
Stabilizing Dynamical Systems via Policy Gradient Methods

Juan C. Perdomo*
University of California, Berkeley

Jack Umenberger
MIT

Max Simchowitz
MIT

Abstract

Stabilizing an unknown control system is one of the most fundamental problems in control systems engineering. In this paper, we provide a simple, model-free algorithm for stabilizing fully observed dynamical systems. While model-free methods have become increasingly popular in practice due to their simplicity and flexibility, stabilization via direct policy search has received surprisingly little attention. Our algorithm proceeds by solving a series of discounted LQR problems, where the discount factor is gradually increased. We prove that this method efficiently recovers a stabilizing controller for linear systems, and for smooth, nonlinear systems within a neighborhood of their equilibria. Our approach overcomes a significant limitation of prior work, namely the need for a pre-given stabilizing control policy. We empirically evaluate the effectiveness of our approach on common control benchmarks.

1 Introduction

Stabilizing an unknown control system is one of the most fundamental problems in control systems engineering. A wide variety of tasks - from maintaining a dynamical system around a desired equilibrium point, to tracking a reference signal (e.g a pilot's input to a plane) - can be recast in terms of stability. More generally, synthesizing an initial stabilizing controller is often a necessary first step towards solving more complex tasks, such as adaptive or robust control design [Sontag, 1999, 2009].

In this work, we consider the problem of finding a stabilizing controller for an unknown dynamical system via direct policy search methods. We introduce a simple procedure based off policy gradients which provably stabilizes a dynamical system around an equilibrium point. Our algorithm only requires access to a simulator which can return rollouts of the system under different control policies, and can efficiently stabilize both linear and smooth, nonlinear systems.

Relative to model-based approaches, model-free procedures, such as policy gradients, have two key advantages: they are conceptually simple to implement, and they are easily adaptable; that is, the same method can be applied in a wide variety of domains without much regard to the intricacies of the underlying dynamics. Due to their simplicity and flexibility, direct policy search methods have become increasingly popular amongst practitioners, especially in settings with complex, nonlinear dynamics which may be challenging to model. In particular, they have served as the main workhorse for recent breakthroughs in reinforcement learning and control [Silver et al., 2016, Mnih et al., 2015, Andrychowicz et al., 2020].

Despite their popularity amongst practitioners, model-free approaches for continuous control have only recently started to receive attention from the theory community [Fazel et al., 2018, Kakade et al., 2020, Tu and Recht, 2019]. While these analyses have begun to map out the computational and statistical tradeoffs that emerge in choosing between model-based and model-free approaches,

*Correspondence to jcperdomo@berkeley.edu

they all share a common assumption: that the unknown dynamical system in question is stable, or that an initial stabilizing controller is known. As such, they do not address the perhaps more basic question, *how do we arrive at a stabilizing controller in the first place?*

1.1 Contributions

We establish a reduction from stabilizing an unknown dynamical system to solving a series of discounted, infinite-horizon LQR problems via policy gradients, for which no knowledge of an initial stable controller is needed. Our approach, which we call *discount annealing*, gradually increases the discount factor and yields a control policy which is near optimal for the undiscounted LQR objective. To the best of our knowledge, our algorithm is the first model-free procedure shown to provably stabilize unknown dynamical systems, thereby solving an open problem from Fazel et al. [2018].

We begin by studying linear, time-invariant dynamical systems with full state observation and assume access to *inexact* cost and gradient evaluations of the discounted, infinite-horizon LQR cost of a state-feedback controller K . Previous analyses (e.g., [Fazel et al., 2018]) establish how such evaluations can be implemented with access to (finitely many, finite horizon) trajectories sampled from a simulator. We show that our method recovers the controller K_* which is the optimal solution of the *undiscounted* LQR problem in a bounded number of iterations, up to optimization and simulator error. The stability of the resulting K_* is guaranteed by known stability margin results for LQR. In short, we prove the following guarantee:

Theorem 1 (informal). *For linear systems, discount annealing returns a stabilizing state-feedback controller which is also near-optimal for the LQR problem. It uses at most polynomially many, ε -inexact gradient and cost evaluations, where the tolerance ε also depends polynomially on the relevant problem parameters.*

Since both the number of queries and error tolerance are polynomial, discount annealing can be efficiently implemented using at most polynomially many samples from a simulator.

Furthermore, our results extend to smooth, *nonlinear* dynamical systems. Given access to a simulator that can return damped system rollouts, we show that our algorithm finds a controller that attains near-optimal LQR cost for the Jacobian linearization of the nonlinear dynamics at the equilibrium. We then show that this controller stabilizes the nonlinear system within a neighborhood of its equilibrium.

Theorem 2 (informal). *Discount annealing returns a state-feedback controller which is exponentially stabilizing for smooth, nonlinear systems within a neighborhood of their equilibrium, using again only polynomially many samples drawn from a simulator.*

In each case, the algorithm returns a near optimal solution K to the relevant LQR problem (or local approximation thereof). Hence, the stability properties of K are, in theory, no better than those of the optimal LQR controller K_* . Importantly, the latter may have worse stability guarantees than the optimal solution of a corresponding robust control objective (e.g. \mathcal{H}_∞ synthesis). Nevertheless, we focus on the LQR subroutine in the interest of simplicity, clarity, and in order to leverage prior analyses of model-free methods for LQR. Extending our procedure to robust-control objectives is an exciting direction for future work.

Lastly, while our theoretical analysis only guarantees that the resulting controller will be stabilizing within a small neighborhood of the equilibrium, our simulations on nonlinear systems, such as the nonlinear cartpole, illustrate that discount annealing produces controllers that are competitive with established robust control procedures, such as \mathcal{H}_∞ synthesis, without requiring any knowledge of the underlying dynamics.

1.2 Related work

Given its central importance to the field, stabilization of unknown and uncertain dynamical systems has received extensive attention within the controls literature. We review some of the relevant literature and point the reader towards classical texts for a more comprehensive treatment [Sontag, 2009, Sastry and Bodson, 2011, Zhou et al., 1996, Callier and Desoer, 2012, Zhou and Doyle, 1998].

Model-based approaches. Model-based methods construct approximate system models in order to synthesize stabilizing control policies. Traditional analyses consider stabilization of both linear and

nonlinear dynamical systems in the asymptotic limit of sufficient data [Sastry and Bodson, 2011, Sastry, 2013]. More recent, non-asymptotic studies have focused almost entirely on *linear* systems, where the controller is generated using data from multiple independent trajectories [Fiechter, 1997, Dean et al., 2019, Faradonbeh et al., 2018a, 2019]. Assuming the model is known, stabilizing policies may also be synthesized via convex optimization [Prajna et al., 2004] by combining a ‘dual Lyapunov theorem’ [Rantzer, 2001] with sum-of-squares programming [Parrilo, 2003]. Relative to these analysis our focus is on strengthening the theoretical foundations of model-free procedures and establishing rigorous guarantees that policy gradient methods can also be used to generate stabilizing controllers.

Online control. Online control studies the problem of adaptively *fine-tuning* the performance of an already-stabilizing control policy on a *single* trajectory [Dean et al., 2018, Faradonbeh et al., 2018b, Cohen et al., 2019, Mania et al., 2019, Simchowitz and Foster, 2020, Hazan et al., 2020, Simchowitz et al., 2020, Kakade et al., 2020]. Though early papers in this direction consider systems without pre-given stabilizing controllers [Abbasi-Yadkori and Szepesvári, 2011], their guarantees degrade exponentially in the system dimension (a penalty ultimately shown to be unavoidable by Chen and Hazan [2021]). Rather than fine-tuning an already stabilizing controller, we focus on the more basic problem of finding a controller which is stabilizing in the first place, and allow for the use of multiple independent trajectories.

Model-free approaches. Model-free approaches eschew trying to approximate the underlying dynamics and instead directly search over the space of control policies. The landmark paper of Fazel et al. [2018] proves that, despite the non-convexity of the problem, direct policy search on the infinite-horizon LQR objective efficiently converges to the globally optimal policy, assuming the search is initialized at an already stabilizing controller. Fazel et al. [2018] pose the synthesis of this initial stabilizing controller via policy gradients as an open problem; one that we solve in this work.

Following this result, there have been a large number of works studying policy gradients procedures in continuous control, see for example Feng and Lavaei [2020], Malik et al. [2019], Mohammadi et al. [2020, 2021], Zhang et al. [2021] just to name a few. Relative to our analysis, these papers consider questions of policy finite-tuning, derivative-free methods, and robust (or distributed) control which are important, yet somewhat orthogonal to the stabilization question considered herein. The recent analysis by Lamperski [2020] is perhaps the most closely related piece of prior work. It proposes a model-free, off-policy algorithm for computing a stabilizing controller for deterministic LQR systems. Much like discount annealing, the algorithm also works by alternating between policy optimization (in their case by a closed-form policy improvement step based on the Riccati update) and increasing a damping factor. However, whereas we provide precise finite-time convergence guarantees to a stabilizing controller for both linear and nonlinear systems, the guarantees in Lamperski [2020] are entirely asymptotic and restricted to linear systems. Furthermore, we pay special attention to quantifying the various error tolerances in the gradient and cost queries to ensure that the algorithm can be efficiently implemented in finite samples.

1.3 Background on stability of dynamical systems

Before introducing our results, we first review some of the basic concepts and definitions regarding stability of dynamical systems. In this paper, we study discrete-time, noiseless, time-invariant dynamical systems with states $\mathbf{x}_t \in \mathbb{R}^{d_x}$ and control inputs $\mathbf{u}_t \in \mathbb{R}^{d_u}$. In particular, given an initial state \mathbf{x}_0 , the dynamics evolves according to $\mathbf{x}_{t+1} = G(\mathbf{x}_t, \mathbf{u}_t)$ where $G : \mathbb{R}^{d_x} \times \mathbb{R}^{d_u} \rightarrow \mathbb{R}^{d_x}$ is a state transition map. An equilibrium point of a dynamical system is a state \mathbf{x}_* such that $G(\mathbf{x}_*, 0) = \mathbf{x}_*$. As per convention, we assume that the origin $\mathbf{x}_* = 0$ is the desired equilibrium point around which we wish to stabilize the system.

This paper restricts its attention to static state-feedback policies of the form $\mathbf{u}_t = K\mathbf{x}_t$ for a fixed matrix $K \in \mathbb{R}^{d_u \times d_x}$. Abusing notation slightly, we conflate the matrix K with its induced policy. Our aim is to find a policy K which is *exponentially stabilizing* around the equilibrium point.

Time-invariant, linear systems, where $G(\mathbf{x}, \mathbf{u}) = A\mathbf{x} + B\mathbf{u}$ are stabilizable if and only if there exists a K such that $A+BK$ is a stable matrix [Callier and Desoer, 2012]. That is if $\rho(A+BK) < 1$, where $\rho(X)$ denotes the spectral radius, or the largest eigenvalue magnitude, of a matrix X . For general nonlinear systems, our goal is to find controllers which satisfy the following general, quantitative

definition of exponential stability (e.g Chapter 5.2 in Sastry [2013]). Throughout, $\|\cdot\|$ denotes the Euclidean norm.

Definition 1.1. A controller K is (m, α) -exponentially stable for dynamics G if there exist constants $m, \alpha > 0$ such that if inputs are chosen according to $\mathbf{u}_t = K\mathbf{x}_t$, the sequence of states $\mathbf{x}_{t+1} = G(\mathbf{x}_t, \mathbf{u}_t)$ satisfy

$$\|\mathbf{x}_t\| \leq m \cdot \exp(-\alpha \cdot t) \|\mathbf{x}_0\|. \quad (1.1)$$

Likewise, K is (m, α) -exponentially stable on radius $r > 0$ if (1.1) holds for all \mathbf{x}_0 such that $\|\mathbf{x}_0\| \leq r$.

For linear systems, a controller K is stabilizing if and only if it is stable over the entire state space, however, the restriction to stabilization over a particular radius is in general needed for nonlinear systems. Our approach for stabilizing nonlinear systems relies on analyzing their *Jacobian linearization* about the origin equilibrium. Given a continuously differentiable transition operator G , the local dynamics can be approximated by the Jacobian linearization $(A_{\text{jac}}, B_{\text{jac}})$ of G about the zero equilibrium; that is

$$A_{\text{jac}} := \nabla_{\mathbf{x}} G(\mathbf{x}, \mathbf{u}) \Big|_{(\mathbf{x}, \mathbf{u})=(0,0)}, \quad B_{\text{jac}} := \nabla_{\mathbf{u}} G(\mathbf{x}, \mathbf{u}) \Big|_{(\mathbf{x}, \mathbf{u})=(0,0)}. \quad (1.2)$$

In particular, for \mathbf{x} and \mathbf{u} sufficiently small, $G(\mathbf{x}, \mathbf{u}) = A_{\text{jac}}\mathbf{x} + B_{\text{jac}}\mathbf{u} + f_{\text{nl}}(\mathbf{x}, \mathbf{u})$, where $f_{\text{nl}}(\mathbf{x}, \mathbf{u})$ is a nonlinear remainder from the Taylor expansion of G . To ensure stabilization via state-feedback is feasible, we assume throughout our presentation that the linearized dynamics $(A_{\text{jac}}, B_{\text{jac}})$ are stabilizable.

2 Stabilizing Linear Dynamical Systems

We now present our main results establishing how our algorithm, discount annealing, provably stabilizes linear dynamical systems via a reduction to direct policy search methods. We begin with the following preliminaries on the Linear Quadratic Regulator (LQR).

Definition 2.1 (LQR Objective). For a given starting state \mathbf{x} , we define the LQR problem J_{lin} with discount factor $\gamma \in (0, 1]$, dynamic matrices (A, B) , and state feedback controller K as,

$$J_{\text{lin}}(K \mid \mathbf{x}, \gamma, A, B) := \sum_{t=0}^{\infty} \gamma^t (\mathbf{x}_t^\top Q \mathbf{x}_t + \mathbf{u}_t^\top R \mathbf{u}_t) \text{ s.t. } \mathbf{u}_t = K\mathbf{x}_t, \mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t, \mathbf{x}_0 = \mathbf{x}.$$

Here, $\mathbf{x}_t \in \mathbb{R}^{d_x}$, $\mathbf{u}_t \in \mathbb{R}^{d_u}$, and Q, R are positive definite matrices. Slightly overloading notation, we define

$$J_{\text{lin}}(K \mid \gamma, A, B) := \mathbb{E}_{\mathbf{x}_0 \sim \sqrt{d_x} \cdot \mathcal{S}^{d_x-1}} [J_{\text{lin}}(K \mid \mathbf{x}, \gamma, A, B)],$$

to be the same as the problem above, but where the initial state is now drawn from the uniform distribution over the sphere in \mathbb{R}^{d_x} of radius $\sqrt{d_x}$.²

To simplify our presentation, we adopt the shorthand $J_{\text{lin}}(K \mid \gamma) := J_{\text{lin}}(K \mid \gamma, A, B)$ in cases where the system dynamics (A, B) are understood from context. Furthermore, we assume that (A, B) is stabilizable and that $\lambda_{\min}(Q), \lambda_{\min}(R) \geq 1$. It is a well-known fact that $K_{*, \gamma} := \arg \min_K J_{\text{lin}}(K \mid \gamma, A, B)$ achieves the minimum LQR cost over all possible control laws. We begin our analysis with the observation that the discounted LQR problem is equivalent to the undiscounted LQR problem with damped dynamics matrices.³

Lemma 2.1. For all controllers K such that $J_{\text{lin}}(K \mid \gamma, A, B) < \infty$,

$$J_{\text{lin}}(K \mid \gamma, A, B) = J_{\text{lin}}(K \mid 1, \sqrt{\gamma} \cdot A, \sqrt{\gamma} \cdot B).$$

From this equivalence, it follows from basic facts about LQR that a controller K satisfies $J_{\text{lin}}(0 \mid \gamma, A, B) < \infty$ if and only if $\sqrt{\gamma}(A + BK)$ is stable. Consequently, for $\gamma < \rho(A)^{-2}$, the zero

²This scaling is chosen so that the initial state distribution has identity covariance, and yields cost equivalent to $\mathbf{x}_0 \sim \mathcal{N}(0, I)$.

³This lemma is folklore within the controls community, see e.g. Lamperski [2020].

Discount Annealing

Initialize: Objective $J(\cdot | \cdot)$, $\gamma_0 \in (0, \rho(A)^{-2})$, $K_0 \leftarrow 0$, and $Q \leftarrow I, R \leftarrow I$

For $t = 0, 1, \dots$

1. If $\gamma_t = 1$, run policy gradients once more as in Step 2, break, and return the resulting K' .

2. Using policy gradients (see Eq. (2.1)) initialized at K_t , find K' such that:

$$J_{\text{lin}}(K' | \gamma_t) - \min_K J_{\text{lin}}(K | \gamma_t) \leq d_x. \quad (2.2)$$

3. Update initial controller $K_{t+1} \leftarrow K'$.

4. Using binary or random search, find a discount factor $\gamma' \in [\gamma_t, 1]$ such that

$$2.5J(K_{t+1} | \gamma_t) \leq J(K_{t+1} | \gamma') \leq 8J(K_{t+1} | \gamma_t). \quad (2.3)$$

5. Update the discount factor $\gamma_{t+1} \leftarrow \gamma'$.

Figure 1: Discount annealing algorithm. The procedure is identical for both linear and nonlinear systems. For linear, we initialize $J = J_{\text{lin}}(\cdot | \gamma_0)$ and for nonlinear $J = J_{\text{nl}}(\cdot | \gamma_0, r_*)$ where r_* is chosen as in Theorem 2. See Theorem 1, Theorem 2, and Appendix C for details regarding policy gradients and binary (or random) search. The constants above are chosen for convenience, any constants c_1, c_2 such that $1 < c_1 < c_2$ suffice.

controller is stabilizing and one can solve the discounted LQR problem via direct policy search initialized at $K = 0$ [Fazel et al., 2018]. At this point, one may wonder whether the solution to this highly discounted problem yields a controller which stabilizes the undiscounted system. If this were true, running policy gradients (defined in Eq. (2.1)) to convergence, on a single discounted LQR problem, would suffice to find a stabilizing controller.

$$K_{t+1} = K_t - \eta \nabla_K J_{\text{lin}}(K_t | \gamma), \quad K_0 = 0, \quad \eta > 0 \quad (2.1)$$

Unfortunately, the following proposition shows that this is not the case.

Proposition 2.2 (Impossibility of Reward Shaping). *Fix $A = \text{diag}(0, 2)$. For any positive definite cost matrices Q, R and discount factor γ such that $\sqrt{\gamma}A$ is stable, there exists a matrix B such that (A, B) is controllable (and thus stabilizable), yet the optimal controller $K_{*,\gamma} := \arg \min_K J(K | \gamma, A, B)$ on the discounted problem is such that $A + BK_{*,\gamma}$ is unstable.*

We now describe the discount annealing procedure for linear systems (Figure 1), which provably recovers a stabilizing controller K . For simplicity, we present the algorithm assuming access to noisy, bounded cost and gradient evaluations which satisfy the following definition. Employing standard arguments from [Fazel et al., 2018, Flaxman et al., 2005], we illustrate how these evaluations can be efficiently implemented using polynomially many samples drawn from a simulator in Appendix C.

Definition 2.2 (Gradient and Cost Queries). Given an error parameter $\varepsilon > 0$ and a function $J : \mathbb{R}^d \rightarrow \mathbb{R}$, ε -Grad(J, \mathbf{z}) returns a vector $\tilde{\nabla}$ such that $\|\tilde{\nabla} - \nabla J(\mathbf{z})\|_F \leq \varepsilon$. Similarly, ε -Eval(J, \mathbf{z}, c) returns a scalar v such that $|v - \min\{J(\mathbf{z}), c\}| \leq \varepsilon$.

The procedure leverages the equivalence (Lemma 2.1) between discounted costs and damped dynamics for LQR, and the consequence that the zero controller is stabilizing if we choose γ_0 sufficiently small. Hence, for this discount factor, we may apply policy gradients initialized at the zero controller in order to recover a controller K_1 which is near-optimal for the γ_0 discounted objective.

Our key insight is that, due to known stability margins for LQR controllers, K_1 is stabilizing for the γ_1 discounted dynamics for some discount factor $\gamma_1 > (1 + c)\gamma_0$, where c is a small constant that has a uniform lower bound. Therefore, K_1 has finite cost on the γ_1 discounted problem, so that we may again use policy gradients initialized at K_1 to compute a near-optimal controller K_2 for this larger discount factor. By iterating, we have that $\gamma_t \geq (1 + c)^t \gamma_0$ and can increase the discount factor up to 1, yielding a near-optimal stabilizing controller for the undiscounted LQR objective.

The rate at which we can increase the discount factors γ_t depends on certain properties of the (unknown) dynamical system. Therefore, we opt for binary search to compute the desired γ in the absence of system knowledge. This yields the following guarantee, which we state in terms of properties of the matrix P_* , the optimal value function for the undiscounted LQR problem, which satisfies $\min_K J_{\text{lin}}(K | 1) = \text{tr}[P_*]$ (see [Appendix A](#) for further details).

Theorem 1 (Linear Systems). *Let $M_{\text{lin}} := \max\{16\text{tr}[P_*], J_{\text{lin}}(K_0 | \gamma_0)\}$. The following statements are true regarding the discount annealing algorithm when run on linear dynamical systems:*

- a) *Discount annealing returns a controller \hat{K} which is $(\sqrt{2\text{tr}[P_*]}, (4\text{tr}[P_*])^{-1})$ -exponentially stable.*
- b) *If $\gamma_0 < 1$, the algorithm is guaranteed to halt whenever t is greater than $64\text{tr}[P_*]^4 \log(1/\gamma_0)$.*

Furthermore, at each iteration t :

- c) *Policy gradients as defined in [Eq. \(2.1\)](#) achieves the guarantee in [Eq. \(2.2\)](#) using only $\text{poly}(M_{\text{lin}}, \|A\|_{\text{op}}, \|B\|_{\text{op}})$ many queries to $\varepsilon\text{-Grad}(\cdot, J_{\text{lin}}(\cdot | \gamma))$ as long as ε is less than $\text{poly}(M_{\text{lin}}^{-1}, \|A\|_{\text{op}}^{-1}, \|B\|_{\text{op}}^{-1})$.*
- c) *The noisy binary search algorithm (see [Figure 2](#)) returns a discount factor γ' satisfying [Eq. \(2.3\)](#) using at most $\lceil 4 \log(\text{tr}[P_*]) \rceil + 10$ many queries to $\varepsilon\text{-Eval}(\cdot, J_{\text{lin}}(\cdot | \gamma))$ for $\varepsilon = .1d_x$.*

We remark that since ε need only be polynomially small in the relevant problem parameters, each call to $\varepsilon\text{-Grad}$ and $\varepsilon\text{-Eval}$ can be carried out using only polynomially many samples from a simulator which returns finite horizon system trajectories under various control policies. We make this claim formal in [Appendix C](#).

