

Benign Nonconvex Landscapes in Optimal and Robust Control

Yang Zheng

Assistant Professor,
ECE Department, UC San Diego


Seminar at Systems Control Group, University of Toronto

March 25, 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


Data-driven Decision Making

- ❑ Data-driven decision-making for complex tasks in dynamical systems, e.g., game playing, robotic manipulation/ locomotion, networked systems, ChatGPT, etc.
- ❑ **Reinforcement learning (RL)** has served as one backbone of the recent successes of data-driven decision-making.
- ❑ **Policy optimization** as one of the major workhorses of modern RL.



Today's topic: Policy optimization

Feedback control



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


Policy Optimization


Policy parametrization

$$\begin{aligned} \min_{\mathbf{K}} \quad & J(\mathbf{K}) \\ \text{s.t.} \quad & \mathbf{K} \in \mathcal{C} \end{aligned}$$


- 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



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

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



Opportunities

- Easy-to-implement
- **Scalable** to high-dimensional problems
- Enable **model-free search** with rich observations (e.g. images)

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] [Fatkullin & 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**!

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

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

$u_t = Kx_t \longrightarrow$ 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 is nonconvex in K

$$\begin{aligned} \min_K \quad & J(K) \\ \text{s.t.} \quad & K \in \mathcal{C} \end{aligned}$$



Non-convexity in control

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

- **Stabilization** is already a **non-convex** constraint!

✓ A policy K is said to **stabilize** the system, if the system trajectory converges to the origin.

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

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


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

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

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

$$\min_K J(K)$$

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



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


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

Non-convexity in control


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

$$\begin{aligned} \min_{\mathbf{K}} \quad & J(\mathbf{K}) \\ \text{s.t.} \quad & \mathbf{K} \in \mathcal{C} \end{aligned}$$

- 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!**

- 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: $J_{\text{LQR}}(K) := \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \int_0^T x^\top(t) Q x(t) + u^\top(t) R u(t) dt \right]$

$$\begin{aligned} \min_K \quad & J(K) \\ \text{s.t.} \quad & K \in \mathcal{C} \end{aligned}$$



Benign non-convexity in LQR

□ Linear Quadratic Regulator (LQR)

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**


$$\{K \in \mathbb{R}^{m \times n} \mid A + BK \text{ is stable}\}$$

$$\iff \{K \in \mathbb{R}^{m \times n} \mid \exists P \succ 0, P(A + BK)^\top + (A + BK)P \prec 0\}$$

$$\iff \{K \in \mathbb{R}^{m \times n} \mid \exists P \succ 0, PA^\top + L^\top B^\top + AP + BL \prec 0, L = KP\}$$

$$\iff \{K = LP^{-1} \in \mathbb{R}^{m \times n} \mid \exists P \succ 0, PA^\top + L^\top B^\top + AP + BL \prec 0\}.$$

$$\begin{aligned} \min_K \quad & J(K) \\ \text{s.t.} \quad & K \in \mathcal{C} \end{aligned}$$



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

Benign non-convexity in LQR

□ Linear Quadratic Regulator (LQR)

$$J_{\text{LQR}}(K) := \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \int_0^T x^\top(t) Q x(t) + u^\top(t) R u(t) dt \right]$$

$$\begin{aligned} \min_K \quad & J(K) \\ \text{s.t.} \quad & K \in \mathcal{C} \end{aligned}$$

Fact 2: the LQR cost is real analytical and thus **infinitely differentiable** over $\mathcal{C}$.

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

✓ Another form $J_{\text{LQR}}(K) = \text{Tr} [(Q + K^\top R K) X]$


Lyapunov equation $(A + BK)X + X(A + BK)^\top + W = 0$

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)


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

$$\begin{aligned} \min_K \quad & J(K) \\ \text{s.t.} \quad & K \in \mathcal{C} \end{aligned}$$

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

- 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 H_{∞} 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: $J_{\infty}(K) := \sup_{\|w(t)\|_2 \leq 1} \int_0^{\infty} x^{\top}(t)Qx(t) + u^{\top}(t)Ru(t) dt$

Benign non-convexity in robust control



□ State-feedback Hinf robust control

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

$$\begin{aligned} \min_K \quad & J(K) \\ \text{s.t.} \quad & K \in \mathcal{C} \end{aligned}$$

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)

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

Benign non-convexity in LQG

□ Linear Quadratic Gaussian (LQG) Control



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

$$y(t) = Cx(t) + D_v v(t)$$

Performance:

$$J_{\text{LQG}} := \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \int_0^T x^\top(t) Q x(t) + u^\top(t) R u(t) dt \right]$$

➤ We consider **dynamic policies**:

$$\dot{\xi}(t) = A_K \xi(t) + B_K y(t)$$

$$u(t) = C_K \xi(t)$$

$$K = (A_K, B_K, C_K) \in \mathbb{R}^{n \times n} \times \mathbb{R}^{n \times p} \times \mathbb{R}^{m \times n}$$

Policy optimization for LQG

$$\min_K J(K)$$

$$\text{s.t. } K = (A_K, B_K, C_K) \in \mathcal{C}_{\text{full}}$$


Benign nonconvexity in LQG

Linear Quadratic Gaussian (LQG) Control

$$\mathcal{C}_{\text{full}} = \left\{ K \mid K = (A_K, B_K, C_K) \text{ is full order} \right. \\ \left. \begin{bmatrix} A & BC_K \\ B_K C & A_K \end{bmatrix} \text{ is Hurwitz stable} \right\}$$


