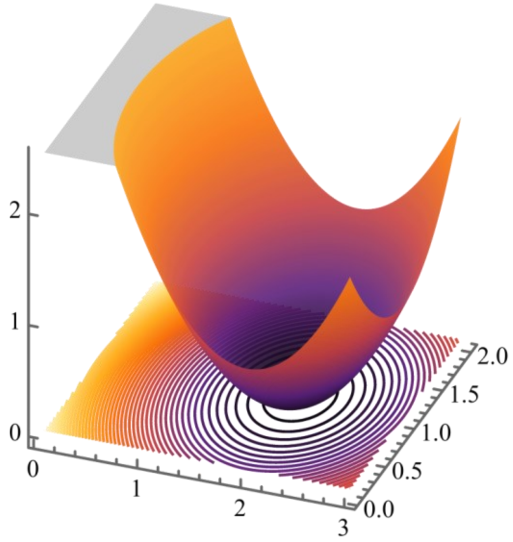

- Our idea: Exploit **LMI-based convex reformulations** of control problems
- They reveal the **hidden convexity** in policy optimization landscapes.
- Under suitable **nondegenerate** conditions, **all Clarke stationary points are globally optimal** and there is **no spurious local minimum**.

Example 1

❑ **Nonconvex and smooth function**

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

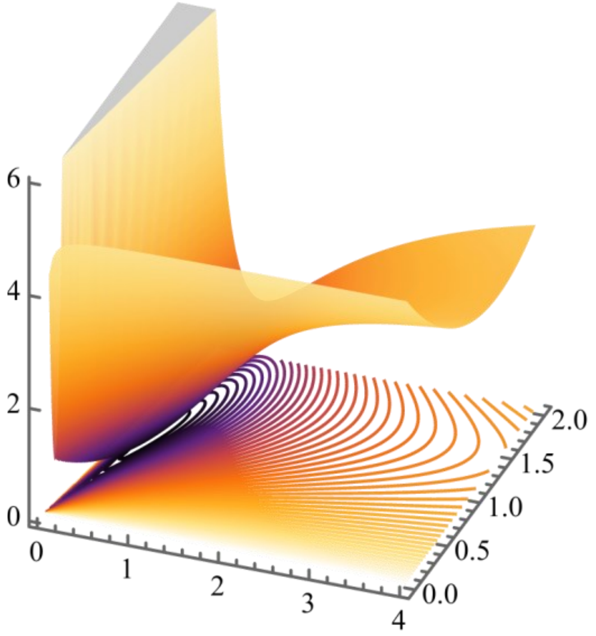

Its global **minimizer** is

$$x_1^* = 0.5, x_2^* = 1$$

Define an **invertible map**

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

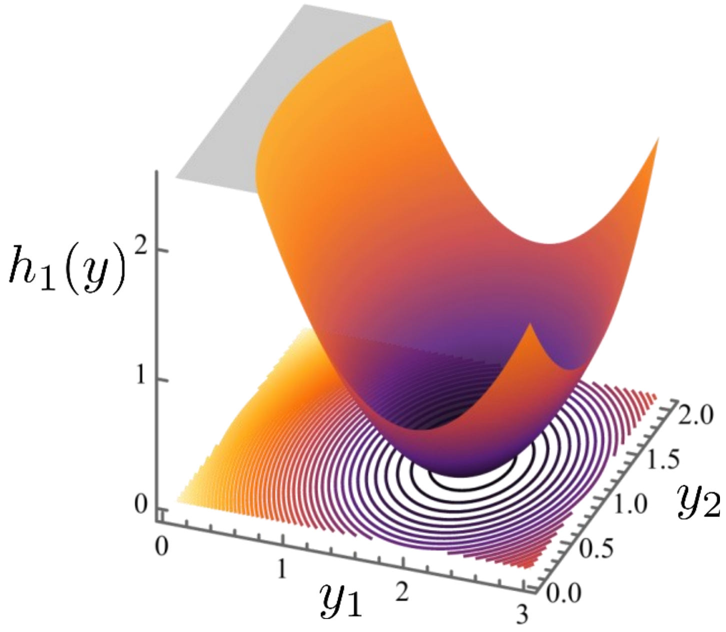

$$h_1(y) := f_1(g^{-1}(y)) = (y_1 - 2)^2 + (y_2 - 1)^2, \qquad \forall y_1 > 0, y_2 > 0.$$

Example 2

- Nonconvex and non-smooth function

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

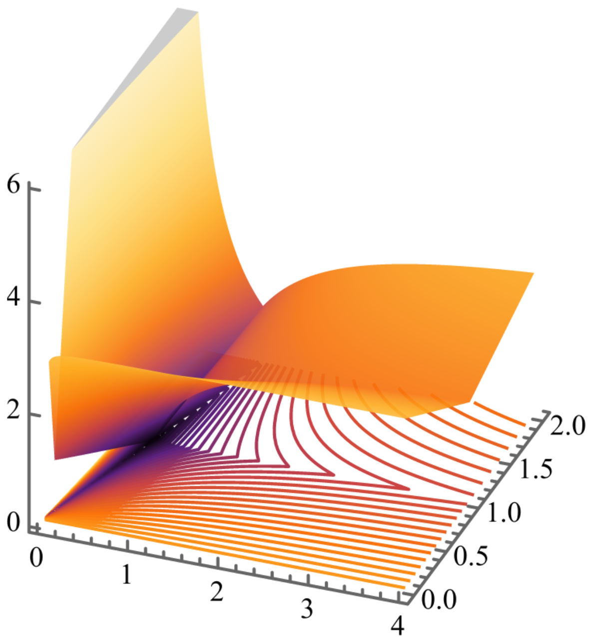

Its global **minimizer** is

$$x_1^* = 0.5, x_2^* = 1$$

Define an **invertible map**

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

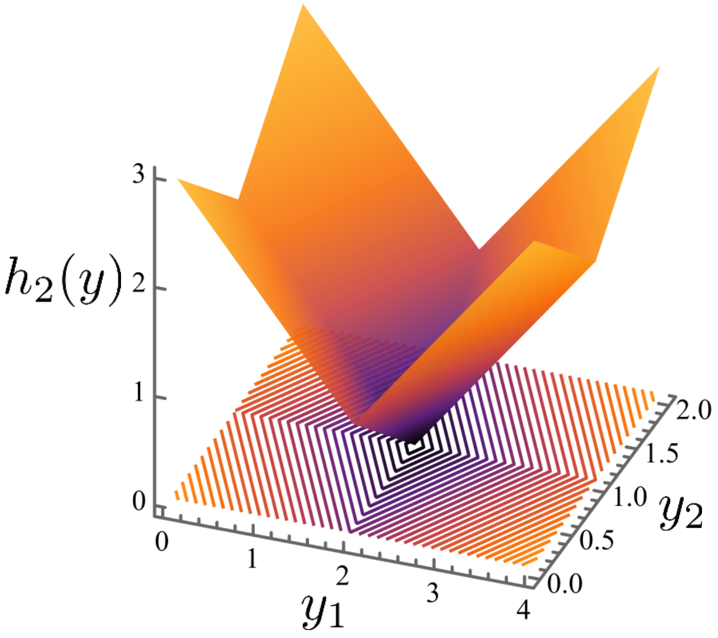

$$h_2(y) := f_2(g^{-1}(y)) = |y_1 - 2| + |y_2 - 1|, \qquad \forall y_1 > 0, y_2 > 0,$$

Example 3

- Linear Quadratic Regulator (LQR)

$$J(k_1,k_2) = \frac{1-2k_2+3k_2^2-2k_2^3-2k_1^2k_2}{k_2^2-1}, \qquad \forall k_1 \in \mathbb{R}, k_2 < -1.$$

- This cost function comes from an LQR instance

$$A = \begin{bmatrix} -2 & 0 \\ 0 & 1 \end{bmatrix}, \ B = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \ Q = I_2, R = 1$$

- There exists an invertible mapping

$$g(k) := \left(\frac{k_1}{1-k_2}, \frac{2k_2 - k_1^2 - 2k_2^2}{k_2^2-1} \right) \qquad \forall k_1 \in \mathbb{R}, k_2 < -1.$$