Proof. We prove part b) of the theorem and defer the proofs of the remaining parts of [Appendix A](#). Define $P_{K,\gamma}$ to be the solution to the discrete-time Lyapunov equation. That is for $\sqrt{\gamma}(A + BK)$ stable, $P_{K,\gamma}$ solves:

$$P_{K,\gamma} := Q + K^\top R K + \gamma(A + BK)^\top P_{K,\gamma}(A + BK). \quad (2.4)$$

Using this notation, $P_* = P_{K_*,1}$ is the solution to the above Lyapunov equation with $\gamma = 1$. The key step of the proof is [Proposition A.4](#), which uses Lyapunov theory to verify the following: given the current discount factor γ_t , an idealized discount factor γ'_{t+1} defined by

$$\gamma'_{t+1} := \left(\frac{1}{8\|P_{K_{t+1},\gamma_t}\|_{\text{op}}^4} + 1 \right)^2 \gamma_t,$$

satisfies $J_{\text{lin}}(K_{t+1} | \gamma'_{t+1}) = \text{tr}[P_{K_{t+1},\gamma'_{t+1}}] \leq 2\text{tr}[P_{K_{t+1},\gamma_t}] = 2J_{\text{lin}}(K_{t+1} | \gamma_t)$. Since the control cost is non-decreasing in γ , the binary search update in Step 4 ensures that the actual γ_{t+1} also satisfies

$$\gamma_{t+1} \geq \left(\frac{1}{8\|P_{K_{t+1},\gamma_t}\|_{\text{op}}^4} + 1 \right)^2 \gamma_t \geq \left(\frac{1}{128\text{tr}[P_*]^4} + 1 \right)^2 \gamma_t,$$

The following calculation (which uses $d_x \leq \text{tr}[P_*]$ for $\lambda_{\min}(Q) \geq 1$) justifies the second inequality above:

$$\|P_{K_{t+1},\gamma_t}\|_{\text{op}}^4 \leq \text{tr}[P_{K_{t+1},\gamma_t}]^4 = J_{\text{lin}}(K_{t+1} | \gamma_t)^4 \leq (\min_K J_{\text{lin}}(K | \gamma_t) + d_x)^4 \leq 16\text{tr}[P_*]^4.$$

Therefore, $\gamma_t \geq (1/(128\text{tr}[P_*]^4) + 1)^{2t} \gamma_0$. The precise bound follows from taking logs of both sides and using the numerical inequality $\log(1+x) \leq x$ to simplify the denominator. \square

3 Stabilizing Nonlinear Dynamical Systems

We now extend the guarantees of the discount annealing algorithm to smooth, nonlinear systems. Whereas our study of linear systems explicitly leveraged the equivalence of discounted costs and damped dynamics, our analysis for nonlinear systems *requires* access to system rollouts under damped dynamics, since the previous equivalence between discounting and damping breaks down in nonlinear settings.

More specifically, in this section, we assume access to a simulator which given a controller K , returns trajectories generated according to $\mathbf{x}_{t+1} = \sqrt{\gamma}G_{\text{nl}}(\mathbf{x}_t, K\mathbf{x}_t)$ for any damping factor $\gamma \in (0, 1]$, where G_{nl} is the transition operator for the nonlinear system. While such trajectories may be infeasible to generate on a physical system, we believe these are reasonable to consider when dynamics are represented using software simulators, as is often the case in practice [Lewis et al., 2003, Peng et al., 2018].

The discount annealing algorithm for nonlinear systems is almost identical to the algorithm for linear systems. It again works by repeatedly solving a series of quadratic cost objectives on the nonlinear dynamics as defined below, and progressively increasing the damping factor γ .

Definition 3.1 (Nonlinear Objective). For a state feedback controller $K : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_u}$, damping factor $\gamma \in (0, 1]$, and an initial state \mathbf{x} , we define:

$$J_{\text{nl}}(K \mid \mathbf{x}, \gamma) := \sum_{t=0}^{\infty} \mathbf{x}_t^\top Q \mathbf{x}_t + \mathbf{u}_t^\top R \mathbf{u}_t \quad (3.1)$$

$$\text{s.t } \mathbf{u}_t = K\mathbf{x}_t, \quad \mathbf{x}_{t+1} = \sqrt{\gamma} \cdot G_{\text{nl}}(\mathbf{x}_t, \mathbf{u}_t), \quad \mathbf{x}_0 = \mathbf{x}. \quad (3.2)$$

Overloading notation as before, we let $J_{\text{nl}}(K \mid \gamma, r) := \mathbb{E}_{\mathbf{x} \sim r \cdot \mathcal{S}^{d_x-1}} [J_{\text{nl}}(K \mid \gamma, \mathbf{x})] \times \frac{d_x}{r^2}$.

The normalization by d_x/r^2 above is chosen so that the nonlinear objective coincides with the LQR objective when G_{nl} is in fact linear. Relative to the linear case, the only algorithmic difference for nonlinear systems is that we introduce an extra parameter r which determines the radius for the initial state distribution. As established in [Theorem 2](#), this parameter must be chosen small enough to ensure that discount annealing succeeds. Our analysis pertains to dynamics which satisfy the following smoothness definition.

Assumption 1 (Local Smoothness). The transition map G_{nl} is continuously differentiable. Furthermore, there exist $r_{\text{nl}}, \beta_{\text{nl}} > 0$ such that for all $(\mathbf{x}, \mathbf{u}) \in \mathbb{R}^{d_x+d_u}$ with $\|\mathbf{x}\| + \|\mathbf{u}\| \leq r_{\text{nl}}$,

$$\|\nabla_{\mathbf{x}, \mathbf{u}} G_{\text{nl}}(\mathbf{x}, \mathbf{u}) - \nabla_{\mathbf{x}, \mathbf{u}} G_{\text{nl}}(0, 0)\|_{\text{op}} \leq \beta_{\text{nl}}(\|\mathbf{x}\| + \|\mathbf{u}\|).$$

For simplicity, we assume $\beta_{\text{nl}} \geq 1$ and $r_{\text{nl}} \leq 1$. Using [Assumption 1](#), we can apply Taylor's theorem to rewrite G_{nl} as its Jacobian linearization around the equilibrium point, plus a nonlinear remainder term.

Lemma 3.1. *If G_{nl} satisfies [Assumption 1](#), then all \mathbf{x}, \mathbf{u} for which $\|\mathbf{x}\| + \|\mathbf{u}\| \leq r_{\text{nl}}$,*

$$G_{\text{nl}}(\mathbf{x}, \mathbf{u}) = A_{\text{jac}}\mathbf{x} + B_{\text{jac}}\mathbf{u} + f_{\text{nl}}(\mathbf{x}, \mathbf{u}), \quad (3.3)$$

where $\|f_{\text{nl}}(\mathbf{x}, \mathbf{u})\| \leq \beta_{\text{nl}}(\|\mathbf{x}\|^2 + \|\mathbf{u}\|^2)$, $\|\nabla f_{\text{nl}}(\mathbf{x}, \mathbf{u})\| \leq \beta_{\text{nl}}(\|\mathbf{x}\| + \|\mathbf{u}\|)$, and where $(A_{\text{jac}}, B_{\text{jac}})$ are the system's Jacobian linearization matrices defined in [Eq. \(1.2\)](#).

Rather than trying to directly understand the behavior of stabilization procedures on the nonlinear system, the key insight of our nonlinear analysis is that we can reason about the performance of a state-feedback controller on the nonlinear system via its behavior on the system's Jacobian linearization. In particular, the following lemma establishes how any controller which achieves finite discounted LQR cost for the Jacobian linearization is guaranteed to be exponentially stabilizing on the damped nonlinear system for initial states that are small enough. Throughout the remainder of this section, we define $J_{\text{lin}}(\cdot \mid \gamma) := J_{\text{lin}}(\cdot \mid \gamma, A_{\text{jac}}, B_{\text{jac}})$ as the LQR objective from [Definition 2.1](#) where $(A, B) = (A_{\text{jac}}, B_{\text{jac}})$.

Lemma 3.2 (Restatement of [Lemma B.2](#)). *Suppose that $\mathcal{C}_J = J_{\text{lin}}(K \mid \gamma) < \infty$, then K is $(\mathcal{C}_J^{1/2}, (4\mathcal{C}_J)^{-1})$ exponentially stable on the damped system $\sqrt{\gamma}G_{\text{nl}}$ over radius $r = r_{\text{nl}} / (\beta_{\text{nl}}\mathcal{C}_J^{3/2})$.*

The second main building block of our nonlinear analysis is the observation that if the dynamics are locally smooth around the equilibrium point, then by [Lemma 3.1](#), decreasing the radius r of the initial state distribution $\mathbf{x}_0 \sim r \cdot \mathcal{S}^{d_x-1}$ reduces the magnitude of the nonlinear remainder term f_{nl} . Hence, the nonlinear system smoothly approximates its Jacobian linearization. More precisely, we establish that the difference in gradients and costs between $J_{\text{nl}}(K \mid \gamma, r)$ and $J_{\text{lin}}(K \mid \gamma)$ decrease linearly with the radius r .

Proposition 3.3. *Assume $J_{\text{lin}}(K \mid \gamma) < \infty$. Then, for $P_{K,\gamma}$ defined as in [Eq. \(2.4\)](#):*

- a) *If $r \leq \frac{r_{\text{nl}}}{2\beta_{\text{nl}}\|P_{K,\gamma}\|_{\text{op}}^2}$, then $|J_{\text{nl}}(K \mid \gamma, r) - J_{\text{lin}}(K \mid \gamma)| \leq 8d_x\beta_{\text{nl}}\|P_{K,\gamma}\|_{\text{op}}^4 \cdot r$.*
- b) *If $r \leq \frac{1}{12\beta_{\text{nl}}\|P_{K,\gamma}\|_{\text{op}}^{5/2}}$, then, $\|\nabla_K J_{\text{nl}}(K \mid \gamma, r) - \nabla_K J_{\text{lin}}(K \mid \gamma)\|_{\text{F}} \leq 48d_x\beta_{\text{nl}}(1 + \|B\|_{\text{op}})\|P_{K,\gamma}\|_{\text{op}}^7 \cdot r$.*

Lastly, because policy gradients on linear dynamical systems is robust to inexact gradient queries, we show that for r sufficiently small, running policy gradients on J_{nl} converges to a controller which has performance close to the optimal controller for the LQR problem with dynamic matrices $(A_{\text{jac}}, B_{\text{jac}})$. As noted previously, we can then use [Lemma 3.2](#) to translate the performance of the optimal LQR controller for the Jacobian linearization to an exponential stability guarantee for the nonlinear dynamics. Using these insights, we establish the following theorem regarding discount annealing for nonlinear dynamics.

Theorem 2 (Nonlinear Systems). *Let $M_{\text{nl}} := \max\{2\text{tr}[P_\star], J_{\text{lin}}(K_0 \mid \gamma_0)\}$. The following statements are true regarding the discount annealing algorithm for nonlinear dynamical systems when r_\star is less than a fixed quantity that is $\text{poly}(1/M_{\text{nl}}, 1/\|A\|_{\text{op}}, 1/\|B\|_{\text{op}}, r_{\text{nl}}/\beta_{\text{nl}})$*

- a) *Discount annealing returns a controller \hat{K} which is $(\sqrt{2\text{tr}[P_\star]}, (8\text{tr}[P_\star])^{-1})$ -exponentially stable over a radius $r = r_{\text{nl}} / (8\beta_{\text{nl}}\text{tr}[P_\star]^2)$*
- b) *If $\gamma_0 < 1$, the algorithm is guaranteed to halt whenever t is greater than $64\text{tr}[P_\star]^4 \log(1/\gamma_0)$.*

Furthermore, at each iteration t :

- c) *Policy gradients achieves the guarantee in [Eq. \(2.2\)](#) using only $\text{poly}(M_{\text{nl}}, \|A\|_{\text{op}}, \|B\|_{\text{op}})$ many queries to $\varepsilon\text{-Grad}(\cdot, J_{\text{nl}}(\cdot \mid \gamma))$ as long as ε is less than some fixed polynomial $\text{poly}(M_{\text{nl}}^{-1}, \|A\|_{\text{op}}^{-1}, \|B\|_{\text{op}}^{-1})$.*
- c) *Let c_0 denote a universal constant. With probability $1 - \delta$, the noisy random search algorithm (see [Figure 2](#)) returns a discount factor γ' satisfying [Eq. \(2.3\)](#) using at most $c_0 \cdot \text{tr}[P_\star]^4 \log(1/\delta)$ queries to $\varepsilon\text{-Eval}(\cdot, J_{\text{nl}}(\cdot \mid \gamma, r_\star))$ for $\varepsilon = .01d_x$.*

We note that while our theorem only guarantees that the controller is stabilizing around a polynomially small neighborhood of the equilibrium, in experiments, we find that the resulting controller successfully stabilizes the dynamics for a wide range of initial conditions. Relative to the case of linear systems where we leveraged the monotonicity of the LQR cost to search for discount factors using binary search, this monotonicity breaks down in the case of nonlinear systems and we instead analyze a random search algorithm to simplify the analysis.

4 Experiments

In this section, we evaluate the ability of the discount annealing algorithm to stabilize a simulated nonlinear system. Specifically, we consider the familiar cart-pole, with $d_x = 4$ (positions and velocities of the cart and pole), and $d_u = 1$ (horizontal force applied to the cart). The goal is to stabilize the system with the pole in the unstable ‘upright’ equilibrium position. For further details, including the precise dynamics, see [Appendix D.1](#). The system was simulated in discrete-time with a simple forward Euler discretization, i.e., $\mathbf{x}_{t+1} = \mathbf{x}_t + T_s \dot{\mathbf{x}}_t$, where $\dot{\mathbf{x}}_t$ is given by the continuous time dynamics, and $T_s = 0.05$ (20Hz). Simulations were carried out in PyTorch [Paszke et al., 2019] and run on a single GPU.

Table 1: Final region of attraction radius r_{roa} as a function of the initial state radius r used during training (discount annealing). We report the [min, max] values of r_{roa} over 5 independent trials. The optimal LQR policy for the linearized system achieved $r_{\text{roa}} = 0.703$ when applied to the nonlinear system. We also synthesized an \mathcal{H}_∞ optimal controller for the linearized dynamics, which achieved $r_{\text{roa}} = 0.506$.

r	0.1	0.3	0.5	0.6	0.7
r_{roa}	[0.702, 0.704]	[0.711, 0.713]	[0.727, 0.734]	[0.731, 0.744]	[0.769, 0.777]

Setup. The discounted annealing algorithm of Figure 1 was implemented as follows. In place of the true infinite horizon discounted cost $J_{\text{nl}}(K | \gamma, r)$ in Eq. (C.4) we use a finite horizon, finite sample Monte Carlo approximation as described in Appendix C,

$$J_{\text{nl}}(K | \gamma) \approx \frac{1}{N} \sum_{i=1}^N J_{\text{nl}}^{(H)}(K | \gamma, \mathbf{x}^{(i)}), \quad \mathbf{x}^{(i)} \sim r \cdot \mathcal{S}^{d_x-1}.$$

Here, $J_{\text{nl}}^{(H)}(K | \gamma, \mathbf{x}) = \sum_{j=0}^{H-1} \mathbf{x}_t^\top Q \mathbf{x}_t + \mathbf{u}_t^\top R \mathbf{u}_t$, is the length H , finite horizon cost of a controller K in which the states evolve according to the $\sqrt{\gamma}$ damped dynamics from Eq. (C.5) and $\mathbf{u}_t = K \mathbf{x}_t$. We used $N = 5000$ and $H = 1000$ in our experiments. For the cost function, we used $Q = T_s \cdot I$ and $R = T_s$. We compute unbiased approximations of the gradients using automatic differentiation on the finite horizon objective $J_{\text{nl}}^{(H)}$.

Instead of using SGD updates for policy gradients, we use Adam [Kingma and Ba, 2014] with a learning rate of $\eta = 0.01/r$. Furthermore, we replace the policy gradient termination criteria in Step 2 (Eq. (2.2)) by instead halting after a fixed number ($M = 200$) of gradient descent steps. We wish to emphasize that the hyperparameters (N, H, η, M) were not optimized for performance. In particular, for $r = 0.1$, we found that as few as $M = 40$ iterations of policy gradient and horizons as short as $H = 400$ were sufficient. Finally, we used an initial discount factor $\gamma_0 = 0.9 \cdot \|A_{\text{jac}}\|_2^{-2}$, where A_{jac} denotes the linearization of the (discrete-time) cart-pole about the vertical equilibrium.

Results. We now proceed to discuss the performance of the algorithm, focusing on three main properties of interest: i) the number of iterations of discount annealing required to find a stabilizing controller (that is, increase γ_t to 1), ii) the maximum radius r of the ball of initial conditions $\mathbf{x}_0 \sim r \cdot \mathcal{S}^{d_x-1}$ for which discount annealing succeeds at stabilizing the system, and iii) the radius r_{roa} of the largest ball contained within the region of attraction (ROA) for the policy returned by discount annealing. Although the true ROA (the set of all initial conditions such that the closed-loop system converges asymptotically to the equilibrium point) is not necessarily shaped like a ball (as the system is more sensitive to perturbations in the position and velocity of the pole than the cart), we use the term region of attraction radius to refer to the radius of the largest ball contained in the ROA.

Concerning (i), discount annealing reliably returned a stabilizing policy in less than 9 iterations. Specifically, over 5 independent trials for each initial radius $r \in \{0.1, 0.3, 0.5, 0.7\}$ (giving 20 independent trials, in total) the algorithm never required more than 9 iterations to return a stabilizing policy.

Concerning (ii), discount annealing reliably stabilized the system for $r \leq 0.7$. For $r \approx 0.75$, we observed trials in which the state of the damped system ($\gamma < 1$) diverged to infinity. For such a rollout, the gradient of the cost is not well-defined, and policy gradient is unable to improve the policy, which prevents discount annealing from finding a stabilizing policy.

Concerning (iii), in Table 1 we report the final radius r_{roa} for the region of attraction of the final controller returned by discount annealing as a function of the training radius r . We make the following observations. Foremost, the policy returned by discount annealing extends the radius of the ROA beyond the radius used during training, i.e. $r_{\text{roa}} > r$. Moreover, for each $r > .1$, the r_{roa} achieved by discount annealing is greater than the $r_{\text{roa}} = 0.703$ achieved by the exact optimal LQR controller and the $r_{\text{roa}} = 0.506$ achieved by the exact optimal \mathcal{H}_∞ controller for the system’s Jacobian linearization (see Table 1). (The \mathcal{H}_∞ optimal controller mitigates the effect of worst-case additive state disturbances on the cost; cf. Appendix D.2 for details).

One may hypothesize that this is due to the fact that discount annealing directly operates on the true nonlinear dynamics whereas the other baselines (LQR and \mathcal{H}_∞ control), find the optimal controller

for an idealized linearization of the dynamics. Indeed, there is evidence to support this hypothesis. In Figure 3 presented in Appendix D, we plot the error $\|K_{\text{pg}}^*(\gamma_t) - K_{\text{lin}}^*(\gamma_t)\|_F$ between the policy $K_{\text{pg}}^*(\gamma_t)$ returned by policy gradients, and the optimal LQR policy $K_{\text{lin}}^*(\gamma_t)$ for the (damped) linearized system, as a function of the discount factor γ_t used in each iteration of the discount annealing algorithm. For small training radius, such as $r = 0.05$, $K_{\text{pg}}^*(\gamma_t) \approx K_{\text{lin}}^*(\gamma_t)$ for all γ_t . However, for larger radii (i.e $r = 0.7$), we see that $\|K_{\text{pg}}^*(\gamma_t) - K_{\text{lin}}^*(\gamma_t)\|_F$ steadily increases as γ_t increases.

That is, as discount annealing increases the discount factor γ and the closed-loop trajectories explore regions of the state space where the dynamics are increasingly nonlinear, K_{pg}^* begins to diverge from K_{lin}^* . Moreover, at the conclusion of discount annealing $K_{\text{pg}}^*(1)$ achieves a lower cost, namely [15.2, 15.4] vs [16.5, 16.8] (here $[a, b]$ denotes $[\min, \max]$ over 5 trials) and larger r_{roa} , namely [0.769, 0.777] vs [0.702, 0.703], than $K_{\text{lin}}^*(1)$, suggesting that the method has indeed adapted to the nonlinearity of the system. Similar observations as to the behavior of controllers fine tuned via policy gradient methods are predicted by the theoretical results from Qu et al. [2020].

5 Discussion

This work illustrates how one can provably stabilize a broad class of dynamical systems via a simple model-free procedure based off policy gradients. In line with the simplicity and flexibility that have made model-free methods so popular in practice, our algorithm works under relatively weak assumptions and with little knowledge of the underlying dynamics. Furthermore, we solve an open problem from previous work [Fazel et al., 2018] and take a step towards placing model-free methods on more solid theoretical footing. We believe that our results raise a number of interesting questions and directions for future work.

In particular, our theoretical analysis states that discount annealing returns a controller whose stability properties are similar to those of the optimal LQR controller for the system’s Jacobian linearization. We were therefore quite surprised when in experiments, the resulting controller had a significantly better radius of attraction than the exact optimal LQR and \mathcal{H}_∞ controllers for the linearization of the dynamics. It is an interesting and important direction for future work to gain a better understanding of exactly when and how model-free procedures are adaptive to the nonlinearities of the system and improve upon these model-based baselines. Furthermore, for our analysis of nonlinear systems, we require access to damped system trajectories. It would be valuable to understand whether this is indeed necessary or whether our analysis could be extended to work without access to damped trajectories.

As a final note, in this work we reduce the problem of stabilizing dynamical systems to running policy gradients on a discounted LQR objective. This choice of reducing to LQR was in part made for simplicity to leverage previous analyses. However, it is possible that overall performance could be improved if rather than reducing to LQR, we instead attempted to run a model-free method that directly tries to optimize a robust control objective (which explicitly deals with uncertainty in the system dynamics). We believe that understanding these tradeoffs in objectives and their relevant sample complexities is an interesting avenue for future inquiry.

Acknowledgments

We would like to thank Peter Bartlett for many helpful conversations and comments throughout the course of this project, and Russ Tedrake for support with the numerical experiments. JCP is supported by an NSF Graduate Research Fellowship. MS is supported by an Open Philanthropy Fellowship grant. JU is supported by the National Science Foundation, Award No. EFMA-1830901, and the Department of the Navy, Office of Naval Research, Award No. N00014-18-1-2210.