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**



Policy optimization for LQG

$$\min_K J(K) \\ \text{s.t. } K = (A_K, B_K, C_K) \in \mathcal{C}_{\text{full}}$$



Two connected components

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

Benign nonconvexity in LQG

□ Linear Quadratic Gaussian (LQG) Control

Fact 2: For a **controllable and observable** policy, if it is a **stationary point** $\nabla J(K) = 0$, 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 Foundation for Learning Control Policies

Bin Hu¹, Kaiqing Zhang^{2,3}, Na Li⁴, Mehran Mesbahi⁵, Maryam Fazel

Policy Optimization in Control: Geometry and Algorithmic Implications

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


^aHarvard University, School of Engineering and Applied Sciences,
150 Western Ave, Boston, 02134, MA, US

^bUniversity of California San Diego, Department of Electrical and Computer Engineering,
9500 Gilman Drive, La Jolla, 92093, CA, US


^cUniversity 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



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** 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

$$f_1(x_1, x_2) = \left(\frac{x_2}{x_1} - 2 \right)^2 + (x_2 - 1)^2, \quad \text{dom}(f_1) = \{x \in \mathbb{R}^2 \mid x_1 > 0, x_2 > 0\}.$$





Its global **minimizer** is

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

Define an **invertible map**

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



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

Example 2

□ Nonconvex and non-smooth function

$$f_2(x_1, x_2) = \left| \frac{x_2}{x_1} - 2 \right| + |x_2 - 1|, \quad \text{dom}(f_2) = \{x \in \mathbb{R}^2 \mid x_1 > 0, x_2 > 0\}.$$

$\xrightarrow{\quad y_1 \quad}$ $\xrightarrow{\quad y_2 \quad}$

Its global **minimizer** is



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

Define an **invertible map**

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



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



Direct Convex Reformulation


□ Direct convex reformulation (no lifting)

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

Guarantee 1: Minimizing J is equivalent to a convex problem


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

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



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



Formal ECL Definition



- Consider a continuous function $J(K) : \mathcal{D} \rightarrow \mathbb{R}$ where $\mathcal{D} \subseteq \mathbb{R}^d$.
- Denote its strict and non-strict epigraph as $\text{epi}_{>}(J) := \{(K, \gamma) \in \mathcal{D} \times \mathbb{R} \mid \gamma > J(K)\}$, $\text{epi}_{\geq}(J) := \{(K, \gamma) \in \mathcal{D} \times \mathbb{R} \mid \gamma \geq J(K)\}$.


Extended Convex Lifting (ECL)

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

- A lifted set $\mathcal{L}_{\text{lift}}$ satisfying $\text{epi}_{>}(J) \subseteq \pi_{K, \gamma}(\mathcal{L}_{\text{lift}}) \subseteq \text{cl } \text{epi}_{\geq}(J)$
- A diffeomorphism $\Phi : \mathcal{L}_{\text{lift}} \rightarrow \mathcal{F}_{\text{cvx}} \times \mathcal{G}_{\text{aux}}$ such that

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

A special ECL



A more intuitive condition

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





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

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



Extended Convex Lifting:

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

$$\text{epi}_{>}(J) \subseteq \pi_{K,\gamma}(\mathcal{L}_{\text{lift}}) \subseteq \text{cl epi}_{\geq}(J)$$
- A diffeomorphism $\Phi : \mathcal{L}_{\text{lift}} \rightarrow \mathcal{F}_{\text{cvx}} \times \mathcal{G}_{\text{aux}}$ such that

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


- By construction, some points in $\text{epi}_{\geq}(J)$ 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 $(K, J(K)) \in \pi_{K,\gamma}(\mathcal{L}_{\text{lift}})$

➡ **well-behaved**

ECL Guarantees



Guarantee 1: Convex Reformulation

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


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

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:

$$J_{\text{LQR}}(K) := \lim_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \int_0^T x^\top(t) Q x(t) + u^\top(t) R u(t) dt \right]$$



❑ Nonconvex and smooth landscape



Linear Quadratic Regulator (LQR)

□ Construction of ECL

Step 1: Lifting

$$\mathcal{L}_{\text{LQR}} := \left\{ (K, \gamma, \mathbf{X}) : \mathbf{X} \succ 0, (A + BK)\mathbf{X} + \mathbf{X}(A + BK)^\top + W = 0, \gamma \geq \text{Tr} \left[(Q + K^\top R K) \mathbf{X} \right] \right\}.$$

Step 2: Convex set

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

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

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

 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: $J_\infty(K) := \sup_{\|w(t)\|_2 \leq 1} \int_0^\infty x^\top(t) Q x(t) + u^\top(t) R u(t) dt$



□ Building an ECL

Step 1: Lifting

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

State-feedback Robust Control

□ Building an ECL

Step 2: Convex set

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

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

- No auxiliary set
 - Lifted set satisfies $\pi_{K,\gamma}(\mathcal{L}_\infty) = \text{epi}_{\geq}(J_\infty)$
- ➡ 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 $\lim_{t \rightarrow \infty} \mathbb{E}[x(t)\xi(t)^T]$ 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, 2026].




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

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

- **Subgradient method** for non-smooth convex functions

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

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




Very careful step
size tuning

Non-trivial to extend for
nonsmooth/nonconvex case

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

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

- Valid for any positive step size

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

Weakly convex functions

Proximal bundle procedure

Proximal Descent Method





□ Weakly convex function

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

- 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 $k = 1, 2, \dots, T$ **do**
 $x_{k+1} = \text{ProxDescent}(x_k)$;
 end for




Convergence guarantees

$$\min_x f(x)$$

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

- ❑ 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






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**?



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