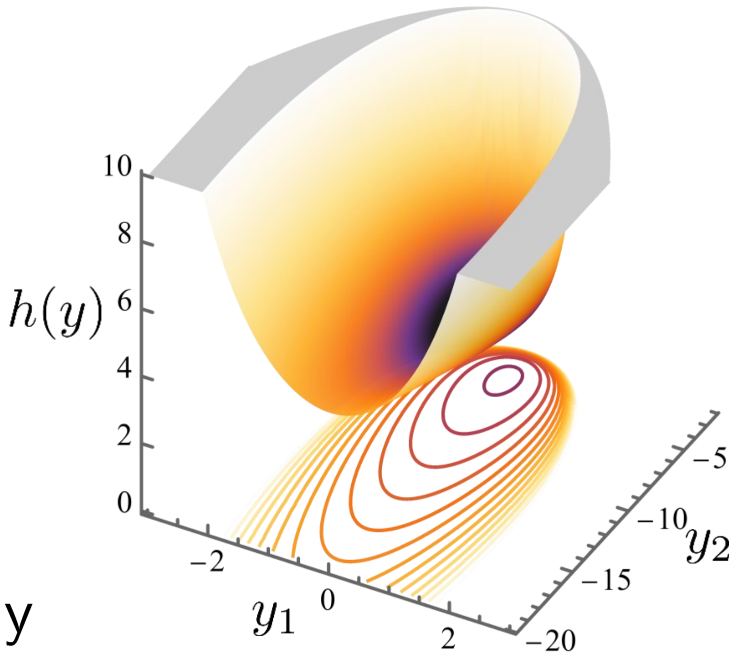

- We get a convex function in terms of the new variable y

$$h(y) := J(g^{-1}(y)) = -y_2 - 1 + y^\mathsf{T} \begin{bmatrix} 1 & y_1 \\ y_1 & -y_2-2 \end{bmatrix}^{-1} y, \qquad \forall \begin{bmatrix} 1 & y_1 \\ y_1 & -y_2-2 \end{bmatrix} \succ 0.$$

Direct convex reformulation

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

    Denote its epigraph as $\mathrm{epi}_{\geq}(f) := \{(\mathsf{K}, \gamma) \in \mathcal{D} \times \mathbb{R} \mid \gamma \geq J(\mathsf{K})\}$.
  - Suppose there exists a **smooth and invertible map** Φ between

    $\mathrm{epi}_{\geq}(J)$ and a **convex set** $\mathcal{F}_{\mathrm{cvx}}$
  - and we further have $(y, \gamma) = \Phi(\mathsf{K}, \gamma), \ \forall (\mathsf{K}, \gamma) \in \mathrm{epi}_{\geq}(J)$

Guarantee 1: Equivalence 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 is globally optimal;

in other words, $0 \in \partial J(\mathsf{K}^*)$ implies global optimality

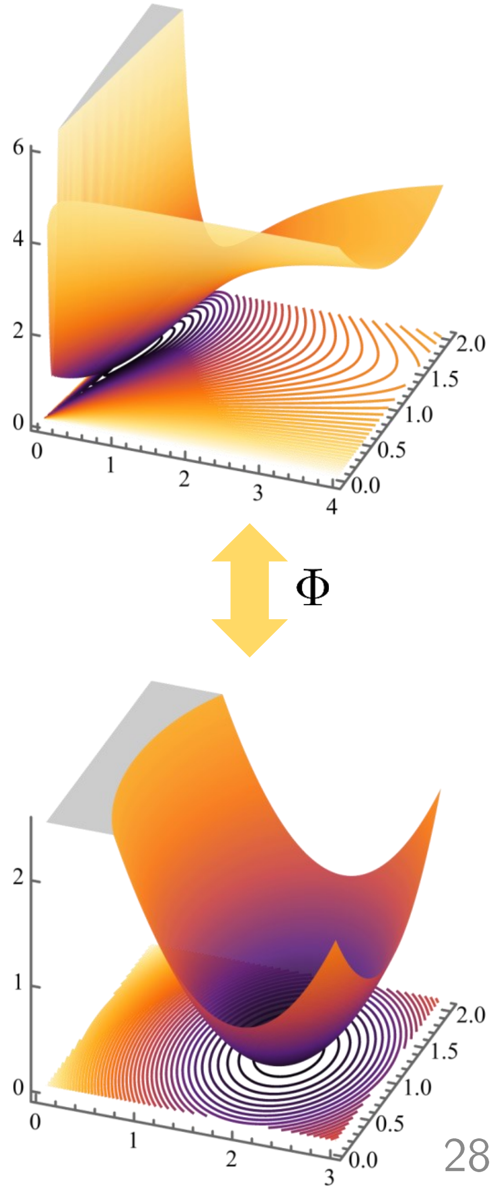

Extended Convex Lifting (ECL)

A schematic illustration of ECL:

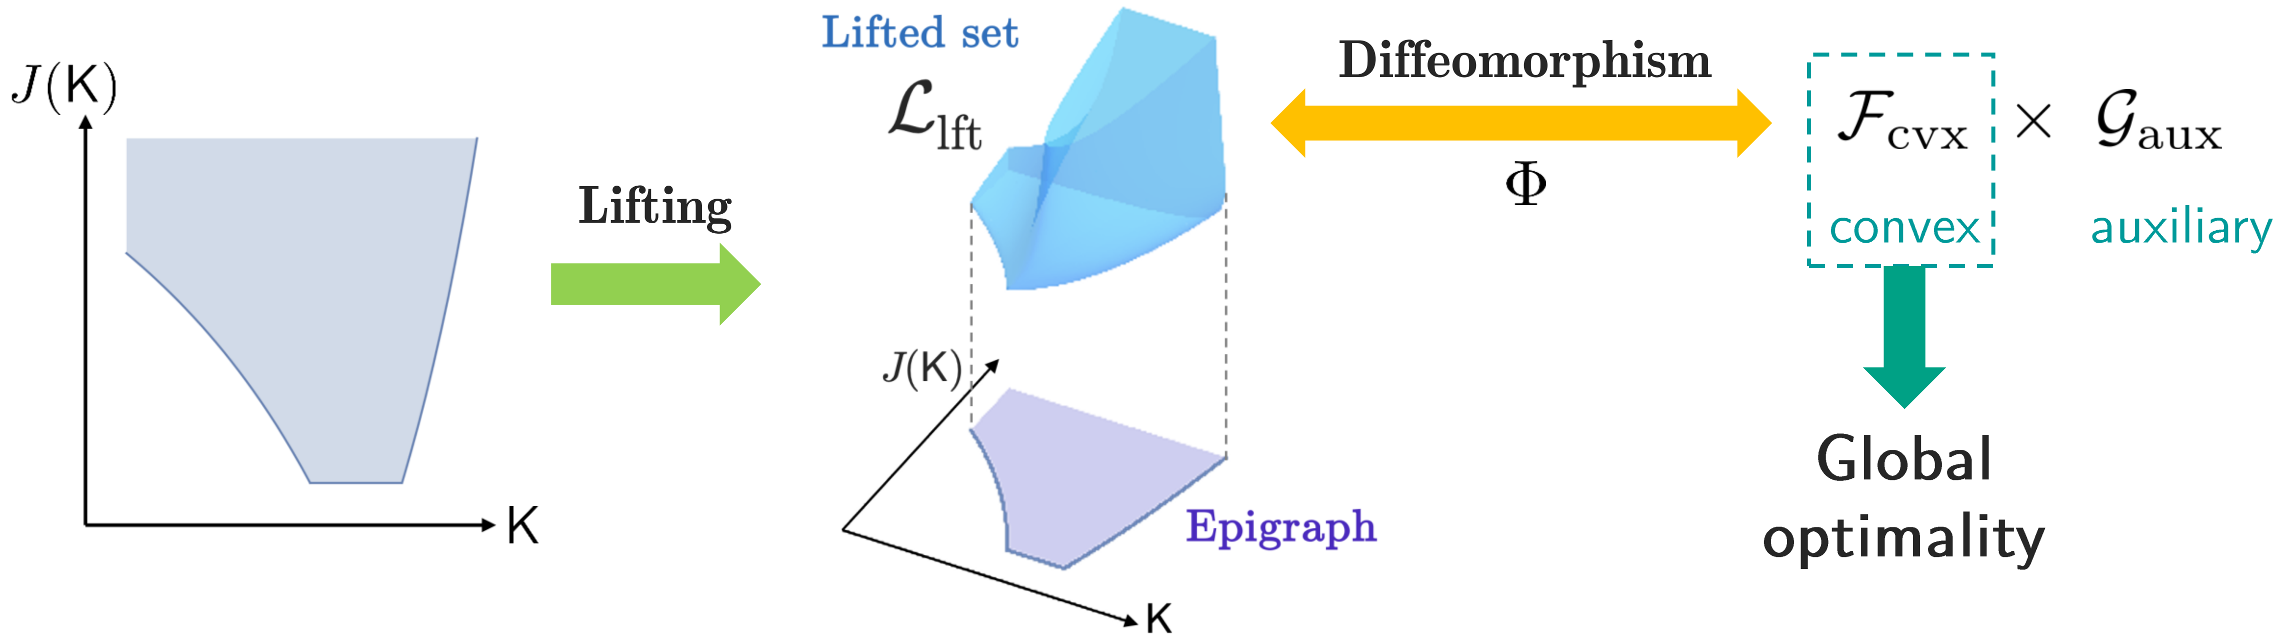

Two key features

Feature 1: a lifting procedure

Feature 2: an auxiliary set

Extended Convex Lifting (ECL)