References

- Yasin Abbasi-Yadkori and Csaba Szepesvári. Regret bounds for the adaptive control of linear quadratic systems. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 1–26, 2011.
- OpenAI: Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, et al. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1):3–20, 2020.
- Frank M Callier and Charles A Desoer. *Linear system theory*. Springer Science & Business Media, 2012.
- Xinyi Chen and Elad Hazan. Black-box control for linear dynamical systems. In *Conference on Learning Theory*, pages 1114–1143. PMLR, 2021.
- Alon Cohen, Tomer Koren, and Yishay Mansour. Learning linear-quadratic regulators efficiently with only \sqrt{t} regret. In *International Conference on Machine Learning*, pages 1300–1309. PMLR, 2019.
- Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. Regret bounds for robust adaptive control of the linear quadratic regulator. In *Advances in Neural Information Processing Systems*, pages 4188–4197, 2018.
- Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *Foundations of Computational Mathematics*, pages 1–47, 2019.
- Mohamad Kazem Shirani Faradonbeh, Ambuj Tewari, and George Michailidis. Finite-time adaptive stabilization of linear systems. *IEEE Transactions on Automatic Control*, 64(8):3498–3505, 2018a.
- Mohamad Kazem Shirani Faradonbeh, Ambuj Tewari, and George Michailidis. On optimality of adaptive linear-quadratic regulators. *arXiv preprint arXiv:1806.10749*, 2018b.
- Mohamad Kazem Shirani Faradonbeh, Ambuj Tewari, and George Michailidis. Randomized algorithms for data-driven stabilization of stochastic linear systems. In *2019 IEEE Data Science Workshop (DSW)*, pages 170–174. IEEE, 2019.
- Maryam Fazel, Rong Ge, Sham Kakade, and Mehran Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. In *International Conference on Machine Learning*, pages 1467–1476. PMLR, 2018.
- Han Feng and Javad Lavaei. Escaping locally optimal decentralized control policies via damping. In *2020 American Control Conference (ACC)*, pages 50–57. IEEE, 2020.
- Claude-Nicolas Fiechter. Pac adaptive control of linear systems. In *Proceedings of the tenth annual conference on Computational learning theory*, pages 72–80, 1997.
- Abraham D. Flaxman, Adam Tauman Kalai, and H. Brendan McMahan. Online convex optimization in the bandit setting: gradient descent without a gradient. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '05, page 385–394, USA, 2005. Society for Industrial and Applied Mathematics.
- Elad Hazan, Sham Kakade, and Karan Singh. The nonstochastic control problem. In *Algorithmic Learning Theory*, pages 408–421. PMLR, 2020.
- Sham Kakade, Akshay Krishnamurthy, Kendall Lowrey, Motoya Ohnishi, and Wen Sun. Information theoretic regret bounds for online nonlinear control. In *Advances in Neural Information Processing Systems*, volume 33, pages 15312–15325, 2020.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

- Andrew Lamperski. Computing stabilizing linear controllers via policy iteration. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 1902–1907. IEEE, 2020.
- Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, pages 1302–1338, 2000.
- Frank L Lewis, Darren M Dawson, and Chaouki T Abdallah. *Robot manipulator control: theory and practice*. CRC Press, 2003.
- Dhruv Malik, Ashwin Pananjady, Kush Bhatia, Koulik Khamaru, Peter Bartlett, and Martin Wainwright. Derivative-free methods for policy optimization: Guarantees for linear quadratic systems. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2916–2925. PMLR, 2019.
- Horia Mania, Stephen Tu, and Benjamin Recht. Certainty equivalence is efficient for linear quadratic control. In *Advances in Neural Information Processing Systems*, pages 10154–10164, 2019.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- Hesameddin Mohammadi, Mahdi Soltanolkotabi, and Mihailo R Jovanović. On the linear convergence of random search for discrete-time lqr. *IEEE Control Systems Letters*, 5(3):989–994, 2020.
- Hesameddin Mohammadi, Armin Zare, Mahdi Soltanolkotabi, and Mihailo R Jovanovic. Convergence and sample complexity of gradient methods for the model-free linear quadratic regulator problem. *IEEE Transactions on Automatic Control*, 2021.
- Pablo A Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Mathematical programming*, 96(2):293–320, 2003.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703*, 2019.
- Xue Bin Peng, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. Sim-to-real transfer of robotic control with dynamics randomization. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3803–3810, 2018. doi: 10.1109/ICRA.2018.8460528.
- Juan C. Perdomo, Max Simchowitz, Alekh Agarwal, and Peter L. Bartlett. Towards a dimension-free understanding of adaptive linear control. In *Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pages 3681–3770. PMLR, 2021.
- Stephen Prajna, Pablo A Parrilo, and Anders Rantzer. Nonlinear control synthesis by convex optimization. *IEEE Transactions on Automatic Control*, 49(2):310–314, 2004.
- Guannan Qu, Chenkai Yu, Steven Low, and Adam Wierman. Combining model-based and model-free methods for nonlinear control: A provably convergent policy gradient approach. *arXiv preprint arXiv:2006.07476*, 2020.
- Anders Rantzer. A dual to lyapunov’s stability theorem. *Systems & Control Letters*, 42(3):161–168, 2001.
- Shankar Sastry. *Nonlinear systems: analysis, stability, and control*, volume 10. Springer Science & Business Media, 2013.
- Shankar Sastry and Marc Bodson. *Adaptive control: stability, convergence and robustness*. Courier Corporation, 2011.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- Max Simchowitz and Dylan Foster. Naive exploration is optimal for online lqr. In *International Conference on Machine Learning*, pages 8937–8948. PMLR, 2020.

- Max Simchowitz, Karan Singh, and Elad Hazan. Improper learning for non-stochastic control. In *Conference on Learning Theory*, pages 3320–3436. PMLR, 2020.
- Eduardo D Sontag. Nonlinear feedback stabilization revisited. In *Dynamical Systems, Control, Coding, Computer Vision*, pages 223–262. Springer, 1999.
- Eduardo D. Sontag. *Stability and feedback stabilization*. Springer New York, New York, NY, 2009. ISBN 978-0-387-30440-3. doi: 10.1007/978-0-387-30440-3_515.
- Stephen Tu and Benjamin Recht. The gap between model-based and model-free methods on the linear quadratic regulator: An asymptotic viewpoint. In *Conference on Learning Theory*, pages 3036–3083. PMLR, 2019.
- Kaiqing Zhang, Xiangyuan Zhang, Bin Hu, and Tamer Başar. Derivative-free policy optimization for risk-sensitive and robust control design: implicit regularization and sample complexity. *arXiv preprint arXiv:2101.01041*, 2021.
- Kemin Zhou and John Comstock Doyle. *Essentials of robust control*, volume 104. Prentice Hall, 1998.
- Kemin Zhou, John Comstock Doyle, Keith Glover, et al. *Robust and optimal control*, volume 40. Prentice hall New Jersey, 1996.

Table of Contents: Appendix

A	Deferred Proofs and Analysis for the Linear Setting	14
A.1	Discount annealing on linear systems: Proof of Theorem 1	15
A.2	Convergence of policy gradients for LQR: Proof of Lemma A.3	17
A.3	Impossibility of reward shaping: Proof of Proposition 2.2	18
A.4	Auxiliary results for linear systems	19
B	Deferred Proofs and Analysis for the Nonlinear Setting	20
B.1	Discount annealing on nonlinear systems: Proof of Theorem 2	22
B.2	Relating costs and gradients to the linear system: Proof of Proposition 3.3	24
B.3	Establishing Lyapunov functions: Proof of Lemma B.1	30
B.4	Bounding the nonlinearity: Proof of Lemma 3.1	31
C	Implementing ε-Eval, ε-Grad, and Search Algorithms	32
C.1	Implementing ε -Eval	32
C.2	Implementing ε -Grad	33
C.3	Auxiliary results	33
C.4	Search analysis	38
D	Additional Experiments	39
D.1	Cart-pole dynamics	39
D.2	\mathcal{H}_∞ synthesis	39
D.3	Difference between LQR and discount annealing as a function of discount	39

A Deferred Proofs and Analysis for the Linear Setting

Preliminaries on Linear Quadratic Control. The cost a state-feedback controller $J_{\text{lin}}(K \mid \gamma)$ is intimately related to the solution of the discrete-time Lyapunov equation. Given a stable matrix A_{cl} and a symmetric positive definite matrix Σ , we define $\text{dlyap}(A_{\text{cl}}, \Sigma)$ to be the unique solution (over X) to the matrix equation,

$$X = \Sigma + A_{\text{cl}}^\top X A_{\text{cl}}.$$

A classical result in Lyapunov theory states that $\text{dlyap}(A_{\text{cl}}, \Sigma) = \sum_{j=0}^{\infty} (A_{\text{cl}}^j)^\top X A_{\text{cl}}^j$. Recalling our earlier definition, for a controller K such that $\sqrt{\gamma}(A+BK)$ is stable, we let $P_{K,\gamma} := \text{dlyap}(\sqrt{\gamma}(A+BK), Q+K^\top RK)$, where Q, R are the cost matrices for the LQR problem defined in [Definition 2.1](#). As a special case, we let $P_\star := \text{dlyap}(A+BK_\star, Q+K_\star^\top RK_\star)$ where $K_\star := K_{\star,1}$ is the optimal controller for the undiscounted LQR problem. Using these definitions, we have the following facts:

Fact A.1. $J_{\text{lin}}(K \mid \gamma, A, B) < \infty$ if and only if $\sqrt{\gamma} \cdot (A+BK)$ is a stable matrix.

Fact A.2. If $J_{\text{lin}}(K \mid \gamma, A, B) < \infty$ then for all $\mathbf{x} \in \mathbb{R}^{d_x}$, $J_{\text{lin}}(K \mid \mathbf{x}, \gamma, A, B) = \mathbf{x}^\top P_{K,\gamma} \mathbf{x}$. Furthermore,

$$J_{\text{lin}}(K \mid \gamma, A, B) = \mathbb{E}_{\mathbf{x}_0 \sim \sqrt{d_x} S^{d_x-1}} \mathbf{x}_0^\top P_{K,\gamma} \mathbf{x}_0 = \text{tr}[P_{K,\gamma}].$$

Employing these identities, we can now restate and prove [Lemma 2.1](#):

Lemma 2.1. For all K such that $J(K \mid \gamma, A, B) < \infty$, $J_{\text{lin}}(K \mid \gamma, A, B) = J_{\text{lin}}(K \mid 1, \sqrt{\gamma} \cdot A, \sqrt{\gamma} \cdot B)$.

Proof. From the definition of the LQR cost and linear dynamics in [Definition 2.1](#), for $A_{\text{cl}} := A + BK$,

$$\begin{aligned} J_{\text{lin}}(K \mid \gamma, A, B) &= \mathbb{E}_{\mathbf{x}_0} \left[\sum_{t=0}^{\infty} \gamma^t \cdot \mathbf{x}_0^\top (A_{\text{cl}}^t)^\top (Q + K^\top RK) A_{\text{cl}}^t \mathbf{x}_0 \right] \\ &= \mathbb{E}_{\mathbf{x}_0} \left[\sum_{t=0}^{\infty} \mathbf{x}_0^\top ((\sqrt{\gamma} A_{\text{cl}})^t)^\top (Q + K^\top RK) (\sqrt{\gamma} A_{\text{cl}})^t \mathbf{x}_0 \right] \\ &= J_{\text{lin}}(K \mid 1, \sqrt{\gamma} \cdot A, \sqrt{\gamma} \cdot B). \end{aligned}$$

□

Therefore, since LQR with a discount factor is equivalent to LQR with damped dynamics, it follows that for $\gamma_0 < \rho(A)^{-2}$, (noisy) policy gradients initialized at the zero controller converges to the global optimum of the discounted LQR problem. The following lemma is essentially a restatement of the finite sample convergence result for gradient descent on LQR (Theorem 31) from Fazel et al. [2018], where we have set $Q = R = I$ and $\mathbb{E}[\mathbf{x}_0 \mathbf{x}_0^\top] = I$ as in the discount annealing algorithm. We include the proof of this result in [Appendix A.2](#) for the sake of completeness.

Lemma A.3 (Fazel et al. [2018]). *For K_0 such that $J_{\text{lin}}(K_0 \mid \gamma) < \infty$, define (noisy) policy gradients as the procedure which computes updates according to,*

$$K_{t+1} = K_t - \eta \tilde{\nabla}_t,$$

for some matrix $\tilde{\nabla}_t$. There exists a choice of a constant step size $\eta > 0$ and a fixed polynomial,

$$\mathcal{C}_{\text{PG}} := \text{poly} \left(\frac{1}{J_{\text{lin}}(K_0 \mid \gamma)}, \frac{1}{\|A\|_{\text{op}}}, \frac{1}{\|B\|_{\text{op}}} \right),$$

such that if the following inequality holds for all $t = 1 \dots T$,

$$\left\| \nabla_K J_{\text{lin}}(K_t \mid \gamma) - \tilde{\nabla}_t \right\|_{\text{F}} \leq \mathcal{C}_{\text{PG}} \cdot \varepsilon, \quad (\text{A.1})$$

then $J_{\text{lin}}(K_T \mid \gamma) - \min_K J_{\text{lin}}(K \mid \gamma) \leq \varepsilon$ whenever

$$T \geq \text{poly} \left(\|A\|_{\text{op}}, \|B\|_{\text{op}}, J_{\text{lin}}(K_0 \mid \gamma) \right) \log \left(\frac{J_{\text{lin}}(K_0 \mid \gamma) - \min_K J_{\text{lin}}(K \mid \gamma)}{\varepsilon} \right).$$

With this lemma, we can now present the proof of [Theorem 1](#).

A.1 Discount annealing on linear systems: Proof of [Theorem 1](#)

We organize the proof into two main parts. First, we prove statements *c*) and *d*) by an inductive argument. Then, having already proved part *b*) in the main body of the paper, we finish the proof by establishing the stability guarantees of the resulting controller outlined in part *a*).

A.1.1 Proof of *c*) and *d*)

Base case. Recall that at iteration t of discount annealing, policy gradients is initialized at $K_{t,0} := K_t$ and (in the case of linear systems) computes updates according to:

$$K_{t,j+1} = K_{t,j} - \eta \cdot \varepsilon\text{-Grad} (J_{\text{lin}}(\cdot \mid \gamma_t), K_{t,j}).$$

From [Lemma A.3](#), if $J_{\text{lin}}(K_{t,0} \mid \gamma_t) < \infty$ and

$$\|\varepsilon\text{-Grad} (J_{\text{lin}}(\cdot \mid \gamma_t), K_{t,j}) - \nabla J_{\text{lin}}(K_{t,j} \mid \gamma_t)\|_{\text{F}} \leq \mathcal{C}_{\text{pg},t} d_x \quad \forall j \geq 0, \quad (\text{A.2})$$

where $\mathcal{C}_{\text{pg},t} = \text{poly}(1/J_{\text{lin}}(K_{t,0} \mid \gamma_t), 1/\|A\|_{\text{op}}, 1/\|B\|_{\text{op}})$ then policy gradients will converge to a $K_{t,j}$ such that $J_{\text{lin}}(K_{t,j} \mid \gamma_t) - \min_K J_{\text{lin}}(K \mid \gamma_t) \leq d_x$ after $\text{poly}(\|A\|_{\text{op}}, \|B\|_{\text{op}}, J_{\text{lin}}(K_{t,0}))$ many iterations.

By our choice of discount factor, we have that $J_{\text{lin}}(0 \mid \gamma_0) < \infty$. Furthermore, since $\varepsilon \leq \text{poly}(1/J_{\text{lin}}(0 \mid \gamma_0), 1/\|A\|_{\text{op}}, 1/\|B\|_{\text{op}})$, then the condition outlined in [Eq. \(A.2\)](#) holds and policy gradients achieves the desired guarantee when $t = 0$.

The correctness of binary search at iteration $t = 0$ follows from [Lemma C.6](#). In particular, we instantiate the lemma with $f = J_{\text{lin}}(K_{t+1} | \cdot)$ of the LQR objective, which is nondecreasing in γ by definition, $f_1 = 2.5J_{\text{lin}}(K_{t+1} | \gamma_t)$ and $f_2 = 8J_{\text{lin}}(K_{t+1} | \gamma_t)$. The algorithm requires auxiliary values $\bar{f}_1 \in [2.5J_{\text{lin}}(K_{t+1} | \gamma_t), 3J_{\text{lin}}(K_{t+1} | \gamma_t)]$ and $\bar{f}_2 \in [7J_{\text{lin}}(K_{t+1} | \gamma_t), 7.5J_{\text{lin}}(K_{t+1} | \gamma_t)]$ which we can always compute by using $\varepsilon\text{-Eval}$ to estimate the cost $J_{\text{lin}}(K_{t+1} | \gamma_t)$ to precision $.1d_x$ (recall that $J_{\text{lin}}(K | \gamma) \geq d_x$ for any K and γ). The last step needed to apply the lemma is to lower bound the width of the feasible region of γ 's which satisfy the desired criterion that $J_{\text{lin}}(K_{t+1} | \gamma) \in [\bar{f}_1, \bar{f}_2]$.

Let $\gamma' \geq \gamma_t$ be such that $J_{\text{lin}}(K_{t+1} | \gamma') = 3J_{\text{lin}}(K_{t+1} | \gamma_t)$. Such a γ is guaranteed to exist since $J_{\text{lin}}(K_{t+1} | \cdot)$ is nondecreasing in γ and it is a continuous function for all $\gamma < \bar{\gamma} := \sup\{\gamma : J_{\text{lin}}(K_{t+1} | \gamma) < \infty\}$. By the calculation from the proof of part *b*) presented in the main body of the paper, for

$$\gamma'' := \left(\frac{1}{8\|P_{K_{t+1}, \gamma'}\|_{\text{op}}^4} + 1 \right)^2 \gamma',$$

we have that $J_{\text{lin}}(K_{t+1} | \gamma'') \leq 2J_{\text{lin}}(K_{t+1} | \gamma') = 6J_{\text{lin}}(K_{t+1} | \gamma_t)$. By monotonicity and continuity of $J_{\text{lin}}(K_{t+1} | \cdot)$, when restricted to $\gamma \leq \bar{\gamma}$, all $\gamma \in [\gamma', \gamma'']$ satisfy $J_{\text{lin}}(K_{t+1} | \gamma) \in [3J_{\text{lin}}(K_{t+1} | \gamma_t), 6J_{\text{lin}}(K_{t+1} | \gamma_t)]$. Moreover,

$$\gamma'' - \gamma' = \left[\left(\frac{1}{8\|P_{K_{t+1}, \gamma'}\|_{\text{op}}^4} + 1 \right)^2 - 1 \right] \gamma' \geq \frac{1}{4\|P_{K_{t+1}, \gamma'}\|_{\text{op}}^4} \gamma' \geq \frac{1}{4\text{tr}[P_{K_{t+1}, \gamma'}]^4} \gamma_t,$$

where the last line follows from the fact that $\gamma' \geq \gamma_t$ and that the trace of a PSD matrix is always at least as large as the operator norm. Lastly, since $\text{tr}[P_{K_{t+1}, \gamma'}] = J_{\text{lin}}(K_{t+1} | \gamma') = 3J_{\text{lin}}(K_{t+1} | \gamma_t)$, by the guarantee of policy gradients, we have that for $t = 0$, $J_{\text{lin}}(K_1 | \gamma_0) \leq 2\text{tr}[P_\star]$. Therefore, for $t = 0$:

$$\gamma'' - \gamma' \geq \frac{1}{4(6\text{tr}[P_\star])^4}$$

Hence the width of the feasible region is at least $\frac{1}{5184\text{tr}[P_\star]^4} \gamma_0$.

Inductive step. To show that policy gradients achieves the desired guarantee at iteration $t + 1$, we can repeat the exact same argument as in the base case. The only difference is that the need to argue that the cost of the initial controller $\sup_{t \geq 0} J_{\text{lin}}(K_t | \gamma_t)$ is uniformly bounded across iterations. By the inductive hypothesis on the success of the binary search algorithm at iteration $t - 1$, we have that,

$$\begin{aligned} J_{\text{lin}}(K_t | \gamma_t) &\leq 8J_{\text{lin}}(K_t | \gamma_{t-1}) \\ &\leq 8 \left(\min_K J_{\text{lin}}(K | \gamma_{t-1}) + d_x \right) \\ &\leq 8 \left(\min_K J_{\text{lin}}(K | 1) + d_x \right) \\ &\leq 16\text{tr}[P_\star]. \end{aligned}$$

Hence, by [Lemma A.3](#), policy gradients achieves the desired guarantee using $\text{poly}(M_{\text{lin}}, \|A\|_{\text{op}}, \|B\|_{\text{op}})$ many queries to $\varepsilon\text{-Grad}(\cdot, J_{\text{lin}}(\cdot | \gamma))$ as long as ε is less than $\text{poly}(M_{\text{lin}}^{-1}, \|A\|_{\text{op}}^{-1}, \|B\|_{\text{op}}^{-1})$.

Likewise, the argument for the correctness of the binary search procedure is identical to that of the base case. Because of the success of policy gradients and binary search at the previous iteration, we can upper bound, $\text{tr}[P_{K_{t+1}, \gamma'}]$ by $6\text{tr}[P_\star]$ and get a uniform lower bound on the width of the feasible region.

A.1.2 Proof of *a*)

After halting, we see that discount annealing returns a controller \hat{K} satisfying the stated condition from Step 2 requiring that,

$$J_{\text{lin}}(\hat{K} | 1) - J_{\text{lin}}(K_\star | 1) = \text{tr}[\hat{P} - P_\star] \leq d_x.$$

Here, we have used [Fact A.2](#) to rewrite $J_{\text{lin}}(\widehat{K} \mid 1)$ as $\text{tr}[\widehat{P}]$ for $\widehat{P} := \text{dlyap}(A + B\widehat{K}, Q + \widehat{K}^\top R\widehat{K})$ (and likewise for P_\star). Since $\text{tr}[P_\star] \geq d_x$, we conclude that $\text{tr}[\widehat{P}] \leq 2\text{tr}[P_\star]$. Now, by properties of the Lyapunov equation (see [Lemma A.5](#)) the following holds for $A_{\text{cl}} := A + B\widehat{K}$:

$$\|A_{\text{cl}}^2\|_{\text{op}}^t \leq \|\widehat{P}\|_{\text{op}} \left(1 - \frac{1}{\|\widehat{P}\|_{\text{op}}}\right)^t \leq \text{tr}[\widehat{P}] \exp\left(-t/\text{tr}[\widehat{P}]\right).$$

Hence, we conclude that,

$$\|\mathbf{x}_t\| = \|A_{\text{cl}}^t \mathbf{x}_0\| \leq \|A_{\text{cl}}^t\|_{\text{op}} \|\mathbf{x}_0\| \leq \sqrt{2\text{tr}[P_\star]} \exp\left(-\frac{1}{4\text{tr}[P_\star]} \cdot t\right) \|\mathbf{x}_0\|.$$

A.2 Convergence of policy gradients for LQR: Proof of [Lemma A.3](#)

Proof. Note that, by [Lemma 2.1](#), proving the above result for $J_{\text{lin}}(\cdot \mid \gamma, A, B)$ is the same as proving it for $J_{\text{lin}}(\cdot \mid 1, \sqrt{\gamma}A, \sqrt{\gamma}B)$. We start by defining the following idealized updates,

$$\begin{aligned} K' &= K - \eta \nabla_K J_{\text{lin}}(K_t \mid \gamma) \\ K'' &= K - \eta \widetilde{\nabla}_t. \end{aligned}$$

From [Lemmas 13, 24, and 25](#) from [Fazel et al. \[2018\]](#), there exists a fixed polynomial,

$$\mathcal{C}_\eta := \text{poly}\left(\frac{1}{\sqrt{\gamma}\|A\|_{\text{op}}}, \frac{1}{\sqrt{\gamma}\|B\|_{\text{op}}}, \frac{1}{J_{\text{lin}}(K_0 \mid \gamma)}\right)$$

Such that, for $\eta \leq \mathcal{C}_\eta$, the following inequality holds,

$$J_{\text{lin}}(K' \mid \gamma) - J_{\text{lin}}(K_\star \mid \gamma) \leq \left(1 - \eta \frac{1}{\|\Sigma_{K_\star}\|_{\text{op}}}\right) (J_{\text{lin}}(K \mid \gamma) - J_{\text{lin}}(K_\star \mid \gamma)),$$

where $\Sigma_{K_\star} = \mathbb{E}_{\mathbf{x}_0}[\sum_{t=0}^{\infty} \mathbf{x}_{t,\star}^\top \mathbf{x}_{t,\star}]$ and $\{\mathbf{x}_{t,\star}\}_{t=0}^{\infty}$ is the sequence of states generated by the controller K_\star . Therefore, if $J_{\text{lin}}(K'' \mid \gamma)$ and $J_{\text{lin}}(K' \mid \gamma)$ satisfy,

$$|J_{\text{lin}}(K'' \mid \gamma) - J_{\text{lin}}(K' \mid \gamma)| \leq \frac{1}{2\|\Sigma_{K_\star}\|_{\text{op}}} \eta \varepsilon \tag{A.3}$$

then, as long as $J_{\text{lin}}(K \mid \gamma) - J_{\text{lin}}(K_\star \mid \gamma) \geq \varepsilon$, this following inequality also holds:

$$J_{\text{lin}}(K'' \mid \gamma) - J_{\text{lin}}(K_\star \mid \gamma) \leq \left(1 - \eta \frac{1}{2\|\Sigma_{K_\star}\|_{\text{op}}}\right) (J_{\text{lin}}(K \mid \gamma) - J_{\text{lin}}(K_\star \mid \gamma)).$$

The proof then follows by unrolling the recursion and simplifying. We now focus on establishing [Eq. \(A.3\)](#). By [Lemma 27](#) in [Fazel et al. \[2018\]](#), if

$$\|K'' - K'\|_{\text{op}} = \eta \|\nabla_K J_{\text{lin}}(K_j \mid \gamma) - \widetilde{\nabla}_j\|_{\text{op}} \leq \mathcal{C}_K$$

where \mathcal{C}_K is a fixed polynomial $\mathcal{C}_K := \text{poly}\left(\frac{1}{J_{\text{lin}}(K_0 \mid \gamma)}, \frac{1}{\sqrt{\gamma}\|A\|_{\text{op}}}, \frac{1}{\sqrt{\gamma}\|B\|_{\text{op}}}\right)$, then

$$|J_{\text{lin}}(K'' \mid \gamma) - J_{\text{lin}}(K' \mid \gamma)| \leq \mathcal{C}_{\text{cost}} \|K'' - K'\|_{\text{op}} = \mathcal{C}_{\text{cost}} \cdot \eta \|\nabla_K J_{\text{lin}}(K_j \mid \gamma) - \widetilde{\nabla}_j\|_{\text{op}},$$

where $\mathcal{C}_{\text{cost}} := \text{poly}(d_x, \|R\|_{\text{op}}, \|B\|_{\text{op}}, J_{\text{lin}}(K_0))$. Therefore, [Eq. \(A.3\)](#) holds if

$$\|\nabla_K J_{\text{lin}}(K_j \mid \gamma) - \widetilde{\nabla}_j\|_{\text{op}} \leq \min\left\{\frac{1}{2\|\Sigma_{K_\star}\|_{\text{op}} \mathcal{C}_{\text{cost}}}\varepsilon, \mathcal{C}_K\right\}.$$

The exact statement follows from using $d_x \leq J_{\text{lin}}(K_0 \mid \gamma)$ and $\|\Sigma_{K_\star, \gamma}\|_{\text{op}} \leq J_{\text{lin}}(K_\star, \gamma) \leq J_{\text{lin}}(K_0 \mid \gamma)$ by [Lemma 13](#) in [\[Fazel et al., 2018\]](#) and taking the polynomial in the proposition statement to be the minimum of \mathcal{C}_K and $1/\mathcal{C}_{\text{cost}}$. \square

A.3 Impossibility of reward shaping: Proof of Proposition 2.2

Proof. Consider the linear dynamical system with dynamics matrices,

$$A = \begin{bmatrix} 0 & 0 \\ 0 & 2 \end{bmatrix}, \quad B = \begin{bmatrix} 1 \\ \beta \end{bmatrix}$$

where $\beta > 0$ is a parameter to be chosen later. Note that a linear dynamical system of these dimensions is controllable (and hence stabilizable) [Callier and Desoer, 2012], since the matrix

$$[B \quad AB] = \begin{bmatrix} 1 & 0 \\ \beta & 2\beta \end{bmatrix}$$

is full rank. For any controller $K = [k_1 \quad k_2]$, the closed loop system $A_{\text{cl}} := A + BK$ has the form,

$$A_{\text{cl}} = \begin{bmatrix} k_1 & k_2 \\ \beta k_1 & 2 + \beta k_2 \end{bmatrix}.$$

By Gershgorin's circle theorem, A_{cl} has an eigenvalue λ which satisfies,

$$|\lambda| \geq |2 + \beta k_2| - |\beta k_1| \geq 2 - 2\beta \max\{|k_1|, |k_2|\}.$$

Therefore, any controller K for which the closed-loop system $A + BK$ is stable must have the property that,

$$\max\{|k_1|, |k_2|\} \geq \frac{1}{2\beta}.$$

Using this observation and Fact A.2, for any discount factor γ , a stabilizing controller K must satisfy,

$$\begin{aligned} J_{\text{lin}}(K \mid \gamma) &= \text{tr}[P_{K,\gamma}] \\ &\geq \text{tr}[Q] + \text{tr}[K^\top RK] \\ &\geq (k_1^2 + k_2^2) \cdot \sigma_{\min}(R) \\ &\geq \frac{1}{4\beta^2} \cdot \sigma_{\min}(R). \end{aligned}$$

In the above calculation, we have used the identity $P_{K,\gamma} = Q + K^\top RK + \gamma A_{\text{cl}}^\top P_{K,\gamma} A_{\text{cl}}$ as well as the assumption that R is positive definite. Next, we observe that for a discount factor $\gamma = c^2 \cdot \rho(A)^{-2}$, where $c \in (0, 1)$ as chosen in the initial iteration of our algorithm, the cost of the 0 controller has the following upper bound:

$$\begin{aligned} J_{\text{lin}}(0 \mid \gamma_0) &= \sum_{j=0}^{\infty} c^{2j} \rho(A)^{-2j} \cdot \text{tr}[(A^\top)^j Q A^j] \\ &\leq \|Q\|_{\text{op}} \sum_{j=0}^{\infty} c^{2j} \rho(A)^{-2j} \|A^j\|_{\text{F}}^2 \\ &= \|Q\|_{\text{op}} \sum_{j=0}^{\infty} \left\| \left(\frac{c}{\rho(A)} \cdot A \right)^j \right\|_{\text{F}}^2. \end{aligned}$$

Using standard Lyapunov arguments (see for example Section D.2 in Perdomo et al. [2021]) the sum in the last line is a geometric series and is equal to some function $f(c, A) < \infty$, which depends *only* on c and A , for all $c \in (0, 1)$. Using this calculation, it follows that

$$\min_K J_{\text{lin}}(K \mid \gamma_0) \leq J_{\text{lin}}(0 \mid \gamma) \leq \|Q\|_{\text{op}} f(c, A)$$

Hence, for any Q, R , and discount factor $\gamma \in (0, \rho(A)^{-2})$, we can choose β small enough such that,

$$\|Q\|_{\text{op}} f(c, A) < \frac{1}{4\beta^2} \sigma_{\min}(R)$$

implying that the optimal controller $K_{\star,\gamma}$ for the discounted problem cannot be stabilizing for (A, B) . \square

A.4 Auxiliary results for linear systems

Proposition A.4. Let $\sqrt{\gamma_1}(A+BK)$ be a stable matrix and define $P_1 := \text{dlyap}(\sqrt{\gamma_1}(A+BK), Q + K^\top RK)$, then for c defined as

$$c := \frac{1}{8\|P_1\|_{\text{op}}^4} + 1,$$

the following holds for $\gamma_2 := c^2\gamma_1$ and $P_2 := \text{dlyap}(\sqrt{\gamma_2}(A+BK), Q + K^\top RK)$:

$$\text{tr}[P_2 - P_1] \leq \text{tr}[P_1].$$

Proof. The proof is a direct consequence of Proposition C.7 in Perdomo et al. [2021]. In particular, we use their results for the trace norm and use the following substitutions,

$$\begin{aligned} A_1 &\leftarrow \sqrt{\gamma_1}(A+BK) & \Sigma &\leftarrow Q + K^\top RK \\ A_2 &\leftarrow \sqrt{\gamma_2}(A+BK) & \alpha &\leftarrow 1/2 \end{aligned}$$

where A_1, A_2, Σ , and α are defined as in Perdomo et al. [2021]. Note that for c satisfying,

$$c \leq \frac{1}{8\|\sqrt{\gamma_1}(A+BK)\|_{\text{op}}} \min \left\{ \frac{1}{\|P_1\|_{\text{op}}^{3/2}}, \frac{\text{tr}[P_1]}{d_x\|P_1\|_{\text{op}}^{7/2}} \right\} + 1$$

we get that,

$$\|A_1 - A_2\|_{\text{op}}^2 = \gamma_1(c-1)^2\|A+BK\|_{\text{op}}^2 \leq \frac{1}{64\|P_1\|_{\text{op}}^3} = \frac{\alpha^2}{16\|P_1\|_{\text{op}}^3}.$$

Therefore, Proposition C.7 states that, for $\mathcal{C} := \text{tr}[P_1^{-1/2}(Q + K^\top RK)P_1^{-1/2}] \leq \text{tr}[I] = d_x$,

$$\begin{aligned} \text{tr}[P_2 - P_1] &\leq 8\mathcal{C}\sqrt{\gamma_1}(c-1)\|A+BK\|_{\text{op}}\|P_1\|_{\text{op}}^{7/2} \\ &\leq \text{tr}[P_1]. \end{aligned}$$

Lastly, noting that,

$$P_1 = \text{dlyap}(\sqrt{\gamma_1}(A+BK), Q + K^\top RK) \succeq \gamma(A+BK)^\top(A+BK)$$

we have that $\|\sqrt{\gamma_1}(A+BK)\|_{\text{op}} \leq \|P_1\|_{\text{op}}^{1/2}$ and $\text{tr}[P_1] \geq d_x$. Therefore, since $\|P_1\|_{\text{op}} \geq 1$, in order to apply Proposition C.7 from Perdomo et al. [2021] it suffices for c to satisfy,

$$c \leq \frac{1}{8\|P_1\|_{\text{op}}^4} + 1. \quad \square$$

Lemma A.5 (Lemma D.9 in Perdomo et al. [2021]). Let A be a stable matrix, $Q \succeq I$, and define $P := \text{dlyap}(A, Q)$. Then, for all $j \geq 0$,

$$(A^\top)^j A^j \preceq (A^\top)^j P A^j \preceq P \left(1 - \frac{1}{\|P\|_{\text{op}}}\right)^j.$$

Lemma A.6. Let A be a stable matrix and define $P := \text{dlyap}(A, Q)$ where $Q \succeq I$. Then, for any matrix Δ such that $\|\Delta\|_{\text{op}} \leq \frac{1}{6\|P\|_{\text{op}}^2}$, it holds that for all $j \geq 0$

$$((A + \Delta)^\top)^j P (A + \Delta)^j \preceq P \left(1 - \frac{1}{2\|P\|_{\text{op}}}\right)^j.$$

Proof. Expanding out, we have that

$$\begin{aligned} (A + \Delta)^\top P (A + \Delta) &= A^\top P A + A^\top P \Delta + \Delta^\top P A + \Delta^\top P \Delta \\ &\preceq P \left(1 - \frac{1}{\|P\|_{\text{op}}}\right) + \|A^\top P \Delta + \Delta^\top P A + \Delta^\top P \Delta\|_{\text{op}} I, \end{aligned}$$

where in the second line we have used properties of the Lyapunov function, [Lemma A.5](#). Next, we observe that

$$\|\Delta^\top P A\|_{\text{op}} \leq \|\Delta^\top P^{1/2}\|_{\text{op}} \|P^{1/2} A\|_{\text{op}} \leq \|\Delta^\top P^{1/2}\|_{\text{op}} \|P^{1/2}\|_{\text{op}} \leq \|\Delta\|_{\text{op}} \|P\|_{\text{op}},$$

where we have again used [Lemma A.5](#) to conclude that $A^\top P A \preceq P$. Note that the exact same calculation holds for $\|A^\top P \Delta\|_{\text{op}}$. Hence, we can conclude that for Δ such that $\|\Delta\|_{\text{op}} \leq 1$,

$$(A + \Delta)^\top P (A + \Delta) \preceq P \left(1 - \frac{1}{\|P\|_{\text{op}}}\right) + 3\|P\|_{\text{op}} \|\Delta\|_{\text{op}} \cdot I.$$

Using the fact that, $P \succeq I$ and that $\|\Delta\|_{\text{op}} \leq 1/(6\|P\|_{\text{op}}^2)$, we get that,

$$3\|P\|_{\text{op}} \|\Delta\|_{\text{op}} \cdot I \preceq \frac{1}{2\|P\|_{\text{op}}} P,$$

which finishes the proof. \square

B Deferred Proofs and Analysis for the Nonlinear Setting

Establishing Lyapunov functions. Our analysis for nonlinear systems begins with the observation that any state-feedback controller K which achieves finite cost on the γ -discounted LQR problem has an associated value function $P_{K,\gamma}$ which can be used as a Lyapunov function for the $\sqrt{\gamma}$ -damped nonlinear dynamics, for small enough initial states. We present the proof of this result in [Appendix B.3](#).

Lemma B.1. *Let $J_{\text{lin}}(K | \gamma) < \infty$. Then, for all $\mathbf{x} \in \mathbb{R}_x^d$ such that,*

$$\mathbf{x}^\top P_{K,\gamma} \mathbf{x} \leq \frac{r_{\text{nl}}^2}{4\beta_{\text{nl}}^2 \|P_{K,\gamma}\|_{\text{op}}^3},$$

the following inequality holds:

$$\gamma \cdot G_{\text{nl}}(\mathbf{x}, K\mathbf{x})^\top P_{K,\gamma} G_{\text{nl}}(\mathbf{x}, K\mathbf{x}) \leq \mathbf{x}^\top P_{K,\gamma} \mathbf{x} \cdot \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right).$$

Using this observation, we can then show that any controller which has finite discounted LQR cost is exponentially stabilizing over states in a sufficiently small region of attraction.

Lemma B.2. *Assume $J_{\text{lin}}(K | \gamma) < \infty$ and define $\{\mathbf{x}_{t,\text{nl}}\}_{t=0}^\infty$ be the sequence of states generated according to $\mathbf{x}_{t+1,\text{nl}} = \sqrt{\gamma} G_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})$ where $\mathbf{x}_{t,\text{nl}} = \mathbf{x}_0$. If \mathbf{x}_0 is such that,*

$$V_0 := \mathbf{x}_0^\top P_{K,\gamma} \mathbf{x}_0 \leq \frac{r_{\text{nl}}^2}{4\beta_{\text{nl}}^2 \|P_{K,\gamma}\|_{\text{op}}^3},$$

then for all $t \geq 0$ and for V_0 defined as above,

a) *The norm of the state $\|\mathbf{x}_{t,\text{nl}}\|^2$ is bounded by*

$$\|\mathbf{x}_{t,\text{nl}}\|^2 \leq \mathbf{x}_{t,\text{nl}}^\top P_{K,\gamma} \mathbf{x}_{t,\text{nl}} \leq V_0 \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right)^t. \quad (\text{B.1})$$

b) *The norms of $f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})$ and $\nabla f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})$ are bounded by*

$$\|f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})\| \leq \beta_{\text{nl}}(1 + \|K\|_{\text{op}}^2) V_0 \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right)^t. \quad (\text{B.2})$$

$$\|\nabla f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})\|_{\text{op}} \leq \beta(1 + \|K\|_{\text{op}}) V_0^{1/2} \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right)^{t/2}. \quad (\text{B.3})$$

Proof. The proof of (a) follows by repeatedly applying [Lemma B.1](#). Part (b) follows from the first after using [Lemma 3.1](#). The statement of the lemma in the main body follows from using $\|P_{K,\gamma}\|_{\text{op}} \leq \text{tr}[P_{K,\gamma}] = J_{\text{lin}}(K | \gamma)$ and simplifying $(1 - x)^t \leq \exp(-tx)$. \square

Relating G_{nl} to its Jacobian Linearization. Having established how any controller that achieves finite LQR cost is guaranteed to be stabilizing for the nonlinear system, we now go a step further and illustrate how this stability guarantee can be used to prove that the difference in costs and gradients between G_{nl} and its Jacobian linearization are guaranteed to be small.

Proposition 3.3 (restated). *Assume $J_{\text{lin}}(K | \gamma) < \infty$. Then,*

- a) *If $r \leq \frac{r_{\text{nl}}}{2\beta_{\text{nl}}\|P_{K,\gamma}\|_{\text{op}}^2}$, then $|J_{\text{nl}}(K | \gamma, r) - J_{\text{lin}}(K | \gamma)| \leq 8d_x\beta_{\text{nl}}\|P_{K,\gamma}\|_{\text{op}}^4 \cdot r$.*
- b) *If $r \leq \frac{r_{\text{nl}}}{12\beta_{\text{nl}}\|P_{K,\gamma}\|_{\text{op}}^{5/2}}$, then,*

$$\|\nabla_K J_{\text{nl}}(K | \gamma, r) - \nabla_K J_{\text{lin}}(K | \gamma)\|_{\text{F}} \leq 48d_x\beta_{\text{nl}}(1 + \|B\|_{\text{op}})\|P_{K,\gamma}\|_{\text{op}}^7 r.$$

Proof. Due to our assumption on $r = \|\mathbf{x}_0\|$, we have that,

$$\mathbf{x}_0^\top P_{K,\gamma} \mathbf{x}_0 \leq \|P_{K,\gamma}\|_{\text{op}} \|\mathbf{x}_0\|^2 \leq \frac{1}{4\beta_{\text{nl}}\|P_{K,\gamma}\|_{\text{op}}^3}.$$

Therefore, we can apply [Lemma B.5](#) to conclude that,

$$|J_{\text{nl}}(K | \gamma, \mathbf{x}_0) - J_{\text{lin}}(K | \gamma, \mathbf{x}_0)| \leq 8\beta_{\text{nl}}\|\mathbf{x}_0\|^3\|P_{K,\gamma}\|_{\text{op}}^4.$$

Next, we multiply both sides by d_x/r^2 , take expectations, and apply Jensen's inequality to get that,

$$\left| \frac{d_x}{r^2} \mathbb{E}_{\mathbf{x}_0 \sim r \cdot \mathcal{S}^{d_x-1}} J_{\text{nl}}(K | \gamma, \mathbf{x}_0) - \frac{d_x}{r^2} \mathbb{E}_{\mathbf{x}_0 \sim r \cdot \mathcal{S}^{d_x-1}} J_{\text{lin}}(K | \gamma, \mathbf{x}_0) \right| \leq 8 \frac{d_x}{r^2} \beta_{\text{nl}} \|P_{K,\gamma}\|_{\text{op}}^4 \mathbb{E}_{\mathbf{x}_0 \sim r \cdot \mathcal{S}^{d_x-1}} \|\mathbf{x}_0\|^3.$$

Given our definitions of the linear objective in [Definition 2.1](#), we have that,

$$\frac{d_x}{r^2} \mathbb{E}_{\mathbf{x}_0 \sim r \cdot \mathcal{S}^{d_x-1}} J_{\text{lin}}(K | \gamma, \mathbf{x}_0) = J_{\text{lin}}(K | \gamma),$$

for all $r > 0$. Therefore, we can rewrite the inequality above as,

$$|J_{\text{nl}}(K | \gamma, r) - J_{\text{lin}}(K | \gamma)| \leq 8d_x\beta_{\text{nl}}\|P_{K,\gamma}\|_{\text{op}}^4 \cdot r.$$

The second part of the proposition uses the same argument as part a, but this time employing [Lemma B.6](#) to bound the difference in gradients (pointwise). \square

In short, this previous lemma states that if the cost on the linear system is bounded, then the costs and gradients between the nonlinear objective and its Jacobian linearization are close. We can also prove the analogous statement which establishes closeness while assuming that the cost on the nonlinear system is bounded.

Lemma B.3. *Let $\alpha > 1$ be such that $80d_x^2 J_{\text{nl}}(K | \gamma, r) \leq \alpha$.*

- 1. *If $r \leq \frac{r_{\text{nl}}^2}{64\beta_{\text{nl}}\alpha^2(1+\|K\|_{\text{op}})}$, then $|J_{\text{nl}}(K | \gamma, r) - J_{\text{lin}}(K | \gamma)| \leq 8d_x\beta_{\text{nl}}\alpha^4 r$.*
- 2. *If $r \leq \frac{1}{12\beta_{\text{nl}}\alpha^{5/2}}$, then $\|\nabla_K J_{\text{nl}}(K | \gamma, r) - \nabla_K J_{\text{lin}}(K | \gamma)\|_{\text{F}} \leq 48d_x\beta_{\text{nl}}(1+\|B\|_{\text{op}})\alpha^7 r$.*

Proof. The lemma is a consequence of combining [Proposition 3.3](#) and [Proposition C.5](#). In particular, from [Proposition C.5](#) if $r \leq \min\{\alpha r_{\text{nl}}^2, \frac{d_x}{64\alpha^2\beta_{\text{nl}}(1+\|K\|_{\text{op}})}\}$, then

$$80d_x^2 J_{\text{nl}}(K | \gamma, r) \geq \min\{J_{\text{lin}}(K | \gamma), \alpha\}$$

However, since $\alpha \geq 80d_x^2 J_{\text{nl}}(K | \gamma, r)$, we conclude that $80d_x^2 J_{\text{nl}}(K | \gamma, r) \geq J_{\text{lin}}(K | \gamma)$. Having shown that the linear cost is bounded, we can now plug in [Proposition 3.3](#). In particular, if

$$r \leq \frac{r_{\text{nl}}}{2\beta_{\text{nl}}\alpha^2} \leq \frac{r_{\text{nl}}}{2\beta_{\text{nl}}J_{\text{lin}}(K | \gamma)^2}$$

then, [Proposition 3.3](#) states that

$$|J_{\text{nl}}(K | \gamma, r) - J_{\text{lin}}(K | \gamma)| \leq 8d_x\beta_{\text{nl}}J_{\text{lin}}(K | \gamma)^4 r \leq 8d_x\beta_{\text{nl}}\alpha^4 r.$$

To prove the second part of the statement, we again use [Proposition 3.3](#). In particular, since

$$r \leq \frac{1}{12\beta_{\text{nl}}\alpha^{5/2}} \leq \frac{1}{12\beta_{\text{nl}}J_{\text{lin}}(K | \gamma)^{5/2}}$$

we can hence conclude that

$$\|\nabla_K J_{\text{nl}}(K | \gamma, r) - \nabla_K J_{\text{lin}}(K | \gamma)\|_{\text{F}} \leq 48d_x\beta_{\text{nl}}(1 + \|B\|_{\text{op}})\alpha^7 \cdot r.$$

\square

B.1 Discount annealing on nonlinear systems: Proof of Theorem 2

As in Theorem 1, we first prove parts c) and d) by induction and then prove parts a) and b) separately.