A schematic illustration of ECL:

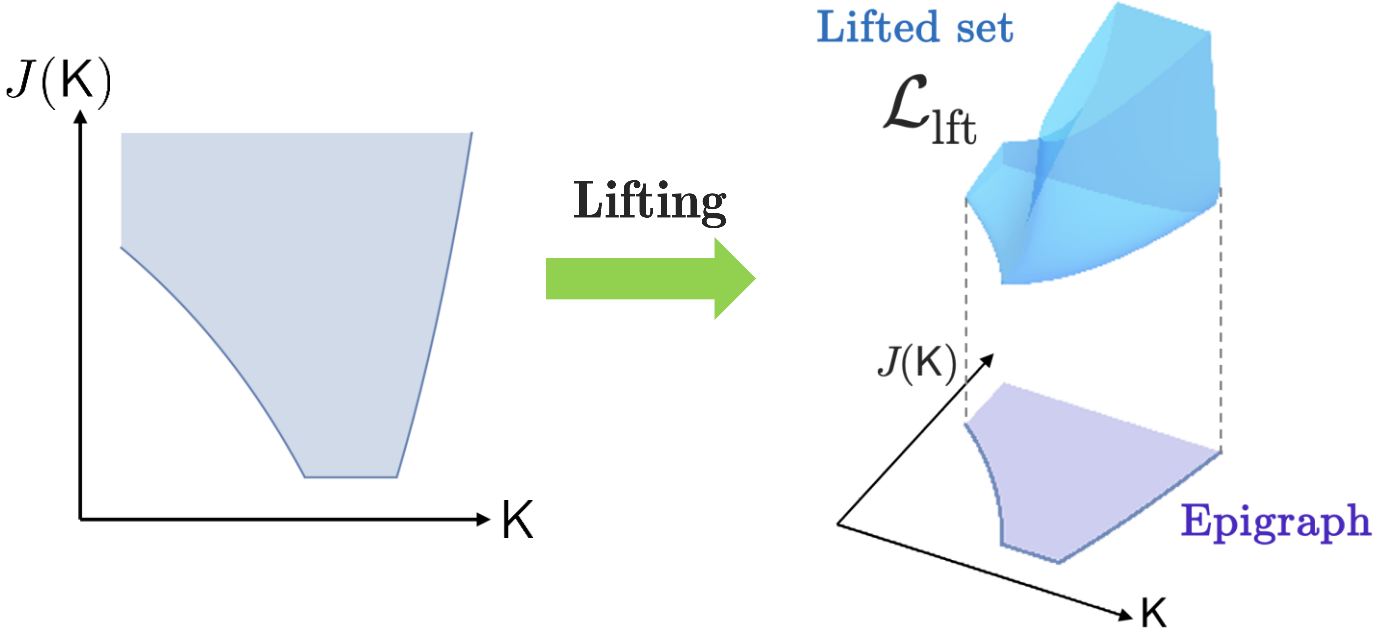

$$\boxed{\begin{array}{c}\text{epi}_{>}(J) \\ (\mathsf{K}, \gamma)\end{array}} \longrightarrow \boxed{\begin{array}{c}\mathcal{L}_{\text{lft}} \\ (\mathsf{K}, \gamma, \xi)\end{array}}$$

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

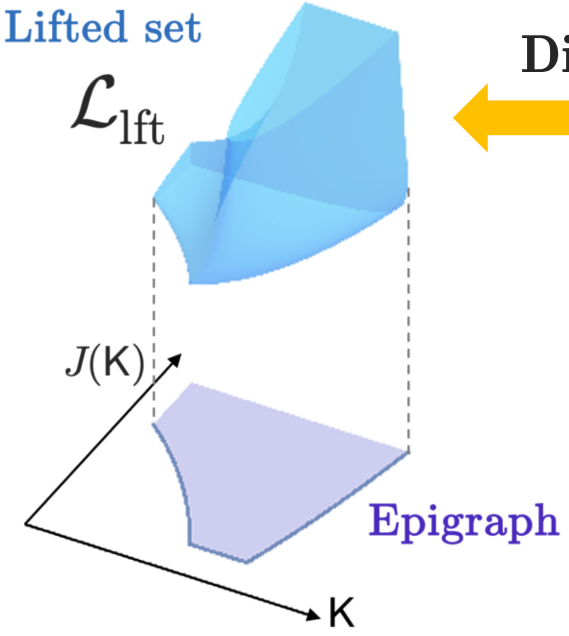

Diffeomorphism

$$\mathcal{L}_{\mathrm{lft}} \overset{\Phi}{\longleftrightarrow} \mathcal{F}_{\mathrm{cvx}} \times \mathcal{G}_{\mathrm{aux}}$$

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

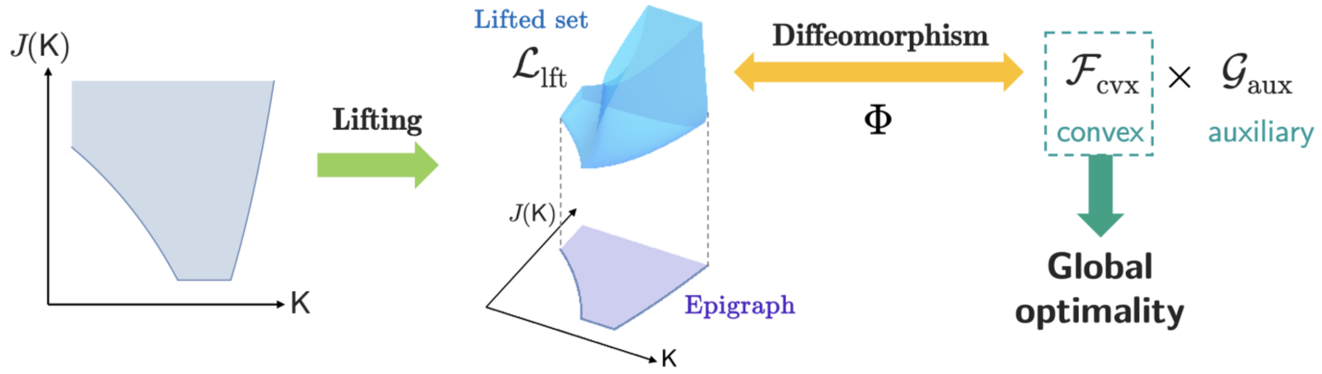

$\mathrm{epi}_{>}(J)$ | (K, γ) → $\mathcal{L}_{\text{lft}}$ | $(\mathsf{K}, \gamma, \xi)$ ↔ $\mathcal{F}_{\text{cvx}} \times \mathcal{G}_{\text{aux}}$ | $(\gamma, \zeta_1, \zeta_2)$

- Consider a continuous function $J : \mathcal{D} \to \mathbb{R}$ 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 $(\mathcal{L}_{\text{lft}}, \mathcal{F}_{\text{cvx}}, \mathcal{G}_{\text{aux}}, \Phi)$ is an ECL of $J : \mathcal{D} \to \mathbb{R}$ if

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

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

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

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

A special ECL

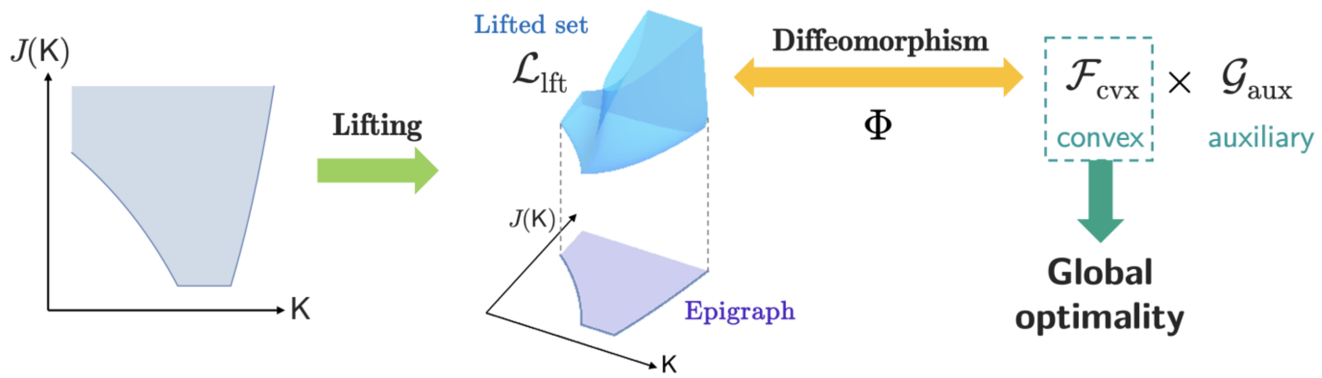

- A lifted set $\mathcal{L}_{\text{lft}}$ satisfying

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

A more intuitive condition

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

$$\pi_{\mathsf{K},\gamma}(\mathcal{L}_{\text{lft}}) = \operatorname{epi}_{\geq}(f).$$

Does this "simpler" lifting condition work?

- The condition on the left is less restrictive
- The simpler condition on the right is only sufficient for state-feedback control (e.g., LQR), but is too restrictive for dynamic output-feedback control (e.g., LQG).
- Some classical LMI formulations are not equivalent to the original control problems.

Non-degenerate points

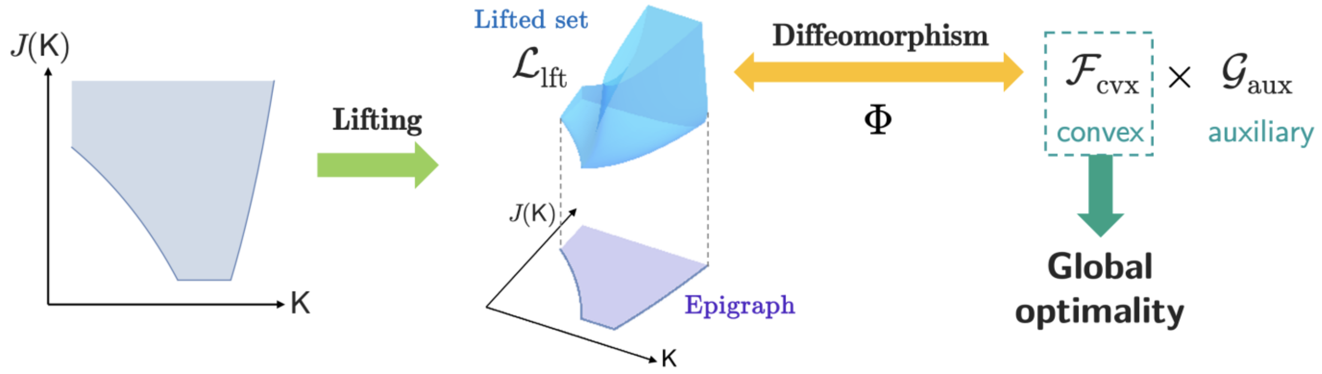

$\mathrm{epi}_>(J)$, (K, γ) → $\mathcal{L}_{\mathrm{lft}}$, $(\mathsf{K}, \gamma, \xi)$ ↔ $\mathcal{F}_{\mathrm{cvx}} \times \mathcal{G}_{\mathrm{aux}}$, $(\gamma, \zeta_1, \zeta_2)$

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** by 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

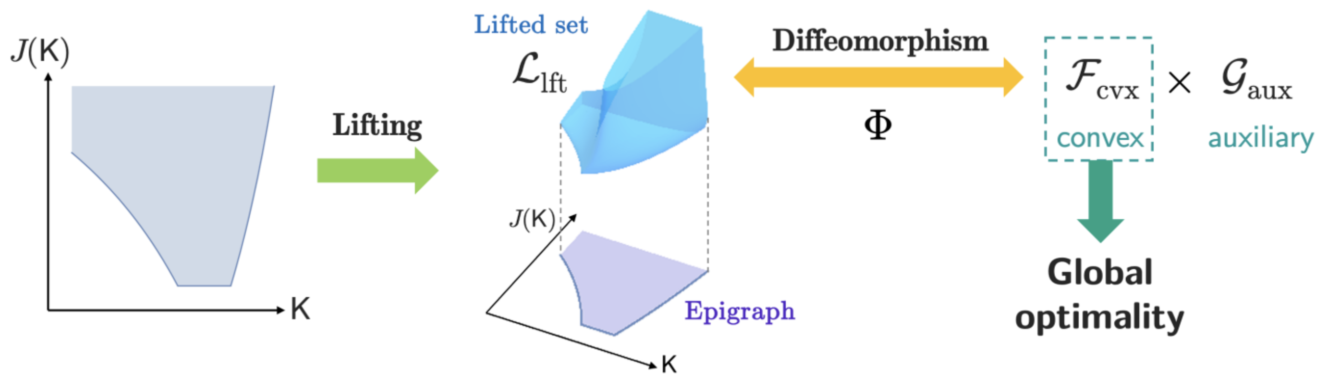

Given an Extended Convex Lifting

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

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

Any non-degenerate Clarke stationary point is 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

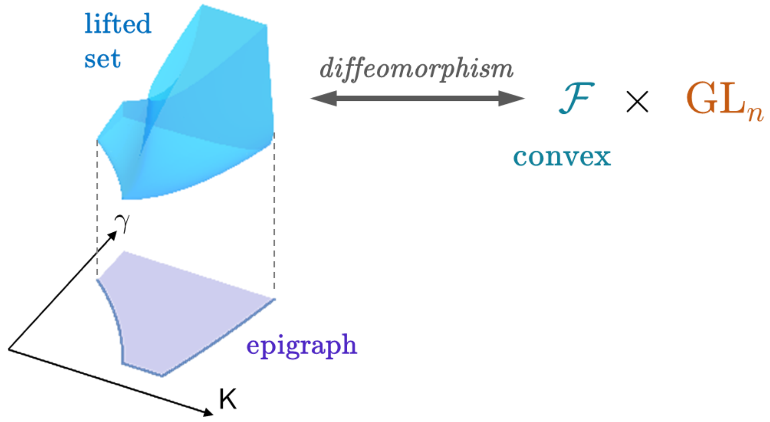

1. The ECL framework
2. **Applications in control**

Linear Quadratic Regulator (LQR)