B.1.1 Proof of c) and d)

Base case. As before, at each iteration t of discount annealing, policy gradients is initialized at $K_{t,0} := K_t$ and computes updates according to,

$$K_{t,j+1} = K_{t,j} - \eta \cdot \varepsilon\text{-Grad}(J_{\text{nl}}(\cdot | \gamma_t, r_\star), K_{t,j}).$$

To prove correctness, we show that the noisy gradients on the nonlinear system are close to the true gradients on the *linear* system. That is,

$$\|\varepsilon\text{-Grad}(J_{\text{nl}}(\cdot | \gamma_t, r_\star), K_{t,j}) - \nabla J_{\text{lin}}(K_{t,j} | \gamma_t)\|_{\text{F}} \leq \mathcal{C}_{\text{pg},t} d_x \quad \forall j \geq 0, \quad (\text{B.4})$$

where $\mathcal{C}_{\text{pg},t} = \text{poly}(1/\|A\|_{\text{op}}, 1/\|B\|_{\text{op}}, 1/J_{\text{lin}}(K_t | \gamma_t))$ is again a fixed polynomial from Lemma A.3.

Consider the first iteration of discount annealing, by choice of γ_0 , we have that $J_{\text{lin}}(K_0 | \gamma_0) < \infty$. Therefore, by Proposition 3.3 if

$$r_\star \leq \min \left\{ \frac{r_{\text{nl}}}{12\beta_{\text{nl}}J_{\text{lin}}(K_{0,0} | \gamma_0)}, \frac{\mathcal{C}_{\text{pg},0}}{100\beta_{\text{nl}}(1 + \|B\|_{\text{op}})J_{\text{lin}}(K_{0,0} | \gamma_0)^7} \right\}$$

it must hold that $\|\nabla J_{\text{nl}}(K_{0,0} | \gamma_0, r_\star) - \nabla J_{\text{lin}}(K_{0,0} | \gamma_0)\|_{\text{F}} \leq .5\mathcal{C}_{\text{pg},0}d_x$. Likewise, if we choose the tolerance parameter $\varepsilon \leq .5\mathcal{C}_{\text{pg},0}d_x$ in $\varepsilon\text{-Grad}$ then we have that

$$\|\varepsilon\text{-Grad}(J_{\text{nl}}(\cdot | \gamma_0, r_\star), K_{0,0}) - J_{\text{nl}}(K_{0,0} | \gamma_0, r_\star)\|_{\text{F}} \leq .5\mathcal{C}_{\text{pg},0}d_x.$$

By the triangle inequality, the inequality in Eq. (B.4) holds for $t = 0$ and $j = 0$. However, because Lemma A.3 shows that policy gradients is a descent method, that is $J_{\text{lin}}(K_{0,j} | \gamma_0) \leq J_{\text{lin}}(K_{0,0} | \gamma_0)$ for all $j \geq 0$, Eq. (B.4) also holds for all $j \geq 0$ for the same choice of r_\star and tolerance parameter for $\varepsilon\text{-Grad}$. By guarantee of Lemma A.3, for $t = 0$, policy gradients achieves the guarantee outlined in Step 2 using at most $\text{poly}(\|A\|_{\text{op}}, \|B\|_{\text{op}}, J_{\text{lin}}(K_0 | \gamma_0))$ many queries.

To prove that random search achieves the guarantee outlines in Step 4 at iteration 0 of discount annealing, we appeal to Lemma C.7. In particular, we instantiate the lemma with $f \leftarrow J_{\text{nl}}(K_1 | \cdot, r_\star)$, $f_1 \leftarrow 8J_{\text{nl}}(K_1 | \gamma_0, r_\star)$, $f_2 \leftarrow 2.5J_{\text{nl}}(K_1 | \gamma_0, r_\star)$. As before, the algorithm requires values $\bar{f}_1 \in [2.9J_{\text{nl}}(K_1 | \gamma_0, r_\star), 3J_{\text{nl}}(K_1 | \gamma_0)]$ and $\bar{f}_2 \in [6J_{\text{nl}}(K_1 | \gamma_0, r_\star), 6.1J_{\text{nl}}(K_1 | \gamma_0)]$. These can be estimated via two calls to $\varepsilon\text{-Eval}$ with tolerance parameter $.01d_x$.

To show the lemma applies we only need to lower bound the width of feasible γ such that

$$2.9J_{\text{nl}}(K_1 | \gamma_0, r_\star) \leq J_{\text{nl}}(K_1 | \gamma, r_\star) \leq 6.1J_{\text{nl}}(K_1 | \gamma_0, r_\star) \quad (\text{B.5})$$

From the guarantee from policy gradients, we know that $J_{\text{lin}}(K_1 | \gamma_0) \leq 2\text{tr}[P_\star]$. Furthermore, from the proof of Theorem 1, we know that there exists $\gamma'', \gamma' \in [0, 1]$ satisfying, $\gamma'' - \gamma' \geq \frac{1}{5200\text{tr}[P_\star]^4} \gamma_0$, such that for all $\gamma \in [\gamma', \gamma'']$

$$3J_{\text{lin}}(K_1 | \gamma_0) \leq J_{\text{lin}}(K_1 | \gamma) \leq 6J_{\text{lin}}(K_1 | \gamma_0). \quad (\text{B.6})$$

To finish the proof of correctness, we show that any γ that satisfies Eq. (B.6) must also satisfy Eq. (B.5). In particular, since $J_{\text{lin}}(K_1 | \gamma_0) \leq 2\text{tr}[P_\star]$ and $J_{\text{lin}}(K_1 | \gamma) \leq 12\text{tr}[P_\star]$ for

$$r_\star \leq \min \left\{ \frac{r_{\text{nl}}}{2\beta_{\text{nl}}(12\text{tr}[P_\star])^2}, \frac{.01}{8\beta_{\text{nl}}(12\text{tr}[P_\star])^4} \right\},$$

it holds that $|J_{\text{nl}}(K_1 | \gamma_0, r_\star) - J_{\text{lin}}(K_1 | \gamma_0)| \leq .01d_x$ and $|J_{\text{nl}}(K_1 | \gamma, r_\star) - J_{\text{lin}}(K_1 | \gamma)| \leq .01d_x$. Using these two inequalities along with Eq. (B.6) implies that Eq. (B.5) must also hold. Therefore, the width of the feasible region is at least $1/(5200\text{tr}[P_\star]^4)$ and random search must return a discount factor using at most 1 over this many iterations by Lemma C.7.

Inductive step. To show that policy gradients converges, we can use the exact same argument as when arguing the base case where instead of referring to $J_{\text{lin}}(K_0 | \gamma_0)$ and $J_{\text{nl}}(K_0 | \gamma_0, r_*)$ we use $J_{\text{lin}}(K_t | \gamma_t)$ and $J_{\text{nl}}(K_t | \gamma_t, r_*)$. In particular, we can reuse the same argument as in the base case, but need to ensure that is that $\sup_{t \geq 1} J_{\text{lin}}(K_t | \gamma_t)$ is uniformly bounded.

To prove this, from the inductive hypothesis on the correctness of binary search at previous iterations, we know that $J_{\text{nl}}(K_t | \gamma_t) \leq 8J_{\text{nl}}(K_t | \gamma_{t-1}, r_*)$. Again by the inductive hypothesis, at time step $t - 1$ policy gradients achieves the desired guarantee from Step 2, implying that $J_{\text{lin}}(K_t | \gamma_{t-1}) \leq 2\text{tr}[P_*]$. By choice of r_* this implies that

$$|J_{\text{lin}}(K_t | \gamma_{t-1}) - J_{\text{nl}}(K_t | \gamma_{t-1}, r_*)| \leq .01d_x$$

and hence $J_{\text{nl}}(K_t | \gamma_t) \leq 20\text{tr}[P_*]$. Now, we can apply [Lemma B.3](#) to conclude that for $\alpha := (80 \times 20)d_x^2 \text{tr}[P_*]$ and

$$r_* \leq \min \left\{ \frac{r_{\text{nl}}^2}{64\beta_{\text{nl}}\alpha^2(1 + \|K\|_{\text{op}})}, \frac{.01}{8\beta_{\text{nl}}\alpha^4} \right\}$$

it holds that $|J_{\text{nl}}(K_t | \gamma_t, r_*) - J_{\text{lin}}(K_t | \gamma_t)| \leq .01d_x$ and hence $J_{\text{lin}}(K_t | \gamma_t) \leq 21\text{tr}[P_*]$. Therefore, $\sup_{t \geq 1} J_{\text{lin}}(K_t | \gamma_t) \leq 21\text{tr}[P_*]$.

Similarly, the inductive step for the random search procedure follows from noting that the exact same argument can be repeated by replacing $J_{\text{nl}}(K_1 | \gamma_0, r_*)$ with $J_{\text{nl}}(K_t | \gamma_{t-1})$ and $J_{\text{lin}}(K_1 | \gamma_0)$ with $J_{\text{lin}}(K_t | \gamma_{t-1})$ since (by the inductive hypothesis) $J_{\text{lin}}(K_t | \gamma_{t-1}) \leq 2\text{tr}[P_*]$.

B.1.2 Proof of a)

By the guarantee from Step 2, the algorithm returns a \widehat{K} which satisfies the following guarantee on the *linear* system:

$$J_{\text{lin}}(\widehat{K} | 1) - \min_K J_{\text{lin}}(K | 1) \leq d_x.$$

Therefore, $J_{\text{lin}}(\widehat{K} | 1) \leq 2\text{tr}[P_*]$. Now by [Lemma B.2](#), the following holds,

$$\begin{aligned} \|\mathbf{x}_{t,\text{nl}}\|^2 &\leq \|\widehat{P}\|_{\text{op}} \|\mathbf{x}_0\|^2 \left(1 - \frac{1}{2\|\widehat{P}\|_{\text{op}}}\right)^t \\ &\leq 2\text{tr}[P_*] \|\mathbf{x}_0\|^2 \exp\left(-\frac{t}{4\text{tr}[P_*]}\right), \end{aligned}$$

for $\widehat{P} := \text{dlyap}(A + B\widehat{K}, Q + \widehat{K}^\top R\widehat{K})$ and all \mathbf{x}_0 such that $\|\mathbf{x}_0\| \leq r_{\text{nl}}/(8\beta_{\text{nl}}\text{tr}[P_*]^2)$.

B.1.3 Proof of b)

The bound for the number of subproblems solved by the discount annealing algorithms is similar to that of the linear case. The crux of the argument for part b) is to show that any $\gamma' \in [\gamma_t, 1]$ such that

$$2.5J_{\text{nl}}(K_{t+1} | \gamma_t, r_*) \leq J_{\text{nl}}(K_{t+1} | \gamma', r_*) \leq 8J_{\text{nl}}(K_{t+1} | \gamma_t, r_*)$$

the following inequality must also hold: $J_{\text{lin}}(K_{t+1} | \gamma') \geq 2J_{\text{lin}}(K_{t+1} | \gamma_t)$. Once we've lower bounded the cost on the *linear* system, we can repeat the same argument as in Theorem 1. Since the cost on the linear system in nondecreasing in γ , it must be the case that γ' satisfies

$$\gamma' \geq \left(\frac{1}{8\|P_{K_{t+1}, \gamma_t}\|_{\text{op}}^4} + 1\right)^2 \gamma_t \geq \left(\frac{1}{128\text{tr}[P_*]^4} + 1\right)^2 \gamma_t.$$

Here, we have again we have used the calculation that,

$$\|P_{K_{t+1}, \gamma_t}\|_{\text{op}}^4 \leq \text{tr}[P_{K_{t+1}, \gamma_t}]^4 = J_{\text{lin}}(K_{t+1} | \gamma_t)^4 \leq (\min_K J_{\text{lin}}(K | \gamma_t) + d_x)^4 \leq 16\text{tr}[P_*]^4,$$

which follows from the guarantee that (for our choice of r_*) policy gradients on the nonlinear system converges to a near optimal controller for the system's Jacobian linearization. Hence, as in the linear

setting, we conclude that $\gamma_t \geq (1/(128\text{tr}[P_\star]^4) + 1)^{2t}\gamma_0$ and discount annealing achieves the same rate as for linear systems.

We now focus on establishing that $J_{\text{lin}}(K_{t+1} | \gamma') \geq 2J_{\text{lin}}(K_{t+1} | \gamma_t)$. By guarantee from policy gradients, we have that $J_{\text{lin}}(K_{t+1} | \gamma_t) \leq \min_K J_{\text{lin}}(K | \gamma_t) + d_x \leq 2\text{tr}[P_\star]$. Therefore, by [Proposition 3.3](#) since

$$r_\star \leq \frac{.01r_{\text{nl}}}{8 \times 2^4 \cdot \text{tr}[P_\star]^4 \beta_{\text{nl}}} \leq \min \left\{ \frac{r_{\text{nl}}}{2\beta_{\text{nl}}\|P_{K_{t+1}, \gamma_t}\|_{\text{op}}^2}, \frac{.01}{8\beta_{\text{nl}}\|P_{K_{t+1}, \gamma_t}\|_{\text{op}}^4} \right\}$$

it holds that $|J_{\text{lin}}(K_{t+1} | \gamma_t) - J_{\text{nl}}(K_{t+1} | \gamma_t, r_\star)| \leq .01d_x$.

Next, we show that $|J_{\text{lin}}(K_{t+1} | \gamma') - J_{\text{nl}}(K_{t+1} | \gamma', r_\star)|$ is also small. In particular, the previous statement, together with the guarantee from Step 4, implies that

$$J_{\text{nl}}(K_{t+1} | \gamma') \leq 8J_{\text{nl}}(K_{t+1} | \gamma_t) \leq 8(J_{\text{lin}}(K_{t+1} | \gamma_t) + .01d_x) \leq 8.08J_{\text{lin}}(K_{t+1} | \gamma_t) \leq 16.16\text{tr}[P_\star].$$

Therefore, for $\alpha := (80 \times 17)\text{tr}[P_\star]d_x^2$, [Lemma B.3](#) implies that if,

$$r \leq \frac{r_{\text{nl}}^2}{64\beta_{\text{nl}}\alpha^2(1 + \|K_{t+1}\|_{\text{op}})},$$

it holds that $|J_{\text{nl}}(K_{t+1} | \gamma', r) - J_{\text{lin}}(K_{t+1} | \gamma')| \leq 8d_x\beta_{\text{nl}}\alpha^4 r$. Hence, if $r \leq .01/(8\beta_{\text{nl}}\alpha^4)$ we get that $|J_{\text{lin}}(K_{t+1} | \gamma') - J_{\text{nl}}(K_{t+1} | \gamma', r_\star)| \leq .01d_x$. Using again the fact that $\min\{J_{\text{lin}}(K | \gamma), J_{\text{nl}}(K | \gamma, r)\} \geq d_x$ for all K, γ, r we hence conclude that

$$\begin{aligned} J_{\text{lin}}(K_{t+1} | \gamma') &\geq .99J_{\text{nl}}(K_{t+1} | \gamma', r_\star) \\ &\geq 2.5 \times .99 \cdot J_{\text{nl}}(K_{t+1} | \gamma_t, r_\star) \\ &\geq 2.5 \times .99^2 \cdot J_{\text{lin}}(K_{t+1} | \gamma_t), \end{aligned}$$

which finished the proof of the fact that $J_{\text{lin}}(K_{t+1} | \gamma') \geq 2J_{\text{lin}}(K_{t+1} | \gamma_t)$.

B.2 Relating costs and gradients to the linear system: Proof of [Proposition 3.3](#)

In order to relate the properties of the nonlinear system to its Jacobian linearization, we employ the following version of the performance difference lemma.

Lemma B.4 (Performance Difference). *Assume $J_{\text{lin}}(K | \gamma) < \infty$ and define $\{\mathbf{x}_{t,\text{nl}}\}_{t=0}^\infty$ be the sequence of states generated according to $\mathbf{x}_{t+1,\text{nl}} = \sqrt{\gamma}G_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})$ where $\mathbf{x}_{t,\text{nl}} = \mathbf{x}_0$. Then,*

$$J_{\text{nl}}(K | \gamma, \mathbf{x}) - J_{\text{lin}}(K | \gamma, \mathbf{x}) = \sum_{t=0}^{\infty} \gamma \cdot f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})^\top P_{K,\gamma} (G_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}) + G_{\text{lin}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})).$$

Proof. From the definition of the relevant objectives, and [Fact A.2](#), we get that,

$$\begin{aligned} J_{\text{nl}}(K | \gamma, \mathbf{x}) - J_{\text{lin}}(K | \gamma, \mathbf{x}) &= \left(\sum_{t=0}^{\infty} \mathbf{x}_{t,\text{nl}}^\top (Q + K^\top RK) \mathbf{x}_{t,\text{nl}} \right) - \mathbf{x}^\top P_{K,\gamma} \mathbf{x} \\ &= \left(\sum_{t=0}^{\infty} \mathbf{x}_{t,\text{nl}}^\top (Q + K^\top RK) \mathbf{x}_{t,\text{nl}} \pm \mathbf{x}_{t,\text{nl}}^\top P_{K,\gamma} \mathbf{x}_{t,\text{nl}} \right) - \mathbf{x}^\top P_{K,\gamma} \mathbf{x} \\ &= \sum_{t=0}^{\infty} \mathbf{x}_{t,\text{nl}}^\top (Q + K^\top RK) \mathbf{x}_{t,\text{nl}} + \mathbf{x}_{t+1,\text{nl}}^\top P_{K,\gamma} \mathbf{x}_{t+1,\text{nl}} - \mathbf{x}_{t,\text{nl}}^\top P_{K,\gamma} \mathbf{x}_{t,\text{nl}}, \end{aligned} \tag{B.7}$$

where in the last line we have used the fact that $\mathbf{x}_{0,\text{nl}} = \mathbf{x}$. The proof then follows from the following two observations. First, by definition of state sequence, $\mathbf{x}_{t,\text{nl}}$,

$$\mathbf{x}_{t+1,\text{nl}} = \sqrt{\gamma} \cdot G_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}).$$

Second, since $P_{K,\gamma}$ is the solution to a Lyapunov equation,

$$\begin{aligned} \mathbf{x}_{t,\text{nl}}^\top P_{K,\gamma} \mathbf{x}_{t,\text{nl}} &= \mathbf{x}_{t,\text{nl}}^\top (Q + K^\top RK) \mathbf{x}_{t,\text{nl}} + \gamma \cdot \mathbf{x}_{t,\text{nl}}^\top (A + BK)^\top P_{K,\gamma} (A + BK) \mathbf{x}_{t,\text{nl}} \\ &= \mathbf{x}_{t,\text{nl}}^\top (Q + K^\top RK) \mathbf{x}_{t,\text{nl}} + \gamma \cdot G_{\text{lin}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})^\top P_{K,\gamma} G_{\text{lin}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}). \end{aligned}$$

Plugging these last two lines into [Eq. \(B.7\)](#), we get that $J_{\text{nl}}(K \mid \gamma, \mathbf{x}) - J_{\text{lin}}(K \mid \gamma, \mathbf{x})$ is equal to,

$$\begin{aligned}
&= \sum_{t=0}^{\infty} \gamma \cdot G_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})^\top P_{K,\gamma} G_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}) - \gamma \cdot G_{\text{lin}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})^\top P_{K,\gamma} G_{\text{lin}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}) \\
&= \sum_{t=0}^{\infty} \gamma \cdot (G_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}) - G_{\text{lin}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})) P_{K,\gamma} (G_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}) + G_{\text{lin}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})) \\
&= \sum_{t=0}^{\infty} \gamma \cdot f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}) P_{K,\gamma} (G_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}) + G_{\text{lin}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})).
\end{aligned}$$

□

B.2.1 Establishing similarity of costs

The following lemma follows by bounding the terms appearing in the performance difference lemma.