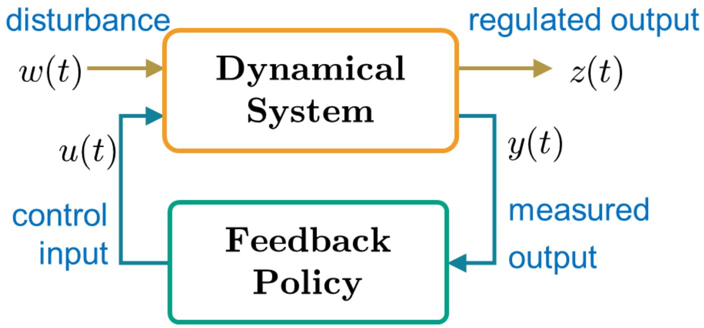

❑ Problem setup

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

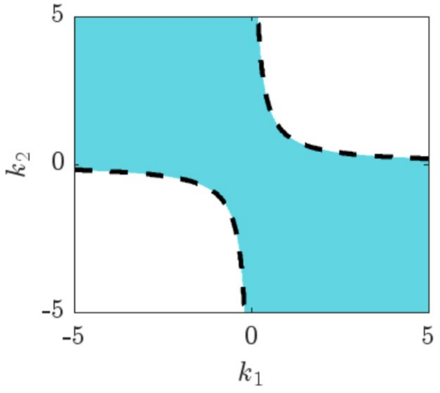

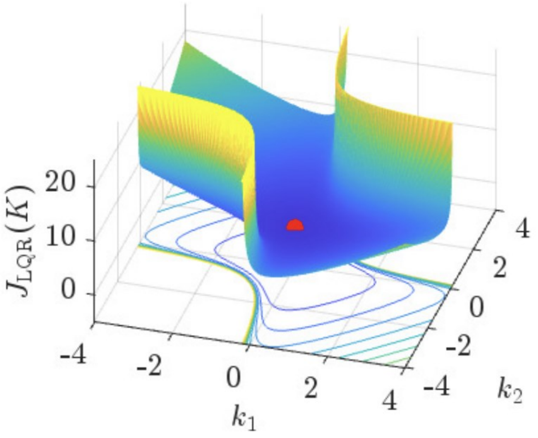

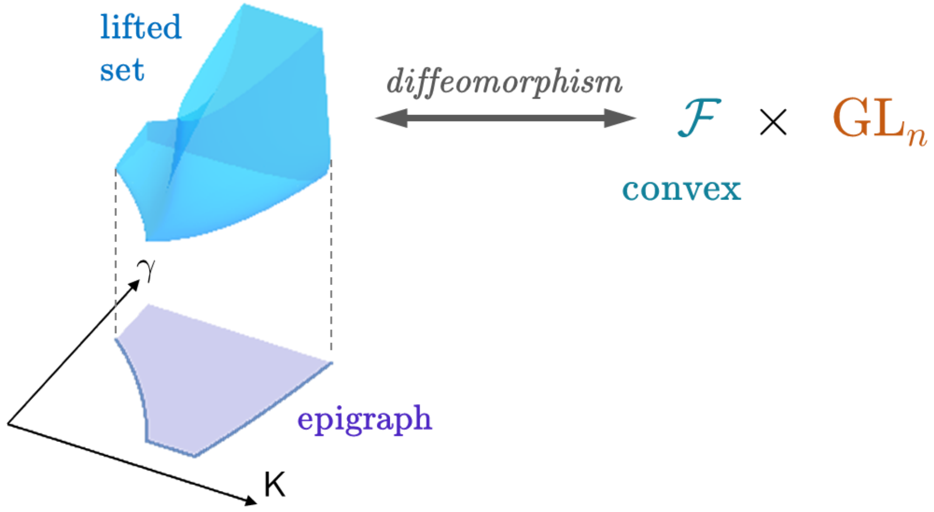

Linear Quadratic Regulator (LQR)

❑ Construction of ECL

Step 1: Lifting (LQR cost by solving a Lyapunov equation [Zhou et al., 1996])

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

Step 2: Convex set (using a change of variables $Y = KX$)

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

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

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

➡ **All policies are non-degenerate**

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

State-feedback robust control

❑ Problem setup

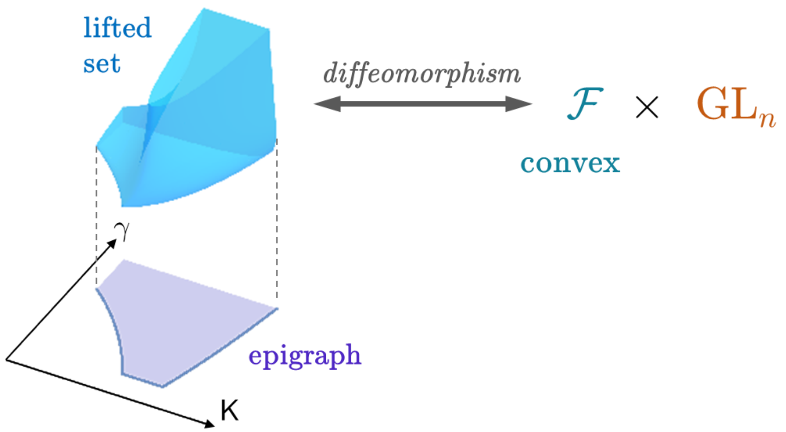

Dynamics: $\dot{x}(t) = Ax(t) + Bu(t) + B_w w(t),$

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 (via the bounded real lemma [Zhou et al., 1996])

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

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

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

➡ All policies are non-degenerate

Corollary: Any Clarke stationary point is 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, whose 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**.

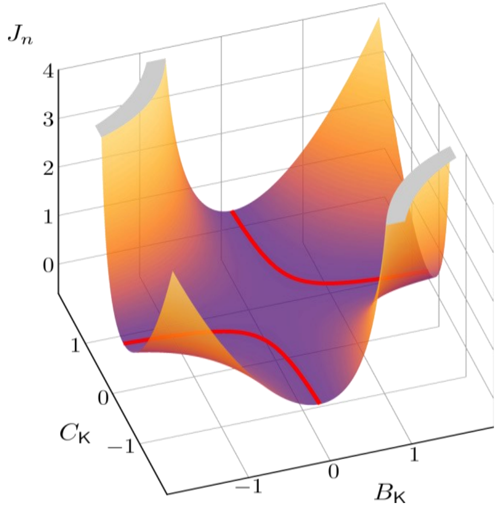

- We also show that **controllable and observable stationary policies are informative**, and are thus globally optimal [ZPT, TAC 2026].
- Similar ECL guarantees hold **for output-feedback robust control** and **other control problems** [ZPT, TAC 2026].

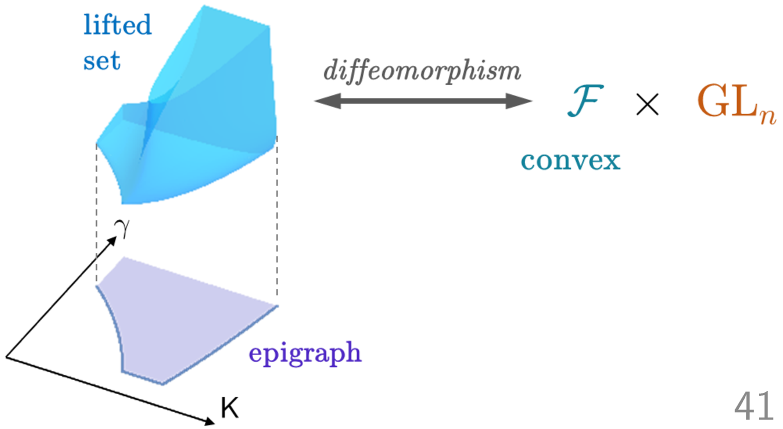

Outline

Geometry:
Benign Nonconvex Landscape

Framework:
Extended Convex Lifting

Algorithm:
Decomposition & Proximal Descent

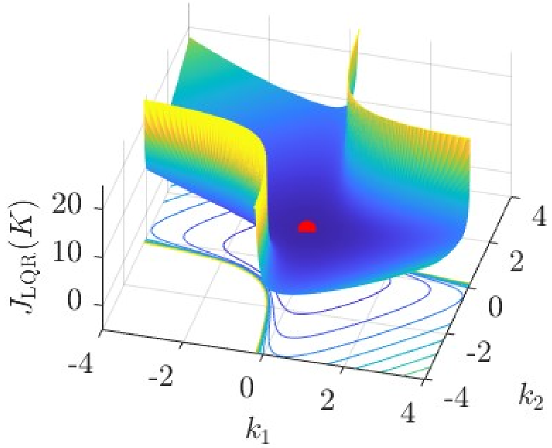

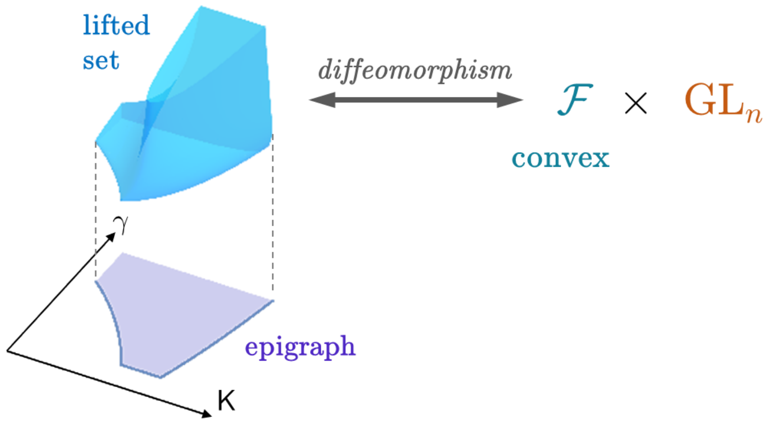

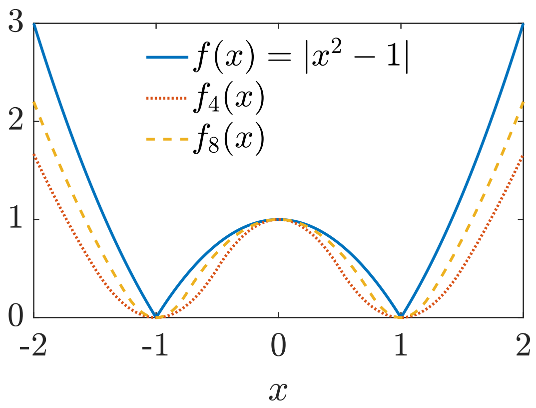

Methods for (non)convex optimization

Convex optimization

- LMIs and SDPs provide explicit stability & robustness guarantees.
- Challenge: large-scale SDP constraints are computationally expensive.
- Decomposition methods exploit problem structure:
  - chordal sparsity [Blair and Peyton, 1993; Vandenberghe and Andersen, 2015];
  - factor-width approximations [Boman et al., 2005; ZSP, TAC 2022];
  - distributed/decomposed computation [Boyd et al., 2011; ZMP, TAC 2018]

Nonconvex policy optimization

- Policy optimization is flexible but often nonconvex and nonsmooth.
- Challenge:
  - Basic subgradient methods cannot ensure descent;
  - Iterates may leave the initial sublevel set.
- Proximal descent provides a unified first-order viewpoint for solving a class of nonconvex and nonsmooth problems [LZ, 2025].

Chordal decomposition for scalable SDPs

❑ Chordal decomposition

- LMIs in control often contain **large sparse positive semidefinite (PSD) constraints**.
- Chordal sparsity decomposes one large PSD cone into **smaller PSD cones**.
- This reduces **computational cost** and enables **distributed/decomposed** computation.

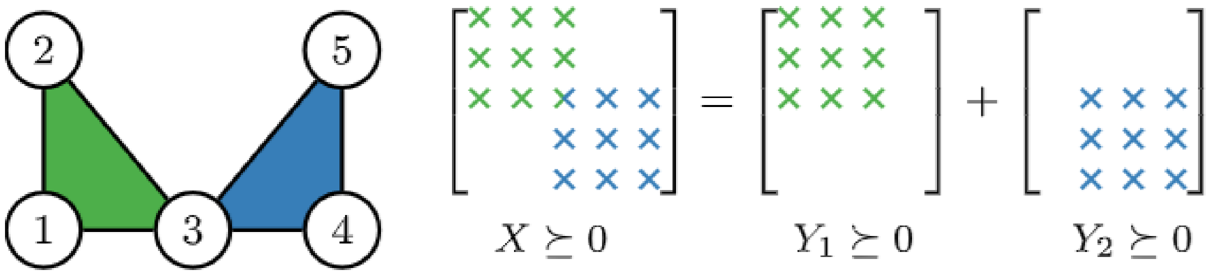

❑ **Semidefinite optimization**: [Fukuda et al., 2001]; [Vandenberghe & Andersen, 2015], [ZFPGW, *Math Prog.* 2020] [Zhang, 2025], etc.

❑ **Control applications**: [Andersen, Pakazad, Hansson & Rantzer, 2014]; [Mason & Papachristodoulou, 2014]; [ZMP, TAC 2018], etc.

CDCS (Cone Decomposition Conic Solver)

https://github.com/oxfordcontrol/CDCS

	Time (s)		
	sedumi	SCS	CDCS
Stability	115.0	108.9	40.6
$\mathcal{H}_2$	805.0	556.1	147.4
$\mathcal{H}_\infty$	3 374.8	2 130.2	172.0

Computational savings from chordal decomposition

Decomposition methods for scalable optimization

❑ Chordal and factor-width decomposition

- **Chordal sparsity**: exact decomposition for sparse PSD constraints;
- **Factor-width approximations**: cheaper conic approximations for scalability.

Promising tools for scalable semidefinite and polynomial optimization

Contents lists available at ScienceDirect

Annual Reviews in Control

journal homepage: www.elsevier.com/locate/arcontrol

ELSEVIER

Vision article

Chordal and factor-width decompositions for scalable semidefinite and polynomial optimization

Check for updates

Yang Zheng [a,*], Giovanni Fantuzzi [b], Antonis Papachristodoulou [c]

[a] *Department of Electrical and Computer Engineering, University of California, San Diego, CA 92093, United States of America*
[b] *Department of Aeronautics, Imperial College London, London, SW7 2AZ, UK*
[c] *Department of Engineering Science, University of Oxford, Parks Road, Oxford OX1 3PJ, UK*

First-order algorithms for nonconvex optimization

$$\min_x f(x)$$

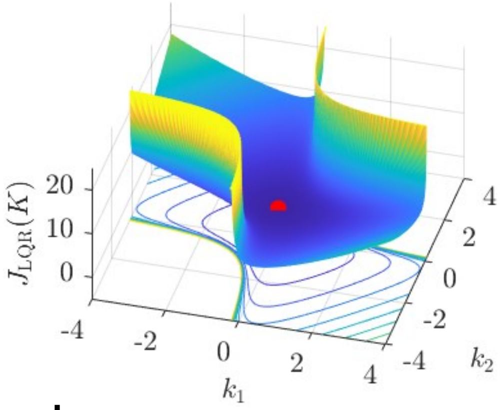

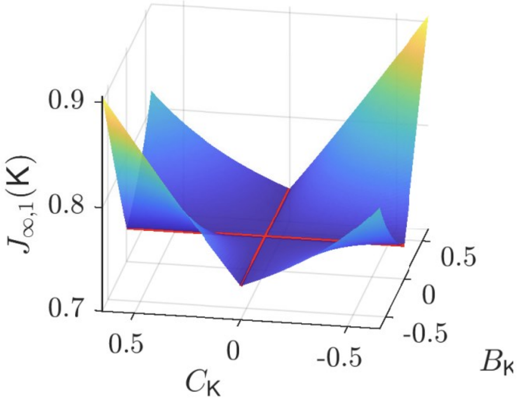

Standard algorithms:

- 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

- Their variants (with acceleration) achieve the best first-order complexity for convex functions [Chaps 2 & 3, Nesterov, 2018];

Proximal descent method

❑ A conceptual descent algorithm

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

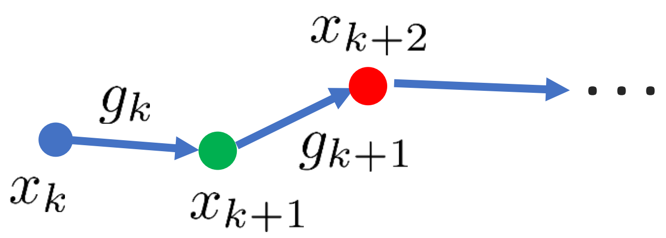

Suppose it satisfies a cost value drop

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

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 is such a cost value drop achievable? And How?