Lemma B.5 (Similarity of Costs). *Assume $J_{\text{lin}}(K \mid \gamma) < \infty$ and define $\{\mathbf{x}_{t,\text{nl}}\}_{t=0}^{\infty}$ be the sequence of states generated according to $\mathbf{x}_{t+1,\text{nl}} = \sqrt{\gamma} G_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})$ where $\mathbf{x}_{t,\text{nl}} = \mathbf{x}_0$. For \mathbf{x}_0 such that,*

$$\mathbf{x}_0^\top P_{K,\gamma} \mathbf{x}_0 \leq \frac{r_{\text{nl}}^2}{4\beta_{\text{nl}}^2 \|P_{K,\gamma}\|_{\text{op}}^3},$$

then,

$$|J_{\text{nl}}(K \mid \gamma, \mathbf{x}_0) - J_{\text{lin}}(K \mid \gamma, \mathbf{x}_0)| \leq 8\beta_{\text{nl}} \|\mathbf{x}_0\|^3 \|P_{K,\gamma}\|_{\text{op}}^4.$$

Proof. We begin with the following observation. Due to our assumption on \mathbf{x}_0 , we can use [Lemma B.2](#) to conclude that for all $t \geq 0$, the following relationship holds for $V_0 := \mathbf{x}_0^\top P_{K,\gamma} \mathbf{x}_0$,

$$\|\mathbf{x}_{t,\text{nl}}\|^2 \leq \mathbf{x}_{t,\text{nl}}^\top P_{K,\gamma} \mathbf{x}_{t,\text{nl}} \leq V_0 \cdot \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right)^t. \quad (\text{B.8})$$

Now, from the performance difference lemma ([Lemma B.4](#)), we get that $J_{\text{nl}}(K \mid \gamma, \mathbf{x}) - J_{\text{lin}}(K \mid \gamma, \mathbf{x})$ is equal to:

$$\sum_{t=0}^{\infty} \gamma \cdot f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})^\top P_{K,\gamma} (G_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}) + G_{\text{lin}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})).$$

Therefore, the difference $|J_{\text{nl}}(K \mid \gamma, \mathbf{x}_0) - J_{\text{lin}}(K \mid \gamma, \mathbf{x}_0)|$ can be bounded by,

$$\sum_{t=0}^{\infty} \underbrace{\|\sqrt{\gamma} \cdot P_{K,\gamma}^{1/2} f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})\|}_{:=T_1} \cdot \underbrace{\|\sqrt{\gamma} \cdot P_{K,\gamma}^{1/2} (G_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}) + G_{\text{lin}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}))\|}_{:=T_2}. \quad (\text{B.9})$$

Now, we analyze each of T_1 and T_2 separately. For T_1 , by [Lemma 3.1](#), and the fact that $\gamma \leq 1$,

$$\begin{aligned}
\|\sqrt{\gamma} \cdot P_{K,\gamma}^{1/2} f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})\| &\leq \|P_{K,\gamma}\|_{\text{op}}^{1/2} \|f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})\| \\
&= \|P_{K,\gamma}\|_{\text{op}}^{1/2} \cdot \beta_{\text{nl}} (\|\mathbf{x}_{t,\text{nl}}\|^2 + \|K\mathbf{x}_{t,\text{nl}}\|^2) \\
&\leq \|P_{K,\gamma}\|_{\text{op}}^{1/2} \beta_{\text{nl}} \cdot (\|K\|_{\text{op}}^2 + 1) V_0 \cdot \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right)^t,
\end{aligned}$$

where in the last line, we have used our assumption on the initial state and [Lemma B.2](#). Moving onto T_2 , we use the triangle inequality to get that

$$T_2 \leq \|\sqrt{\gamma} \cdot P_{K,\gamma}^{1/2} G_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})\| + \|\sqrt{\gamma} \cdot P_{K,\gamma}^{1/2} G_{\text{lin}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})\|.$$

For the second term above, by [Lemma A.5](#) and [Lemma B.2](#), we have that

$$\begin{aligned} \|\sqrt{\gamma} \cdot P_{K,\gamma}^{1/2} G_{\text{lin}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})\| &= \sqrt{\gamma \cdot \mathbf{x}_{t,\text{nl}}^\top (A + BK)^\top P_{K,\gamma} (A + BK) \mathbf{x}_{t,\text{nl}}} \\ &\leq \sqrt{\mathbf{x}_{t,\text{nl}}^\top P_{K,\gamma} \mathbf{x}_{t,\text{nl}}} \\ &\leq V_0^{1/2} \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right)^{t/2}. \end{aligned}$$

Lastly, we bound the first term by again using [Lemma B.2](#),

$$\|\sqrt{\gamma} \cdot P_{K,\gamma}^{1/2} G_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})\| = \|P_{K,\gamma}^{1/2} \mathbf{x}_{t+1,\text{nl}}\| \leq V_0^{1/2} \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right)^{\frac{t+1}{2}}.$$

Therefore, T_2 is bounded by $2V_0^{1/2}(1 - 1/(2\|P_{K,\gamma}\|_{\text{op}}))^{t/2}$. Going back to [Eq. \(B.9\)](#), we can combine our bounds on T_1 and T_2 to conclude that,

$$\begin{aligned} J_{\text{nl}}(K | \gamma, \mathbf{x}) - J_{\text{lin}}(K | \gamma, \mathbf{x}) &\leq V_0^{3/2} \cdot 2\beta_{\text{nl}} \|P_{K,\gamma}\|_{\text{op}}^{1/2} (\|K\|_{\text{op}}^2 + 1) \cdot \sum_{t=0}^{\infty} \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right)^t \\ &= V_0^{3/2} \cdot 4\beta_{\text{nl}} \|P_{K,\gamma}\|_{\text{op}}^{3/2} (\|K\|_{\text{op}}^2 + 1). \end{aligned}$$

Using the fact that $1 + \|K\|_{\text{op}}^2 \leq 2\|P_{K,\gamma}\|_{\text{op}}$ and $V_0 \leq \|\mathbf{x}_0\|^2 \|P_{K,\gamma}\|_{\text{op}}$, we get that

$$V_0^{3/2} \cdot 4\beta_{\text{nl}} \|P_{K,\gamma}\|_{\text{op}}^{3/2} (\|K\|_{\text{op}}^2 + 1) \leq 8\beta_{\text{nl}} \|\mathbf{x}_0\|^3 \|P_{K,\gamma}\|_{\text{op}}^4.$$

□

B.2.2 Establishing similarity of gradients

Much like the previous lemma which bounds the costs between the linear and nonlinear system via the performance difference lemma, this next lemma differentiates the performance difference lemma to bound the difference between gradients.

Lemma B.6. *Assume $J_{\text{lin}}(K | \gamma) < \infty$. If \mathbf{x}_0 is such that*

$$\mathbf{x}_0 P_{K,\gamma} \mathbf{x}_0 \leq \frac{r_{\text{nl}}^2}{144\beta_{\text{nl}}^2 \|P_{K,\gamma}\|_{\text{op}}^4}$$

then,

$$\|\nabla_K J_{\text{nl}}(K | \gamma, \mathbf{x}_0) - \nabla_K J_{\text{lin}}(K | \gamma, \mathbf{x}_0)\|_{\text{F}} \leq 48\beta_{\text{nl}}(1 + \|B\|_{\text{op}}) \|P_{K,\gamma}\|_{\text{op}}^7 \|\mathbf{x}_0\|^3.$$

Proof. Using the variational definition of the Frobenius norm,

$$\begin{aligned} \|\nabla_K J_{\text{nl}}(K | \gamma, \mathbf{x}) - \nabla_K J_{\text{lin}}(K | \gamma, \mathbf{x})\|_{\text{F}} &= \sup_{\|\Delta\|_{\text{F}} \leq 1} \text{tr} [(\nabla_K J_{\text{nl}}(K | \gamma, \mathbf{x}) - \nabla_K J_{\text{lin}}(K | \gamma, \mathbf{x}))^\top \Delta] \\ &= \sup_{\|\Delta\|_{\text{F}} \leq 1} D_K J_{\text{nl}}(K | \gamma, \mathbf{x}) - D_K J_{\text{lin}}(K | \gamma, \mathbf{x}), \end{aligned}$$

where D_K is the directional derivative operator in the direction Δ . The argument follows by bounding the directional derivative appearing above. From the performance difference lemma, [Lemma B.4](#), we have that

$$D_K J_{\text{nl}}(K | \gamma, \mathbf{x}) - D_K J_{\text{lin}}(K | \gamma, \mathbf{x}) = \sum_{t=0}^{\infty} \gamma \cdot D_K [\cdot f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})^\top P_{K,\gamma} \mathbf{x}_{t+1,\text{nl}}].$$

Each term appearing in the sum above can be decomposed into the following three terms,

$$\begin{aligned} &\underbrace{\gamma \cdot (D_K \cdot f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}))^\top P_{K,\gamma} \mathbf{x}_{t+1,\text{nl}}}_{:=T_1} \\ &+ \underbrace{\gamma \cdot f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}) (D_K P_{K,\gamma}) \mathbf{x}_{t+1,\text{nl}}}_{:=T_2} \\ &+ \underbrace{\gamma \cdot f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}) P_{K,\gamma} (D_K \mathbf{x}_{t+1,\text{nl}})}_{:=T_3}. \end{aligned}$$

In order to bound each of these three terms, we start by bounding the directional derivatives appearing above. Throughout the remainder of the proof, we will make repeated use of the following inequalities which follow from [Lemma 3.1](#), [Lemma B.2](#), and our assumption on \mathbf{x}_0 . For all $t \geq 0$,

$$\|\mathbf{x}_{t,\text{nl}}\|^2 \leq V_0 \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right)^t \quad (\text{B.10})$$

$$\|\mathbf{x}_{t,\text{nl}}\| + \|K\mathbf{x}_{t,\text{nl}}\| \leq r_{\text{nl}}. \quad (\text{B.11})$$

Lemma B.7 (Bounding $D_K \mathbf{x}_{t,\text{nl}}$). *Let $\{\mathbf{x}_{t,\text{nl}}\}$ be the sequence of states generated according to $\mathbf{x}_{t+1,\text{nl}} = \sqrt{\gamma}G_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})$. Under the same assumptions as in [Lemma B.6](#), for all $t \geq 0$*

$$\|D_K \mathbf{x}_{t,\text{nl}}\| \leq t \cdot \|P_{K,\gamma}\|_{\text{op}}^{1/2} \left(\frac{1}{6\|P_{K,\gamma}\|_{\text{op}}^2} + \|B\|_{\text{op}} \right) \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right)^{\frac{t-1}{2}} V_0^{1/2}.$$

Proof. Taking derivatives, we get that

$$D_K f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}) = \nabla_{\mathbf{x},\mathbf{u}} f_{\text{nl}}(\mathbf{x}, \mathbf{u}) \Big|_{(\mathbf{x},\mathbf{u})=(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})} \begin{bmatrix} D_K \mathbf{x}_{t,\text{nl}} \\ D_K(K\mathbf{x}_{t,\text{nl}}) \end{bmatrix}.$$

Since, $D_K(K\mathbf{x}_{t,\text{nl}}) = \Delta \mathbf{x}_{t,\text{nl}} + K(D_K \mathbf{x}_{t,\text{nl}})$, we can rewrite the expression above as,

$$\begin{aligned} D_K f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}) &= \nabla_{\mathbf{x},\mathbf{u}} f_{\text{nl}}(\mathbf{x}, \mathbf{u}) \Big|_{(\mathbf{x},\mathbf{u})=(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})} \begin{bmatrix} 0 \\ \Delta \end{bmatrix} \mathbf{x}_{t,\text{nl}} \\ &\quad + \nabla_{\mathbf{x},\mathbf{u}} f_{\text{nl}}(\mathbf{x}, \mathbf{u}) \Big|_{(\mathbf{x},\mathbf{u})=(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})} \begin{bmatrix} I \\ K \end{bmatrix} D_K \mathbf{x}_{t,\text{nl}}. \end{aligned} \quad (\text{B.12})$$

Next, we compute $D_K \mathbf{x}_{t,\text{nl}}$,

$$\begin{aligned} D_K \mathbf{x}_{t,\text{nl}} &= \sqrt{\gamma} \cdot (D_K f_{\text{nl}}(\mathbf{x}_{t-1,\text{nl}}, K\mathbf{x}_{t-1,\text{nl}}) + D_K[(A + BK)\mathbf{x}_{t-1,\text{nl}}]) \\ &= \sqrt{\gamma} \cdot (D_K f_{\text{nl}}(\mathbf{x}_{t-1,\text{nl}}, K\mathbf{x}_{t-1,\text{nl}}) + B\Delta \mathbf{x}_{t-1,\text{nl}} + (A + BK)D_K \mathbf{x}_{t-1,\text{nl}}). \end{aligned}$$

Plugging in our earlier calculation for $D_K f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})$, we get that the following recursion holds for $\psi_t := D_K \mathbf{x}_{t,\text{nl}}$,

$$\psi_t := D_K \mathbf{x}_{t,\text{nl}} = \underbrace{\sqrt{\gamma} \cdot \left[\nabla_{\mathbf{x},\mathbf{u}} f_{\text{nl}}(\mathbf{x}, \mathbf{u}) \Big|_{(\mathbf{x},\mathbf{u})=(\mathbf{x}_{t-1,\text{nl}}, K\mathbf{x}_{t-1,\text{nl}})} \begin{bmatrix} 0 \\ \Delta \end{bmatrix} + B\Delta \right] \mathbf{x}_{t-1,\text{nl}}}_{:=m_{t-1}} \quad (\text{B.13})$$

$$+ \underbrace{\sqrt{\gamma} \cdot \left[\nabla_{\mathbf{x},\mathbf{u}} f_{\text{nl}}(\mathbf{x}, \mathbf{u}) \Big|_{(\mathbf{x},\mathbf{u})=(\mathbf{x}_{t-1,\text{nl}}, K\mathbf{x}_{t-1,\text{nl}})} \begin{bmatrix} I \\ K \end{bmatrix} + (A + BK) \right] \psi_{t-1}}_{:=N_{t-1}}. \quad (\text{B.14})$$

Using the shorthand introduced above, we can re-express the recursion as,

$$\psi_t = m_{t-1} + N_{t-1}\psi_{t-1}.$$

Unrolling this recursion, with the base case that $\psi_0 = D_K \mathbf{x}_{0,\text{nl}} = 0$, we get that

$$\psi_t = \sum_{j=0}^{t-1} \left(\prod_{i=j+1}^{t-1} N_i \right) m_j.$$

Therefore,

$$\|\psi_t\| \leq \sum_{j=0}^{t-1} \left\| \prod_{i=j+1}^{t-1} N_i \right\|_{\text{op}} \|m_j\|. \quad (\text{B.15})$$

Next, we prove that each matrix N_i is stable so that the product of the N_i is small. By [Lemma 3.1](#) and our earlier inequalities, [Eq. \(B.10\)](#) & [Eq. \(B.11\)](#), we have that,

$$\begin{aligned} \left\| \nabla_{\mathbf{x}, \mathbf{u}} f_{\text{nl}}(\mathbf{x}, \mathbf{u}) \Big|_{(\mathbf{x}, \mathbf{u})=(\mathbf{x}_{t-1, \text{nl}}, K\mathbf{x}_{t-1, \text{nl}})} \begin{bmatrix} I \\ K \end{bmatrix} \right\|_{\text{op}} &\leq \left\| \nabla_{\mathbf{x}, \mathbf{u}} f_{\text{nl}}(\mathbf{x}, \mathbf{u}) \Big|_{(\mathbf{x}, \mathbf{u})=(\mathbf{x}_{t-1, \text{nl}}, K\mathbf{x}_{t-1, \text{nl}})} \right\|_{\text{op}} \left\| \begin{bmatrix} I \\ K \end{bmatrix} \right\|_{\text{op}} \\ &\leq \beta_{\text{nl}}(1 + \|K\|_{\text{op}})^2 \|\mathbf{x}_{t-1, \text{nl}}\| \\ &\leq \frac{1}{6\|P_{K, \gamma}\|_{\text{op}}^2}, \end{aligned} \quad (\text{B.16})$$

where we have used our initial condition on V_0 . Therefore, $N_i = \sqrt{\gamma}(A + BK) + \Delta_{N_i}$ where $\|\Delta_{N_i}\|_{\text{op}} \leq 1/(6\|P_{K, \gamma}\|_{\text{op}}^2)$ for all i . By definition of the operator norm, [Lemma A.6](#), and the fact that $P_{K, \gamma} \succeq I$:

$$\begin{aligned} \left\| \prod_{i=j+1}^{t-1} N_i \right\|_{\text{op}}^2 &= \left\| \left(\prod_{i=j+1}^{t-1} N_i \right)^\top \left(\prod_{i=j+1}^{t-1} N_i \right) \right\|_{\text{op}} \\ &\leq \left\| \left(\prod_{i=j+1}^{t-1} N_i \right)^\top P_{K, \gamma} \left(\prod_{i=j+1}^{t-1} N_i \right) \right\|_{\text{op}} \\ &\leq \|P_{K, \gamma}\|_{\text{op}} \left(1 - \frac{1}{2\|P_{K, \gamma}\|_{\text{op}}} \right)^{t-j-1}. \end{aligned}$$

Then, using our bound on $\nabla_{\mathbf{x}, \mathbf{u}} f_{\text{nl}}$ from [Eq. \(B.16\)](#), and [Lemma B.2](#), we bound m_j (defined in [Eq. \(B.13\)](#)) as,

$$\begin{aligned} \|m_j\| &\leq \left\| \sqrt{\gamma} \left[\nabla_{\mathbf{x}, \mathbf{u}} f_{\text{nl}}(\mathbf{x}, \mathbf{u}) \Big|_{(\mathbf{x}, \mathbf{u})=(\mathbf{x}_{j, \text{nl}}, K\mathbf{x}_{j, \text{nl}})} \begin{bmatrix} 0 \\ \Delta \end{bmatrix} + B\Delta \right] \right\|_{\text{op}} \|\mathbf{x}_{j, \text{nl}}\| \\ &\leq \left(\frac{1}{6\|P_{K, \gamma}\|_{\text{op}}^2} + \|B\|_{\text{op}} \right) V_0^{1/2} \left(1 - \frac{1}{2\|P_{K, \gamma}\|_{\text{op}}} \right)^{j/2}. \end{aligned}$$

Returning to our earlier bound on $\|\psi_t\|$ in [Eq. \(B.15\)](#), we conclude that,

$$\begin{aligned} \|\psi_t\| &\leq \sum_{j=0}^{t-1} \|P_{K, \gamma}\|_{\text{op}}^{1/2} \left(1 - \frac{1}{2\|P_{K, \gamma}\|_{\text{op}}} \right)^{(t-j-1)/2} \left(\frac{1}{6\|P_{K, \gamma}\|_{\text{op}}^2} + \|B\|_{\text{op}} \right) V_0^{1/2} \left(1 - \frac{1}{2\|P_{K, \gamma}\|_{\text{op}}} \right)^{j/2} \\ &\leq t \cdot \|P_{K, \gamma}\|_{\text{op}}^{1/2} \left(\frac{1}{6\|P_{K, \gamma}\|_{\text{op}}^2} + \|B\|_{\text{op}} \right) \left(1 - \frac{1}{2\|P_{K, \gamma}\|_{\text{op}}} \right)^{\frac{t-1}{2}} V_0^{1/2}, \end{aligned}$$

concluding the proof. \square

Using this, we can now return to bounding $D_K f_{\text{nl}}(\mathbf{x}_{t, \text{nl}}, K\mathbf{x}_{t, \text{nl}})$.

Lemma B.8 (Bounding $D_K f_{\text{nl}}(\mathbf{x}_{t, \text{nl}}, K\mathbf{x}_{t, \text{nl}})$). *For all $t \geq 0$, the following bound holds:*

$$\|D_K f_{\text{nl}}(\mathbf{x}_{t, \text{nl}}, K\mathbf{x}_{t, \text{nl}})\| \leq 8\beta_{\text{nl}}\|P_{K, \gamma}\|_{\text{op}}^2(1 + \|B\|_{\text{op}}) \cdot t \left(1 - \frac{1}{2\|P_{K, \gamma}\|_{\text{op}}} \right)^{t-1/2} \|\mathbf{x}_0\|^2.$$

Proof. From [Eq. \(B.12\)](#), $\|D_K f_{\text{nl}}(\mathbf{x}_{t, \text{nl}}, K\mathbf{x}_{t, \text{nl}})\|$ is less than,

$$\begin{aligned} &\left\| \nabla_{\mathbf{x}, \mathbf{u}} f_{\text{nl}}(\mathbf{x}, \mathbf{u}) \Big|_{(\mathbf{x}, \mathbf{u})=(\mathbf{x}_{t, \text{nl}}, K\mathbf{x}_{t, \text{nl}})} \right\|_{\text{op}} \left\| \begin{bmatrix} 0 \\ \Delta \end{bmatrix} \right\|_{\text{op}} \|\mathbf{x}_{t, \text{nl}}\| \\ &+ \left\| \nabla_{\mathbf{x}, \mathbf{u}} f_{\text{nl}}(\mathbf{x}, \mathbf{u}) \Big|_{(\mathbf{x}, \mathbf{u})=(\mathbf{x}_{t, \text{nl}}, K\mathbf{x}_{t, \text{nl}})} \right\|_{\text{op}} \left\| \begin{bmatrix} I \\ K \end{bmatrix} \right\|_{\text{op}} \|D_K \mathbf{x}_{t, \text{nl}}\|. \end{aligned}$$

Again using the assumption on the nonlinear dynamics, we can bound the gradient terms as in Eq. (B.16) and get that $\|D_K f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})\|$ is no larger than,

$$\begin{aligned} & \beta_{\text{nl}}(1 + \|K\|_{\text{op}})\|\mathbf{x}_{t,\text{nl}}\|^2 + \beta_{\text{nl}}(1 + \|K\|_{\text{op}})^2\|\mathbf{x}_{t,\text{nl}}\|\|D_K\mathbf{x}_{t,\text{nl}}\| \\ & \leq 2\beta_{\text{nl}}(1 + \|K\|_{\text{op}})^2\|\mathbf{x}_{t,\text{nl}}\|\max\{\|\mathbf{x}_{t,\text{nl}}\|, \|D_K\mathbf{x}_{t,\text{nl}}\|\}. \end{aligned}$$

Seeing as how our upper bound on $\|D_K\mathbf{x}_{t,\text{nl}}\|$ in Lemma B.7 is always larger than the bound for $\|\mathbf{x}_{t,\text{nl}}\|$ in Eq. (B.10), we can bound $\max\{\|\mathbf{x}_{t,\text{nl}}\|, \|D_K\mathbf{x}_{t,\text{nl}}\|\}$ by the former. Consequently,

$$\begin{aligned} \|D_K f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})\| & \leq 2\beta_{\text{nl}}(1 + \|K\|_{\text{op}})^2\|P_{K,\gamma}\|_{\text{op}}^{1/2}V_0 \left(\frac{1}{6\|P_{K,\gamma}\|_{\text{op}}^2} + \|B\|_{\text{op}} \right) \cdot t \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}} \right)^{t-1/2} \\ & \leq 8\beta_{\text{nl}}\|P_{K,\gamma}\|_{\text{op}}^2(1 + \|B\|_{\text{op}}) \cdot t \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}} \right)^{t-1/2} \|\mathbf{x}_0\|^2. \end{aligned}$$

□

Finally, we bound $D_K P_{K,\gamma}$.

Lemma B.9 (Bounding $D_K P_{K,\gamma}$). *The following bound holds:*

$$\|D_K P_{K,\gamma}\|_{\text{op}} \leq 2\|P_{K,\gamma}\|_{\text{op}}^2(\|B\|_{\text{op}} + \|K\|_{\text{op}}).$$

Proof. By definition of the discrete time Lyapunov equation,

$$P_{K,\gamma} = Q + K^\top R K + \gamma(A + BK)^\top P_{K,\gamma}(A + BK).$$

Therefore, the directional derivative $D_K P_{K,\gamma}$ satisfies another Lyapunov equation,

$$\begin{aligned} D_K P_{K,\gamma} & = \underbrace{\Delta^\top R K + K^\top R \Delta + \gamma \cdot (B\Delta)^\top P_{K,\gamma}(A + BK) + \gamma \cdot (A + BK)^\top P_{K,\gamma} B \Delta}_{:=E_K} \\ & \quad + \gamma(A + BK)^\top (D_K P_{K,\gamma})(A + BK), \end{aligned}$$

implying that $D_K P_{K,\gamma} = \text{dlyap}(\sqrt{\gamma}(A + BK), E_K)$. By properties of the Lyapunov equation,

$$\begin{aligned} D_K P_{K,\gamma} & \preceq \|E_K\|_{\text{op}} \text{dlyap}(A + BK, I) \\ & \preceq \|E_K\|_{\text{op}} \text{dlyap}(\sqrt{\gamma}(A + BK), K^\top R K + Q) = \|E_K\|_{\text{op}} P_{K,\gamma}. \end{aligned}$$

Therefore, to bound $\|D_K P_{K,\gamma}\|_{\text{op}}$ it suffices to bound, $\|E_k\|_{\text{op}}$. Using the fact that $\|P_{K,\gamma}^{1/2}\sqrt{\gamma}(A + BK)\|_{\text{op}} \leq \|P_{K,\gamma}^{1/2}\|_{\text{op}}$ and that $\|\Delta\|_{\text{op}} \leq 1$, a short calculation reveals that,

$$\|E_K\|_{\text{op}} \leq 2\|P_{K,\gamma}\|_{\text{op}}(\|B\|_{\text{op}} + \|K\|_{\text{op}}),$$

which together with our previous bound on $D_K P_{K,\gamma}$ implies that,

$$\|D_K P_{K,\gamma}\|_{\text{op}} \leq 2\|P_{K,\gamma}\|_{\text{op}}^2(\|B\|_{\text{op}} + \|K\|_{\text{op}}).$$

□

With Lemmas B.7 to B.9 in place, we now return to bounding terms T_1, T_2, T_3 .