Weakly convex functions

Proximal bundle procedure

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

Weakly convex function

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

- Any convex (nonsmooth) function
- Any L-smooth (nonconvex) function
- **LQR** and **state-feedback robust control** under mild assumptions [WLZ, ACC 2026]

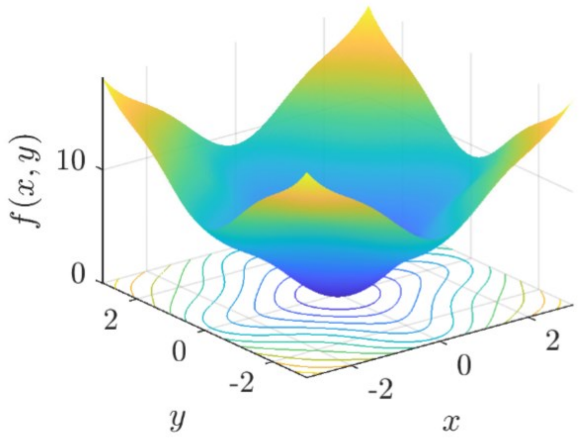

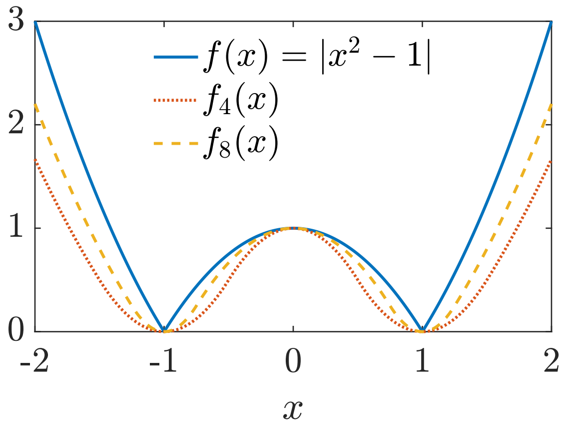

Algorithm 1 Proximal descent method

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

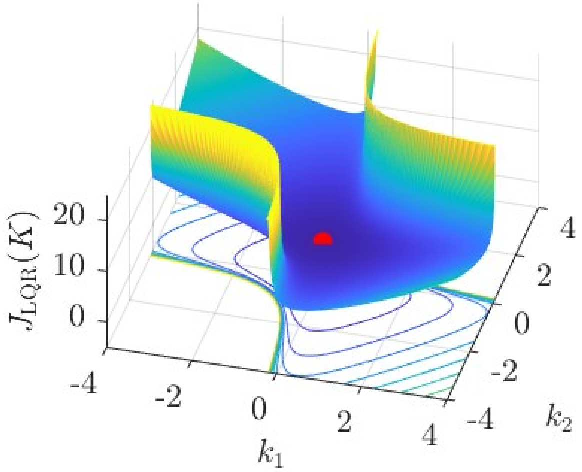

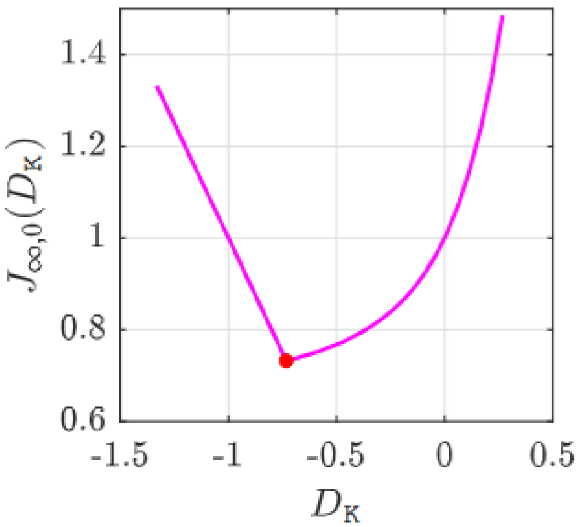

Convergence guarantees

$$\min_x f(x)$$

Liao & Zheng, 2025

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

- 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

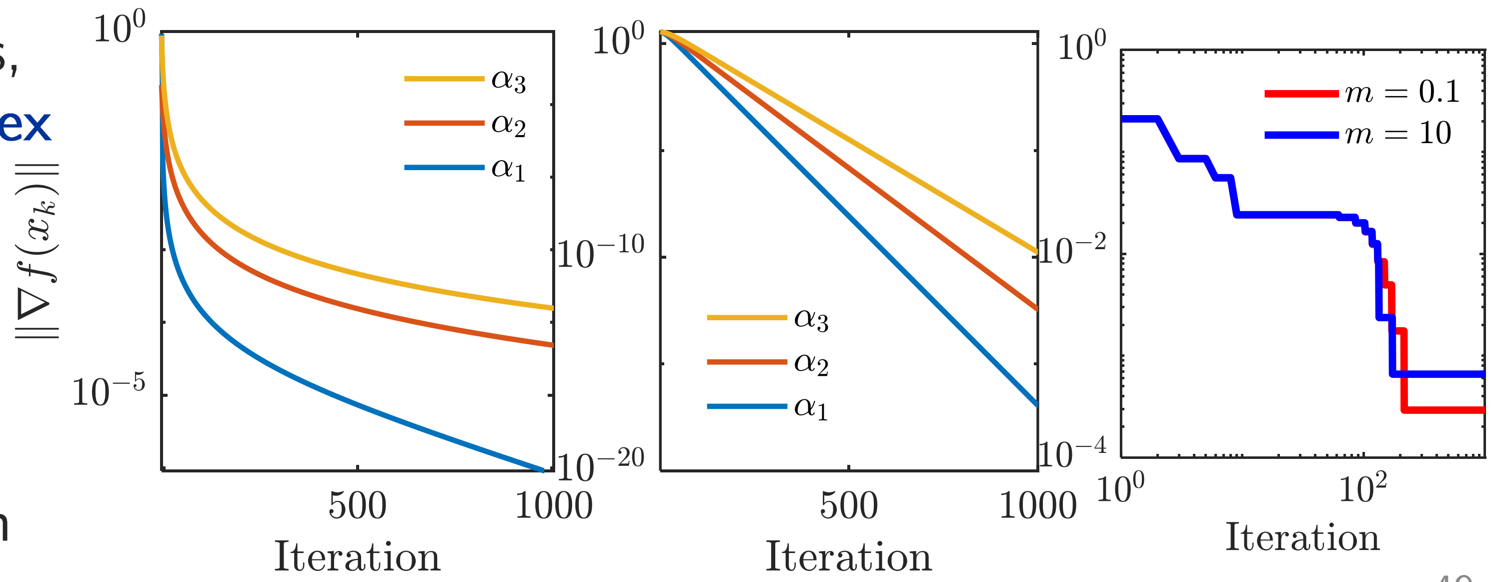

Smooth non-convex examples **Non-smooth** example

Conclusion

Summary

Geometry: Benign Nonconvex Landscape

Framework: Extended Convex Lifting

Algorithm: Decomposition & Proximal Descent

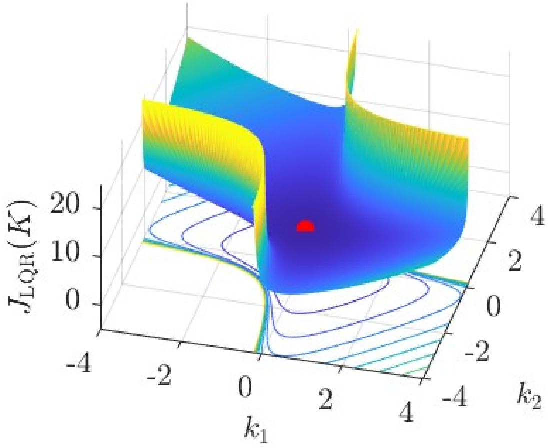

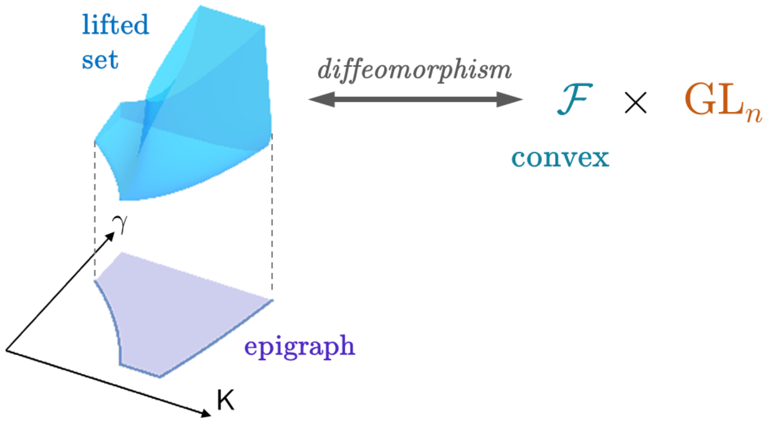

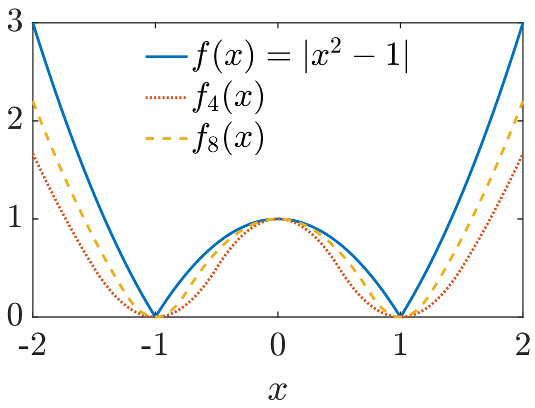

Hidden convexity explains **benign nonconvexity in control** and guides the design of **scalable algorithms**.

More opportunities

- Deal with degenerate points in policy search? Avoid **saddle points** with guarantees?
- Extend the analysis to **distributed control**?
- Establish guarantees **beyond linear control**?

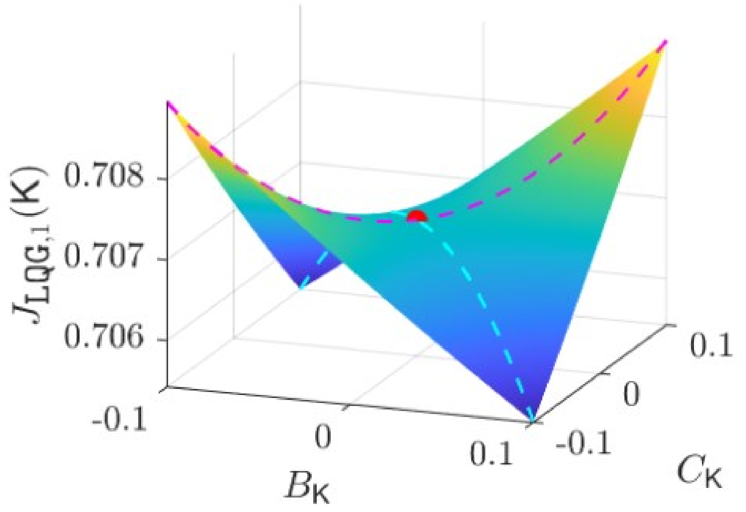

Convex optimization: Reliability and **Global guarantees**

+

Nonconvex policy optimization: Flexibility and scalability

Toward reliable, robust, and learning-enabled control for modern dynamical systems

Respect the unstable[3,4]

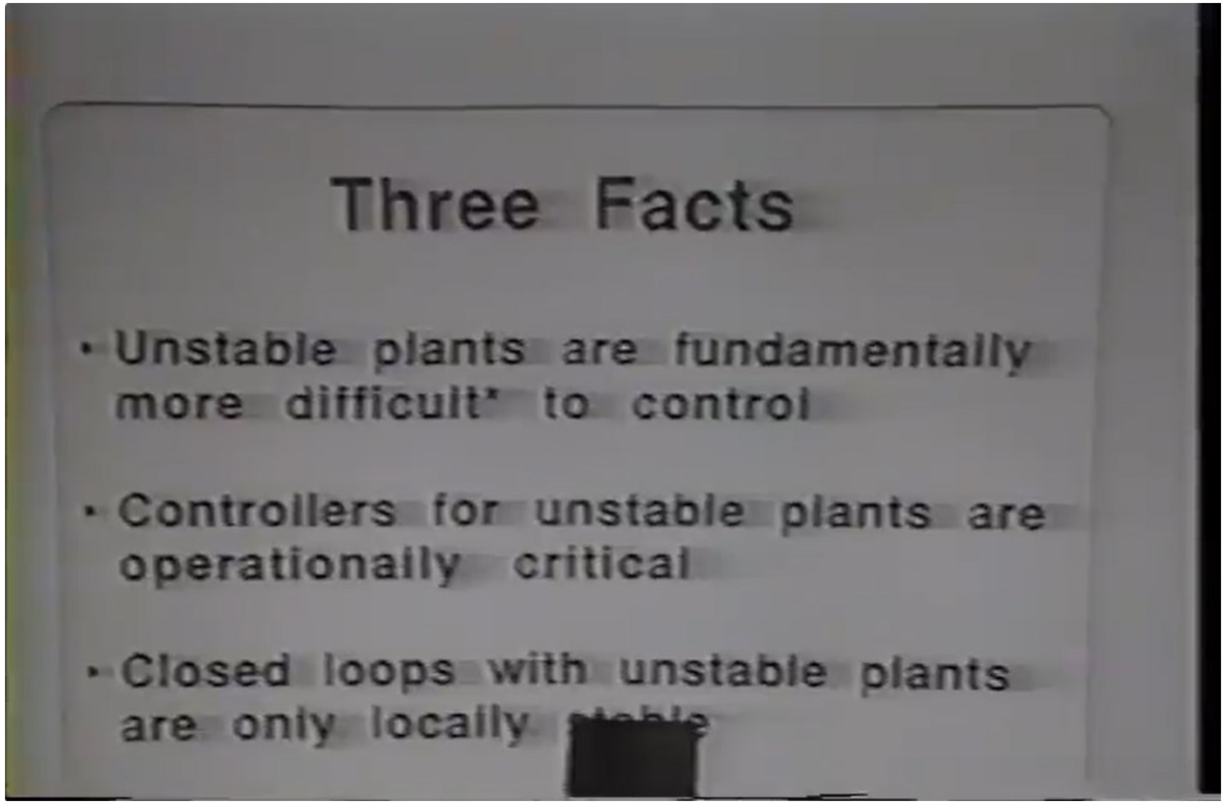

> The practical, physical (and sometimes dangerous) consequences of control must be respected, and **the underlying principles** must be clearly and well taught.
>
> Gunter Stein

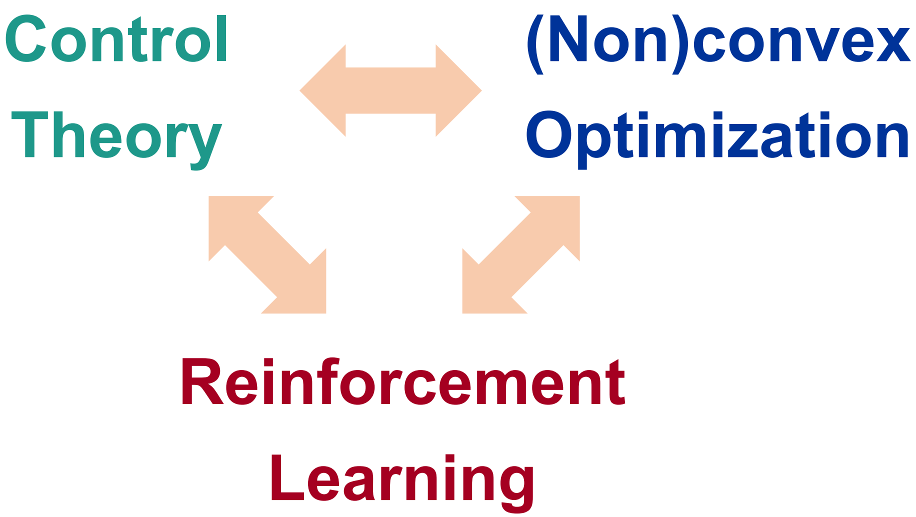

[3] Gunter Stein. "Respect the unstable." *IEEE Control Systems Magazine* 23.4 (2003): 12-25.

[4] "Respect the unstable.", by Gunter Stein, the first Bode Lecture at CDC (1989)

Acknowledgements


Antonis Papachristodoulou


Na Li


Alessio Lomuscio


Behrouz Touri


Jorge Cortes


Mehran Mesbahi


Bahman Gharesifard


Peter Seiler


Feng-Yi Liao

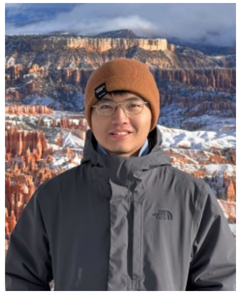
Chih-Fan Pai

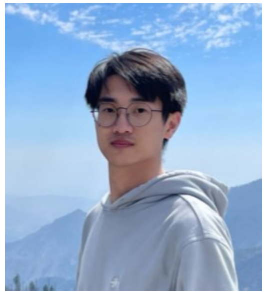
Xu Shang


Yuto Watanabe


Jiachen Qian

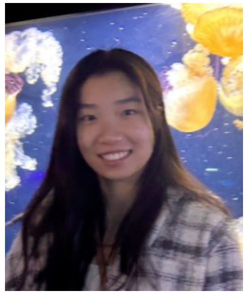
Nicole Li


Thomas Madden


Tianhao Wu


Yujie Tang


Giovanni Fantuzzi


Luca Furieri


Maryam Kamgarpour


Mario Sznaier


Bin Hu

Many others

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