Bounding T_1 Recall $T_1 := \gamma \cdot (D_K \cdot f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}))^\top P_{K,\gamma} \mathbf{x}_{t+1,\text{nl}}$. We then have

$$\|T_1\| \leq \|(D_K f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}}))^\top P_{K,\gamma} \mathbf{x}_{t+1,\text{nl}}\| \leq \|P_{K,\gamma}\|_{\text{op}} \|D_K f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})\| \|\mathbf{x}_{t+1,\text{nl}}\|.$$

Using the bound on $\|\mathbf{x}_{t+1,\text{nl}}\|$ stated in Eq. (B.10) and on $\|D_K f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})\|$ from Lemma B.8, the above simplifies to

$$8\beta_{\text{nl}}\|P_{K,\gamma}\|_{\text{op}}^{7/2}(1 + \|B\|_{\text{op}}) \cdot t \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}} \right)^{1.5t} \|\mathbf{x}_0\|^3.$$

Bounding T_2 Recall $T_2 = f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})D_K P_{K,\gamma}\mathbf{x}_{t+1,\text{nl}}$, so that

$$\|T_2\| \leq \|f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})\| \|D_K P_{K,\gamma}\|_{\text{op}} \|\mathbf{x}_{t+1,\text{nl}}\|.$$

Using the bound on $\|f_{\text{nl}}(\mathbf{x}, \mathbf{u})\|$ from [Lemma 3.1](#),

$$\begin{aligned} \|f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})\| &\leq \beta_{\text{nl}}(1 + \|K\|_{\text{op}}^2) \|\mathbf{x}_{t,\text{nl}}\|^2 \\ &\leq \beta_{\text{nl}}(1 + \|K\|_{\text{op}}^2) \|P_{K,\gamma}\|_{\text{op}} \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right)^t \|\mathbf{x}_0\|^2 \\ &\leq 2\beta_{\text{nl}} \|P_{K,\gamma}\|_{\text{op}}^2 \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right)^t \|\mathbf{x}_0\|^2. \end{aligned} \quad (\text{B.17})$$

Therefore, using [Eq. \(B.10\)](#) again and [Lemma B.9](#),

$$\begin{aligned} \|T_2\| &\leq 4\beta_{\text{nl}} \|P_{K,\gamma}\|_{\text{op}}^{9/2} (\|B\|_{\text{op}} + \|K\|_{\text{op}}) \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right)^{\frac{3}{2}t + \frac{1}{2}} \|\mathbf{x}_0\|^3 \\ &\leq 4\beta_{\text{nl}} \|P_{K,\gamma}\|_{\text{op}}^5 (\|B\|_{\text{op}} + 1) \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right)^{\frac{3}{2}t + \frac{1}{2}} \|\mathbf{x}_0\|^3. \end{aligned}$$

Bounding T_3 . Recall $T_3 := f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})P_{K,\gamma}D_K\mathbf{x}_{t+1,\text{nl}}$. From [Eq. \(B.17\)](#) and [Lemma B.7](#), we have that

$$\begin{aligned} \|T_3\| &\leq \|f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})\| \|P_{K,\gamma}\|_{\text{op}} \|D_K\mathbf{x}_{t+1,\text{nl}}\|_{\text{op}} \\ &\leq 2\beta_{\text{nl}} \|P_{K,\gamma}\|_{\text{op}}^4 (1 + \|B\|_{\text{op}}) \cdot (t+1) \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right)^{1.5t} \|\mathbf{x}_0\|^3. \end{aligned}$$

Wrapping up Therefore,

$$\|T_1\| + \|T_2\| + \|T_3\| \leq 12\beta_{\text{nl}}(1 + \|B\|_{\text{op}}) \|P_{K,\gamma}\|_{\text{op}}^5 (t+1) \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right)^{\frac{3}{2}t} \|\mathbf{x}_0\|^3.$$

And hence,

$$\begin{aligned} \sum_{t=0}^{\infty} D_K [f_{\text{nl}}(\mathbf{x}_{t,\text{nl}}, K\mathbf{x}_{t,\text{nl}})^\top P_{K,\gamma}\mathbf{x}_{t+1,\text{nl}}] &\leq 12\beta_{\text{nl}}(1 + \|B\|_{\text{op}}) \|P_{K,\gamma}\|_{\text{op}}^5 \|\mathbf{x}_0\|^3 \sum_{t=0}^{\infty} (t+1) \left(1 - \frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}\right)^t \\ &= 48\beta_{\text{nl}}(1 + \|B\|_{\text{op}}) \|P_{K,\gamma}\|_{\text{op}}^7 \|\mathbf{x}_0\|^3. \end{aligned}$$

□

B.3 Establishing Lyapunov functions: Proof of [Lemma B.1](#)

Proof. To be concise, we use the shorthand, $\mathbf{z} = (\mathbf{x}, K\mathbf{x})$. We start by expanding out,

$$\gamma \cdot G_{\text{nl}}(\mathbf{z})^\top P_K G_{\text{nl}}(\mathbf{z}) = \gamma (G_{\text{lin}}(\mathbf{z})^\top P_K G_{\text{lin}}(\mathbf{z}) + f_{\text{nl}}(\mathbf{z})^\top P_K f_{\text{nl}}(\mathbf{z}) + 2f_{\text{nl}}(\mathbf{z})^\top P_K G_{\text{lin}}(\mathbf{z})).$$

By the AM-GM inequality for vectors, the following holds for any $\tau > 0$,

$$2f_{\text{nl}}(\mathbf{z})^\top P_K G_{\text{lin}}(\mathbf{z}) \leq \tau \cdot G_{\text{lin}}(\mathbf{z})^\top P_K G_{\text{lin}}(\mathbf{z}) + \frac{1}{\tau} \cdot f_{\text{nl}}(\mathbf{z})^\top P_K f_{\text{nl}}(\mathbf{z}).$$

Combining these two relationships, we get that,

$$\gamma \cdot G_{\text{nl}}(\mathbf{z})^\top P_K G_{\text{nl}}(\mathbf{z}) \leq \gamma \cdot (1 + \tau) \cdot G_{\text{lin}}(\mathbf{z})^\top P_K G_{\text{lin}}(\mathbf{z}) + \gamma \cdot \left(1 + \frac{1}{\tau}\right) \cdot f_{\text{nl}}(\mathbf{z})^\top P_K f_{\text{nl}}(\mathbf{z}). \quad (\text{B.18})$$

Next, by properties of the Lyapunov function [Lemma A.5](#), we have that

$$\gamma G_{\text{lin}}(\mathbf{z})^\top P_K G_{\text{lin}}(\mathbf{z}) = \gamma \cdot \mathbf{x}^\top (A + BK)^\top P_{K,\gamma} (A + BK) \mathbf{x} \leq \mathbf{x}^\top P_{K,\gamma} \mathbf{x} \left(1 - \frac{1}{\|P_{K,\gamma}\|_{\text{op}}}\right).$$

Letting, $V_{\mathbf{x}} := \mathbf{x}^\top P_{K,\gamma} \mathbf{x}$, we can plug in the previous expression into Eq. (B.18) and optimize over τ to get that,

$$\begin{aligned} \gamma \cdot G_{\text{nl}}(\mathbf{z})^\top P_{K,\gamma} G_{\text{nl}}(\mathbf{z}) &\leq V_{\mathbf{x}} \left(1 - \frac{1}{\|P_{K,\gamma}\|_{\text{op}}} \right) + f_{\text{nl}}(\mathbf{z})^\top P_{K,\gamma} f_{\text{nl}}(\mathbf{z}) + 2\sqrt{V_{\mathbf{x}} f_{\text{nl}}(\mathbf{z})^\top P_{K,\gamma} f_{\text{nl}}(\mathbf{z})} \\ &\leq V_{\mathbf{x}} \left(1 - \frac{1}{\|P_{K,\gamma}\|_{\text{op}}} \right) + \|P_{K,\gamma}\|_{\text{op}} \|f_{\text{nl}}(\mathbf{z})\|^2 + 2\sqrt{V_{\mathbf{x}} \|P_{K,\gamma}\|_{\text{op}}} \|f_{\text{nl}}(\mathbf{z})\|, \end{aligned}$$

where we have dropped a factor of γ from the last two terms since $\gamma \leq 1$. Next, the proof follows by noting that this following inequality is satisfied,

$$\|P_{K,\gamma}\|_{\text{op}} \|f_{\text{nl}}(\mathbf{z})\|^2 + 2\sqrt{V_{\mathbf{x}} \|P_{K,\gamma}\|_{\text{op}}} \|f_{\text{nl}}(\mathbf{z})\| \leq \frac{V_{\mathbf{x}}}{2\|P_{K,\gamma}\|_{\text{op}}}, \quad (\text{B.19})$$

whenever,

$$\begin{aligned} \|f_{\text{nl}}(\mathbf{z})\| &\leq \sqrt{\frac{V_{\mathbf{x}}}{2\|P_{K,\gamma}\|_{\text{op}}^2} + \frac{V_{\mathbf{x}}}{\|P_{K,\gamma}\|_{\text{op}}}} - \sqrt{\frac{V_{\mathbf{x}}}{\|P_{K,\gamma}\|_{\text{op}}}} \\ &= \sqrt{\frac{V_{\mathbf{x}}}{\|P_{K,\gamma}\|_{\text{op}}}} \left(\sqrt{\frac{1}{2\|P_{K,\gamma}\|_{\text{op}}}} + 1 - 1 \right). \end{aligned} \quad (\text{B.20})$$

Therefore, assuming the inequality above Eq. (B.20), we get our desired result showing that

$$G_{\text{nl}}(\mathbf{z})^\top P_K G_{\text{nl}}(\mathbf{z}) \leq \mathbf{x}^\top P_K \mathbf{x} \cdot \left(1 - \frac{1}{2\|P_K\|_{\text{op}}} \right).$$

We conclude the proof by showing that Eq. (B.20) is satisfied for all \mathbf{x} small enough. In particular, using our bounds on f_{nl} from Lemma 3.1, if $\|\mathbf{x}\| + \|K\mathbf{x}\| \leq r_{\text{nl}}$,

$$\|f_{\text{nl}}(\mathbf{z})\| \leq \beta_{\text{nl}}(\|\mathbf{x}\|^2 + \|K\mathbf{x}\|^2) \leq \beta_{\text{nl}}(1 + \|K\|_{\text{op}}^2)\|\mathbf{x}\|^2.$$

Since $V_{\mathbf{x}} \geq \|\mathbf{x}\|^2$, in order for Eq. (B.20) to hold, it suffices for $\|\mathbf{x}\|$ to satisfy

$$\|\mathbf{x}\| \leq \min \left\{ \frac{r_{\text{nl}}}{1 + \|K\|_{\text{op}}}, \frac{1}{\beta_{\text{nl}}\|P_{K,\gamma}\|_{\text{op}}^{1/2}(1 + \|K\|_{\text{op}}^2)} \right\}.$$

Using the fact that $\|K\|_{\text{op}}^2 \leq \|P_{K,\gamma}\|_{\text{op}}$, we can simplify upper bound on $\|\mathbf{x}\|$ to be,

$$\|\mathbf{x}\| \leq \frac{r_{\text{nl}}}{2\beta_{\text{nl}}\|P_{K,\gamma}\|_{\text{op}}^{3/2}}.$$

Note that this condition is always implied by the condition on $\mathbf{x}^\top P_{K,\gamma} \mathbf{x}$ in the statement of the proposition. \square

B.4 Bounding the nonlinearity: Proof of Lemma 3.1

Since the origin is an equilibrium point, $G_{\text{nl}}(0, 0) = 0$, we can rewrite G_{nl} as,

$$\begin{aligned} G_{\text{nl}}(\mathbf{x}, \mathbf{u}) &= G_{\text{nl}}(0, 0) + \nabla G_{\text{nl}}(\mathbf{x}, \mathbf{u})|_{(\mathbf{x}, \mathbf{u})=(0,0)} \begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} + f_{\text{nl}}(\mathbf{x}, \mathbf{u}), \\ &= A_{\text{jac}} \mathbf{x} + B_{\text{jac}} \mathbf{u} + f_{\text{nl}}(\mathbf{x}, \mathbf{u}), \end{aligned}$$

for some function f_{nl} . Taking gradients,

$$\nabla f_{\text{nl}}(\mathbf{x}, \mathbf{u}) = \nabla_{\mathbf{x}, \mathbf{u}} G_{\text{nl}}(\mathbf{x}, \mathbf{u})|_{(\mathbf{x}, \mathbf{u})=(\mathbf{x}, \mathbf{u})} - \nabla G_{\text{nl}}(\mathbf{x}, \mathbf{u})|_{(\mathbf{x}, \mathbf{u})=(0,0)}.$$

Hence, for all \mathbf{x}, \mathbf{u} such that the smoothness assumption holds, we get that

$$\|\nabla_{\mathbf{x}, \mathbf{u}} f_{\text{nl}}(\mathbf{x}, \mathbf{u})\|_{\text{op}} \leq \beta_{\text{nl}}(\|\mathbf{x}\| + \|\mathbf{u}\|).$$

Next, by Taylor’s theorem,

$$f_{\text{nl}}(\mathbf{x}, \mathbf{u}) = f_{\text{nl}}(0, 0) + \int_0^1 \nabla f_{\text{nl}}(t \cdot \mathbf{x}, t \cdot \mathbf{u}) \begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} dt,$$

and since $f_{\text{nl}}(0, 0) = 0$, we can bound,

$$\begin{aligned} \|f_{\text{nl}}(\mathbf{x}, \mathbf{u})\| &\leq \int_0^1 \|\nabla f_{\text{nl}}(t \cdot \mathbf{x}, t \cdot \mathbf{u})\|_{\text{op}} \begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} \|dt \\ &\leq \int_0^1 \beta t \cdot (\|\mathbf{x}\| + \|\mathbf{u}\|)^2 dt. \\ &\leq 2\beta(\|\mathbf{x}\|^2 + \|\mathbf{u}\|^2) \int_0^1 t \cdot dt \\ &= \beta(\|\mathbf{x}\|^2 + \|\mathbf{u}\|^2). \end{aligned}$$

C Implementing ε -Eval, ε -Grad, and Search Algorithms

We conclude by discussing the implementations and relevant sample complexities of the noisy gradient and function evaluation methods, ε -Grad and ε -Eval. We then use these to establish the runtime and correctness of the noisy binary and random search algorithms.

We remark that throughout our analysis, we have assumed that ε -Grad and ε -Eval succeed with probability 1. This is purely for the sake of simplifying our presentation. As established in Fazel et al. [2018], the relevant estimators in this section all return ε approximate solutions with probability $1 - \delta$, where δ factors into the sample complexity polynomially. Therefore, it is easy to union bound and get a high probability guarantee. We omit these union bounds for the sake of simplifying the presentation.

C.1 Implementing ε -Eval

Linear setting. From the analysis in Fazel et al. [2018], we know that if $J_{\text{lin}}(K | \gamma) < \infty$, then

$$J_{\text{lin}}(K | \gamma) \approx \frac{1}{N} \sum_{i=1}^N J_{\text{lin}}^{(H)}(K | \gamma, \mathbf{x}^{(i)}), \quad \mathbf{x}^{(i)} \sim r \cdot \mathcal{S}^{d_x-1}, \quad (\text{C.1})$$

where $J_{\text{lin}}^{(H)}$ is the length H , finite horizon cost of K :

$$J_{\text{lin}}^{(H)}(K | \gamma, \mathbf{x}) = \sum_{j=0}^{H-1} \gamma^j \cdot (\mathbf{x}_j^\top Q \mathbf{x}_j + \mathbf{u}_j^\top R \mathbf{u}_j), \quad \mathbf{u}_j = K \mathbf{x}_j, \quad \mathbf{x}_{j+1} = A \mathbf{x}_j + B \mathbf{u}_j, \quad \mathbf{x}_0 = \mathbf{x}.$$

More formally, if $J_{\text{lin}}(K | \gamma)$ is smaller than a constant c , Lemma 26 from Fazel et al. [2018] states that it suffices to set N and H to be polynomials in $\|A\|_{\text{op}}, \|B\|_{\text{op}}, J_{\text{lin}}(K | \gamma)$ and c/ε in order to have an ε -approximate estimate of $J_{\text{lin}}(K | \gamma)$ with high probability.

On the other hand, if $J_{\text{lin}}(K | \gamma)$ is larger than c (recall that the costs should never be larger than some universal constant times M_{lin} during discount annealing), the following lemma proves we can detect that this is the case by setting N and H to be polynomials in $\|A\|_{\text{op}}, \|B\|_{\text{op}}$ and c . This argument follows from the following two lemmas:

Lemma C.1. Fix a constant c , and take $H = 4c$. Then for any K (possibly even with $J_{\text{lin}}(K | \gamma) = \infty$),

$$\mathbb{P}_{\mathbf{x} \sim \sqrt{d_x} \cdot \mathcal{S}^{d_x-1}} \left[J_{\text{lin}}^{(H)}(K | \gamma, \mathbf{x}) \geq \frac{\min \left\{ \frac{1}{2} J_{\text{lin}}(K | \gamma), c \right\}}{d_x} \right] \geq \frac{1}{10}.$$

Setting $c = \alpha M_{\text{lin}}$ for some universal constant α , we get that all calls to the ε -Eval oracle during discount annealing can be implemented with polynomially many samples. Using standard arguments, we can boost this result to hold with high probability by running N independent trials, where N is again a polynomial in the relevant problem parameters.

Nonlinear setting. Next, we sketch why the same estimator described in Eq. (C.1) also works in the nonlinear case if we replace $J_{\text{lin}}^{(H)}(K | \gamma)$ with the analogous finite horizon cost, $J_{\text{nl}}^{(H)}(K | \gamma, r_*)$, for the nonlinear objective. By Lemma B.3 and Proposition 3.3 if the nonlinear cost is small, then costs on the nonlinear system and the linear system are *pointwise* close. Therefore, previous concentration analysis for the linear setting from Fazel et al. [2018] can be easily carried over in order to implement ε -Eval where N and H are depend polynomially on the relevant problem parameters.

On the other hand if the nonlinear cost is large, then the cost on the linear system must also be large. Recall that if the linear cost was bounded by a constant, then the costs of both systems would be pointwise close by Proposition 3.3. By Proposition C.5, we know that the H -step nonlinear cost is lower bounded by the cost on the linear system. Since we can always detect that the cost on the linear system is larger than a constant using polynomially many samples as per Lemma C.1, with high probability we can also detect if the cost on the nonlinear system is large using again only polynomially many samples.

C.2 Implementing ε -Grad

Linear setting. Just like in the case of ε -Eval, Fazel et al. [2018] (Lemma 30) prove that

$$\nabla J_{\text{lin}}(K | \gamma) \approx \frac{1}{N} \sum_{i=1}^N \nabla J_{\text{lin}}^{(H)}(K | \gamma, \mathbf{x}^{(i)}), \quad \mathbf{x}^{(i)} \sim r \cdot \mathcal{S}^{d_x-1} \quad (\text{C.2})$$

where N and H only need to be made polynomial large in the relevant problem parameter and $1/\varepsilon$ in order to get an ε accurate approximation of the gradient. Note that we can safely assume that $J_{\text{lin}}(K | \gamma)$ is finite (and hence the gradient is well defined) since we can always run the test outlined in Lemma C.1 to get a high probability guarantee of this fact. In order to approximate the gradients $\nabla J_{\text{lin}}^{(H)}(K | \gamma, \mathbf{x})$, one can employ standard techniques from Flaxman et al. [2005]. We refer the interested reader to Appendix D in Fazel et al. [2018] for further details.

Nonlinear setting. Lastly, we remark that, as in the case of ε -Eval, Lemma B.6 establishes that the gradients between the linear and nonlinear system are again pointwise close if the cost on the linear system is bounded. In the proof of Theorem 2, we established that during all iterations of discount annealing, the cost on the linear system during executions of policy gradients is bounded by M_{nl} . Therefore, the analysis from Fazel et al. [2018] can be ported over to show that ε -approximate gradients of the nonlinear system can be computed using only polynomially many samples in the relevant problem parameters.

C.3 Auxiliary results

Lemma C.2. Fix any constant $c \geq 1$. Then, for

$$P_H = \sum_{j=0}^{H-1} ((\sqrt{\gamma}(A + BK))^j)^\top (Q + K^\top RK) (\sqrt{\gamma}(A + BK))^j$$

the following relationship holds

$$J_{\text{lin}}^{(H)}(K | \gamma) \geq \min \left\{ \frac{1}{2} J_{\text{lin}}(K | \gamma), c \right\}, \quad \forall H \geq 4c,$$

where $J_{\text{lin}}(K | \gamma)$ may be infinite.

Proof. We have two cases. In the first case $A_{\text{cl}} := \sqrt{\gamma}(A + BK)$ has an eigenvalue $\lambda \geq 1$. Letting \mathbf{v} denote such a corresponding eigenvector of unit norm, one can verify that $\text{tr}(P_H) \geq \mathbf{v}^\top P_H \mathbf{v} \geq \sum_{i=0}^{H-1} \lambda_i^2 \geq H$, which is at least c by assumption.

In the second case, A_{cl} is stable, so $P_{K,\gamma}$ exists and its trace is $J_{\text{lin}}(K | \gamma)$ (see Fact A.2). Then, for $Q_K = Q + K^\top RK$,

$$P_H = \sum_{i=0}^{H-1} A_{\text{cl}}^{i\top} Q_K A_{\text{cl}}^i,$$

we show that if $\text{tr}(P_H) \leq c$, then $\text{tr}(P_H) \geq \frac{1}{2}\text{tr}(P_{K,\gamma})$.

To show this, observe that if $\text{tr}(P_H) \leq c$, then $\text{tr}(P_H) \leq H/4$. Therefore, by the pigeonhole principle (and the fact that $\text{tr}(Q_K) \geq 1$), there exists some $t \in \{1, \dots, H\}$ such that $\text{tr}(A_{\text{cl}}^{t\top} Q_K A_{\text{cl}}^t) \leq 1/4$. Since $Q_K \succeq I$, this means that $\text{tr}(A_{\text{cl}}^{t\top} A_{\text{cl}}^t) = \|A_{\text{cl}}^t\|_{\text{F}}^2 \leq 1/4$ as well. Therefore, letting P_t denote the t -step value function from Eq. (C.3), the identity $P_{K,\gamma} = \sum_{n=0}^{\infty} A_{\text{cl}}^{nt\top} P_t A_{\text{cl}}^{nt}$ means that

$$\begin{aligned} \text{tr}(P_{K,\gamma}) &= \text{tr} \left(\sum_{n=0}^{\infty} A_{\text{cl}}^{nt\top} P_t A_{\text{cl}}^{nt} \right) \\ &\leq \text{tr}(P_t) + \|P_t\| \sum_{n \geq 1} \|A_{\text{cl}}^{nt}\|_{\text{F}}^2 \\ &\leq \text{tr}(P_t) + \|P_t\| \sum_{n \geq 1} \|A_{\text{cl}}^t\|_{\text{F}}^{2n} \\ &\leq \text{tr}(P_t) + \|P_t\| \sum_{n \geq 1} (1/2)^n \\ &\leq 2\text{tr}(P_t) \leq 2\text{tr}(P_H), \end{aligned}$$

where in the last line we used $t \leq H$. This completes the proof. \square

Lemma C.3. Let $\mathbf{u}, \mathbf{v} \stackrel{\text{i.i.d.}}{\sim} \mathcal{S}^{d_x-1}$ then, $\mathbb{P}[(\mathbf{u}^\top \mathbf{v})^2 \geq 1/d_x] \geq .1$.

Proof. The statement is clearly true for $d_x = 1$, therefore we focus on the case where $d_x \geq 2$. Without loss of generality, we can take \mathbf{u} to be the first basis vector \mathbf{e}_1 and let $\mathbf{v} = Z/\|Z\|$ where $Z \sim \mathcal{N}(0, I)$. From these simplifications, we observe that $(\mathbf{v}^\top \mathbf{u})^2$ is equal in distribution to the following random variable,

$$\frac{Z_1^2}{\sum_{i=1}^d Z_i^2},$$

where each Z_i^2 is a chi-squared random variable. Using this equivalency, we have that for arbitrary $M > 0$,

$$\begin{aligned} \mathbb{P}[(\mathbf{u}^\top \mathbf{v})^2 \geq 1/d_x] &= \mathbb{P} \left[Z_1^2 \geq \frac{\sum_{i=2}^{d_x} Z_i^2}{d_x - 1} \right] \\ &= \mathbb{P} \left[Z_1^2 \geq \frac{\sum_{i=2}^{d_x} Z_i^2}{d_x - 1} \mid \sum_{i=2}^{d_x} Z_i^2 \leq M \right] \mathbb{P} \left[\sum_{i=2}^{d_x} Z_i^2 \leq M \right] \\ &\quad + \mathbb{P} \left[Z_1^2 \geq \frac{\sum_{i=2}^{d_x} Z_i^2}{d_x - 1} \mid \sum_{i=2}^{d_x} Z_i^2 > M \right] \mathbb{P} \left[\sum_{i=2}^{d_x} Z_i^2 > M \right] \\ &\geq \mathbb{P} \left[Z_1^2 \geq \frac{\sum_{i=2}^{d_x} Z_i^2}{d_x - 1} \mid \sum_{i=2}^{d_x} Z_i^2 \leq M \right] \mathbb{P} \left[\sum_{i=2}^{d_x} Z_i^2 \leq M \right] \\ &\geq \mathbb{P} \left[Z_1^2 \geq \frac{M}{d_x - 1} \right] \mathbb{P} \left[\sum_{i=2}^{d_x} Z_i^2 \leq M \right]. \end{aligned}$$

Setting $M = 2(d_x - 1)$, we get that,

$$\mathbb{P}[(\mathbf{u}^\top \mathbf{v})^2 \geq 1/d_x] \geq \mathbb{P}[Z_1^2 \geq 2] \mathbb{P} \left[\sum_{i=2}^{d_x} Z_i^2 \leq 2(d_x - 1) \right].$$

From a direct computation,

$$\mathbb{P}[Z_1^2 \geq 2] \geq .15.$$

To bound the last term, if Y is a chi-squared random variable with k degrees of freedom, by Lemma 1 in Laurent and Massart [2000],

$$\mathbb{P}[Y \geq k + 2\sqrt{kx} + x] \leq \exp(-x).$$

Setting $x = 2\sqrt{2(d_x - 1)k} + 2d_x + k - 2$ we get that $k + 2\sqrt{kx} + x = 2(d_x - 1)$. Substituting in $k = d_x - 1$, we conclude that,

$$\mathbb{P} \left[\sum_{i=2}^{d_x} Z_i^2 \leq 2(d_x - 1) \right] \geq 1 - \exp(-2\sqrt{2}(d_x - 1) - 3d_x + 3),$$

which is greater than .99 for $d_x \geq 2$. \square

Lemma C.1 (restated). Fix a constant c , and take $H = 4c$. Then for any K (possibly even with $J_{\text{lin}}(K | \gamma) = \infty$),

$$\mathbb{P}_{\mathbf{x} \sim \sqrt{d_x} \cdot \mathcal{S}^{d_x-1}} \left[J_{\text{lin}}^{(H)}(K | \gamma, \mathbf{x}) \geq \frac{\min \left\{ \frac{1}{2} J_{\text{lin}}(K | \gamma), c \right\}}{d_x} \right] \geq \frac{1}{10}.$$

Proof. Observe that for the finite horizon value matrix $P_H = \sum_{i=0}^{H-1} (A+BK)^{i\top} (Q+K^\top RK)(A+BK)^i$, we have $J_{\text{lin}}^{(H)}(K | \gamma, \mathbf{x}) = \mathbf{x}^\top P_H \mathbf{x}$. We now observe that, since $P_H \succeq 0$, it has a (possibly non-unique) top eigenvector v_1 for which

$$\begin{aligned} J_{\text{lin}}^{(H)}(K | \gamma, \mathbf{x}) &= \mathbf{x}^\top P_H \mathbf{x} \geq \langle v_1, \frac{\mathbf{x}}{\sqrt{d_x}} \rangle^2 \cdot d_x \|P_H\| \\ &\geq \langle v_1, \frac{\mathbf{x}}{\sqrt{d_x}} \rangle^2 \cdot \text{tr}(P_H) \\ &= \underbrace{\langle v_1, \frac{\mathbf{x}}{\sqrt{d_x}} \rangle^2}_{:=Z} \cdot J_{\text{lin}}^{(H)}(K | \gamma) \end{aligned}$$

Since $\mathbf{x}/\sqrt{d_x} \sim \mathcal{S}^{d_x-1}$, [Lemma C.3](#) ensures that $\mathbb{P}[Z \geq 1/d_x] \geq 1/10$. Hence,

$$\mathbb{P}_{\mathbf{x} \sim \sqrt{d_x} \cdot \mathcal{S}^{d_x-1}} \left[J_{\text{lin}}^{(H)}(K | \gamma, \mathbf{x}) \geq \frac{J_{\text{lin}}^{(H)}(K | \gamma)}{d_x} \right] \geq \frac{1}{10}.$$

The bound now follows from invoking [Lemma C.2](#) to lower bound $J_{\text{lin}}^{(H)}(K | \gamma, \mathbf{x}) \geq \min \left\{ \frac{1}{2} J_{\text{lin}}(K | \gamma), c \right\}$ provided $H \geq 4c$. \square

In the following lemmas, we define P_H to be the following matrix where K is any state-feedback controller.

$$P_H = \sum_{j=0}^{H-1} ((\sqrt{\gamma}(A+BK))^j)^\top (Q+K^\top RK)(\sqrt{\gamma}(A+BK))^j \quad (\text{C.3})$$

Similarly, we let $J_{\text{nl}}^{(H)}(K | \gamma, \mathbf{x}_0)$ be the horizon H cost of the nonlinear dynamical system:

$$J_{\text{nl}}^{(H)}(K | \mathbf{x}, \gamma) := \sum_{t=0}^{H-1} \mathbf{x}_t^\top Q \mathbf{x}_t + \mathbf{u}_t^\top R \mathbf{u}_t \quad (\text{C.4})$$

$$\text{s.t } \mathbf{u}_t = K \mathbf{x}_t, \quad \mathbf{x}_{t+1} = \sqrt{\gamma} \cdot G_{\text{nl}}(\mathbf{x}_t, \mathbf{u}_t), \quad \mathbf{x}_0 = \mathbf{x}. \quad (\text{C.5})$$

And again overloading notation like before, we let $J_{\text{nl}}^{(H)}(K | \gamma, r) := \mathbb{E}_{\mathbf{x} \sim r \cdot \mathcal{S}^{d_x-1}} \left[J_{\text{nl}}^{(H)}(K | \gamma, \mathbf{x}) \right] \times \frac{d_x}{r^2}$.

Lemma C.4. Fix a horizon H , constant $\alpha \in (0, d_x)$, $\mathbf{x}_0 \in \mathbb{R}^{d_x}$, and suppose that

$$\|\mathbf{x}_0\|^2 \cdot \frac{\alpha}{2H^2 \beta_{\text{nl}}^2 (1 + \|K\|^2) d_x} \leq r_{\text{nl}}^4.$$

Furthermore, define $\mathcal{X}_\alpha := \{\mathbf{x}_0 \in \mathbb{R}^{d_x} : \langle \mathbf{x}_0, \mathbf{v}_{\max}(P_H) \rangle^2 \geq \alpha \|\mathbf{x}_0\|^2 / d_x\}$. Then, if $\mathbf{x}_0 \in \mathcal{X}_\alpha$, it holds that

$$J_{\text{nl}}^{(H)}(K | \gamma, \mathbf{x}_0) \geq \min \left\{ \frac{\alpha \|\mathbf{x}_0\|^2}{4d_x^2} \cdot J_{\text{lin}}^{(H)}(K | \mathbf{x}_0, \gamma), \sqrt{\frac{\alpha}{d_x}} \frac{\|\mathbf{x}_0\|}{2H \beta_{\text{nl}} (1 + \|K\|)} \right\}.$$

Proof. Fix a constant $r_1 \leq r_{\text{nl}}$ to be selected. Throughout, we use the shorthand $\beta_1^2 = \beta_{\text{nl}}^2 (1 + \|K\|^2)$. We consider two cases.

Case 1: The initial point \mathbf{x}_0 is such that it is always the case that $\|\mathbf{x}_t\| \leq r_1$ for all $t \in \{0, 1, \dots, H-1\}$ and $J_{\text{nl}}^{(H)}(K \mid \gamma, \mathbf{x}_0) \leq r_1^2$. Observe that we can write the nonlinear dynamics as

$$\mathbf{x}_{t+1} = \sqrt{\gamma}(A + BK)\mathbf{x}_t + \mathbf{w}_t,$$

where $\mathbf{w}_t = \sqrt{\gamma} \cdot f_{\text{nl}}(\mathbf{x}_t, K\mathbf{x}_t)$. We now write:

$$\begin{aligned} \mathbf{x}_t &= \mathbf{x}_{t;0} + \mathbf{x}_{t;w}, \text{ where} \\ \mathbf{x}_{t;0} &= \sqrt{\gamma}(A + BK)^t \mathbf{x}_0 \\ \mathbf{x}_{t;w} &= \sum_{i=0}^{t-1} \gamma^{i/2} (A + BK)^i \mathbf{w}_{t-i}. \end{aligned}$$

Then, setting $Q_K = Q + K^\top RK$,

$$\begin{aligned} J_{\text{nl}}^{(H)}(K \mid \gamma, \mathbf{x}_0) &= \sum_{t=0}^{H-1} \mathbf{x}_t^\top Q_K \mathbf{x}_t \\ &= \sum_{t=0}^{H-1} (\mathbf{x}_{t;0}^\top Q_K \mathbf{x}_{t;0} + 2\mathbf{x}_{t;0}^\top Q_K \mathbf{x}_{t;w} + \mathbf{x}_{t;w}^\top Q_K \mathbf{x}_{t;w}) \\ &\geq \frac{1}{2} \sum_{t=0}^{H-1} \mathbf{x}_{t;0}^\top Q_K \mathbf{x}_{t;0} - \sum_{t=0}^{H-1} \mathbf{x}_{t;w}^\top Q_K \mathbf{x}_{t;w} \\ &= \frac{1}{2} \mathbf{x}_0^\top P_H \mathbf{x}_0 - \sum_{t=0}^{H-1} \mathbf{x}_{t;w}^\top Q_K \mathbf{x}_{t;w}, \end{aligned}$$

where (a) the last inequality uses the elementary inequality $\langle \mathbf{v}_1, \Sigma \mathbf{v}_1 \rangle + \langle \mathbf{v}_2, \Sigma \mathbf{v}_2 \rangle + 2\langle \mathbf{v}_1, \Sigma \mathbf{v}_2 \rangle \geq \frac{1}{2} \langle \mathbf{v}_1, \Sigma \mathbf{v}_1 \rangle - \langle \mathbf{v}_2, \Sigma \mathbf{v}_2 \rangle$ for any pair of vectors $\mathbf{v}_1, \mathbf{v}_2$ and $\Sigma \succeq 0$, and (b) the last inequality recognizes how

$$\sum_{t=0}^{H-1} \mathbf{x}_{t;0}^\top Q_K \mathbf{x}_{t;0} = J_{\text{lin}}^{(H)}(K \mid \mathbf{x}_0, \gamma) = \mathbf{x}_0^\top P_H \mathbf{x}_0$$

for P_H defined above in Eq. (C.3). Moreover, for any t ,

$$\begin{aligned} \sum_{t=0}^{H-1} \mathbf{x}_{t;w}^\top Q_K \mathbf{x}_{t;w} &= \sum_{t=0}^{H-1} \sum_{i=0}^{t-1} \sum_{j=0}^{t-1} \mathbf{w}_{t-i}^\top \gamma^{(i+j)/2} ((A + BK)^i)^\top Q_K (A + BK)^j \mathbf{w}_{t-j} \\ &\leq H \sum_{t=0}^{H-1} \sum_{i=0}^{t-1} \mathbf{w}_{t-i}^\top \gamma^i ((A + BK)^i)^\top Q_K (A + BK)^i \mathbf{w}_{t-i} \\ &= H \sum_{t=0}^{H-2} \mathbf{w}_t^\top \left(\sum_{i=0}^{H-t-1} \gamma^i ((A + BK)^i)^\top Q_K (A + BK)^i \right) \mathbf{w}_t \\ &= H \sum_{t=0}^{H-2} \mathbf{w}_t^\top P_{H-t-1} \mathbf{w}_t \leq H^2 \|P_H\|_{\text{op}} \max_{t \in \{0, \dots, H-2\}} \|\mathbf{w}_t\|^2. \end{aligned}$$

Now, because $\|\mathbf{x}_t\| \leq r_1 \leq r_{\text{nl}}$, Lemma 3.1 lets use bound $\|\mathbf{w}_t\|^2 \leq \beta_{\text{nl}}^2 (1 + \|K\|^2) \|\mathbf{x}_t\|^4 \leq \beta_1^2 r_1^4$, where we adopt the previously defined shorthand $\beta_1^2 = \beta_{\text{nl}}^2 (1 + \|K\|^2)$. Therefore,

$$J_{\text{nl}}^{(H)}(K \mid \gamma, \mathbf{x}_0) \geq \frac{1}{2} \mathbf{x}_0^\top P_H \mathbf{x}_0 - H^2 \beta_1^2 r_1^4 \|P_H\|_{\text{op}}.$$

Next, if $\mathbf{x}_0 \in \mathcal{X}_\alpha$,

$$\begin{aligned} J_{\text{nl}}^{(H)}(K \mid \gamma, \mathbf{x}_0) &\geq \frac{\alpha}{2d_x} \|\mathbf{x}_0\|^2 \|P_H\| - H^2 \beta_1^2 r_1^4 \|P_H\|_{\text{op}} \\ &= \|P_H\| \left(\frac{\alpha}{2d_x} \|\mathbf{x}_0\|^2 - H^2 \beta_1^2 r_1^4 \right). \end{aligned}$$

In particular, selecting $r_1^4 = \frac{\alpha}{4\beta_1^2 d_x H^2} \|\mathbf{x}_0\|^2$ (which ensures $r_1 \leq r_*$ by the conditions of the lemma), it holds that,

$$J_{\text{nl}}^{(H)}(K \mid \gamma, \mathbf{x}_0) \geq \frac{\|P_H\| \alpha}{4d_x} \|\mathbf{x}_0\|^2 \geq \frac{\text{tr}(P_H) \alpha}{4d_x^2} \|\mathbf{x}_0\|^2 = \frac{\alpha J_{\text{lin}}^{(H)}(K, \gamma)}{4d_x^2} \|\mathbf{x}_0\|^2.$$

Case 2: The initial point \mathbf{x}_0 is such that it is always the case that either $\|\mathbf{x}_t\| \geq r_1$ for all $t \in \{0, 1, \dots, H-1\}$ or $J_{\text{nl}}^{(H)}(K \mid \gamma, \mathbf{x}_0) \geq r_1^2$. Therefore, in either case, $J_{\text{nl}}^{(H)}(K \mid \gamma, \mathbf{x}_0) \geq r_1^2$. For our choice of r_1 , this gives

$$J_{\text{nl}}^{(H)}(K \mid \gamma, \mathbf{x}_0) \geq \sqrt{\frac{\alpha \|\mathbf{x}_0\|^2}{4d_x H^2 \beta_1^2}} = \sqrt{\frac{\alpha \|\mathbf{x}_0\|^2}{4d_x H^2 \beta_1^2 (1 + \|K\|^2)}} \geq \sqrt{\frac{\alpha}{d_x}} \frac{\|\mathbf{x}_0\|}{2H\beta_{\text{nl}}(1 + \|K\|)}.$$

Combining the cases, we have

$$J_{\text{nl}}^{(H)}(K \mid \gamma, \mathbf{x}_0) \geq \min \left\{ \frac{\alpha \|\mathbf{x}_0\|^2}{4d_x^2} \cdot J_{\text{lin}}^{(H)}(K \mid \mathbf{x}_0, \gamma), \sqrt{\frac{\alpha}{d_x}} \frac{\|\mathbf{x}_0\|}{2H\beta_{\text{nl}}(1 + \|K\|)} \right\}.$$

□

Proposition C.5. Let c be a given (integer) tolerance c and $\mathcal{X}_\alpha := \{\mathbf{x} \in \mathbb{R}^d : \langle \mathbf{x}, \mathbf{v}_{\max}(P_H) \rangle^2 \geq \alpha \|\mathbf{x}\|^2 / d_x\}$ be defined as in [Lemma C.4](#). Then, for $\alpha \in (0, 2]$, $H = 4c$, and $\mathbf{x}_0 \in \mathcal{X}_\alpha$ satisfying,

$$\|\mathbf{x}_0\|^2 \leq c^2 r_{\text{nl}}^4, \quad \text{and} \quad \|\mathbf{x}_0\| \leq \frac{d_x}{64c^2 \beta_{\text{nl}}(1 + \|K\|)}, \quad (\text{C.6})$$

it holds that:

$$\frac{J_{\text{nl}}^{(H)}(K \mid \gamma, \mathbf{x}_0)}{\|\mathbf{x}_0\|^2} \geq \frac{\alpha}{8d_x^2} \min \{J_{\text{lin}}(K \mid \gamma), c\}.$$

Moreover, for $r \leq \min\{cr_{\text{nl}}^2, \frac{d_x}{64c^2 \beta_{\text{nl}}(1 + \|K\|)}\}$,

$$J_{\text{nl}}^{(H)}(K \mid \gamma, r) \geq \frac{1}{80d_x^2} \min \{J_{\text{lin}}(K \mid \gamma), c\}.$$

Proof. From [Lemma C.4](#), it holds that for $H = 4c$,

$$\|\mathbf{x}_0\|^2 \cdot \frac{\alpha}{32c^2 \beta_{\text{nl}}^2 (1 + \|K\|^2) d_x} \leq r_{\text{nl}}^4 \quad (\text{C.7})$$

and hence

$$J_{\text{nl}}^{(H)}(K \mid \gamma, \mathbf{x}_0) \geq \min \left\{ \frac{\alpha \|\mathbf{x}_0\|^2}{4d_x^2} \cdot J_{\text{lin}}^{(H)}(K \mid \mathbf{x}_0, \gamma), \sqrt{\frac{\alpha}{d_x}} \frac{\|\mathbf{x}_0\|}{8c\beta_{\text{nl}}(1 + \|K\|)} \right\}.$$

Note that, due to $\alpha/d_x \leq 1$, $\beta_{\text{nl}}^2 \geq 1$, [Eq. \(C.7\)](#) holds as soon as $\|\mathbf{x}_0\| \leq c^2 r_{\text{nl}}^4$. By [Lemma C.2](#), it then follows that:

$$\begin{aligned} J_{\text{nl}}^{(H)}(K \mid \gamma, \mathbf{x}_0) &\geq \min \left\{ \frac{\alpha \|\mathbf{x}_0\|^2}{4d_x^2} \cdot \min \left\{ \frac{1}{2} J_{\text{lin}}(K \mid \gamma), c \right\}, \sqrt{\frac{\alpha}{d_x}} \frac{\|\mathbf{x}_0\|}{8c\beta_{\text{nl}}(1 + \|K\|)} \right\} \\ &\geq \min \left\{ \frac{\alpha \|\mathbf{x}_0\|^2}{8d_x^2} \min \{J_{\text{lin}}(K \mid \gamma), c\}, \sqrt{\frac{\alpha}{d_x}} \frac{\|\mathbf{x}_0\|}{8c\beta_{\text{nl}}(1 + \|K\|)} \right\} \\ &= \frac{\alpha \|\mathbf{x}_0\|^2}{8d_x^2} \min \left\{ J_{\text{lin}}(K \mid \gamma), c, \left(\frac{\alpha \|\mathbf{x}_0\|^2}{8d_x^2} \right)^{-1} \sqrt{\frac{\alpha}{d_x}} \frac{\|\mathbf{x}_0\|}{8c\beta_{\text{nl}}(1 + \|K\|)} \right\}. \end{aligned}$$

Simplifying, the last term gives

$$\left(\frac{\alpha \|\mathbf{x}_0\|^2}{8d_x^2} \right)^{-1} \sqrt{\frac{\alpha}{d_x}} \frac{\|\mathbf{x}_0\|}{8c\beta_{\text{nl}}(1 + \|K\|)} = \frac{1}{c \|\mathbf{x}_0\|} \cdot \left(\frac{d_x^3}{\alpha} \right)^{1/2} \frac{1}{64\beta_{\text{nl}}(1 + \|K\|)} \geq \frac{1}{c \|\mathbf{x}_0\|} \cdot \frac{d_x}{64\beta_{\text{nl}}(1 + \|K\|)},$$

<u>Noisy Binary Search</u>	<u>Noisy Random Search</u>
<p>Require: \bar{f}_1 and \bar{f}_2 as defined in Lemma C.6.</p> <p>Initialize: $b_0 \leftarrow 0, u_0 \leftarrow 1, c \geq \bar{f}_2 + 2\alpha$ for $\alpha := \alpha = \min\{ f_1 - \bar{f}_1 , f_2 - \bar{f}_2 \}$, $\varepsilon \in (0, \alpha/2)$</p> <p>For $t = 1, \dots$</p> <ol style="list-style-type: none"> 1. Query $a_t \leftarrow \varepsilon\text{-Eval}(f, x_t, c)$ where $x_t = \frac{b_t + u_t}{2}$ 2. If $a_t > \bar{f}_2 + \varepsilon$, update $u_t \leftarrow x_t$ and $b_t \leftarrow b_{t-1}$ 3. Else if, $a_t < \bar{f}_1 + \varepsilon$, update $b_t \leftarrow x_t$ and $u_t \leftarrow u_{t-1}$ 4. Else, break and return x_t 	<p>Require: \bar{f}_1, \bar{f}_2 as defined in Lemma C.7.</p> <p>Initialize: $c \geq \bar{f}_2 + 2\alpha$ for $\alpha := \min\{ f_1 - \bar{f}_1 , f_2 - \bar{f}_2 \}$, $\varepsilon \in (0, \alpha/2)$</p> <p>For $t = 1, \dots$</p> <ol style="list-style-type: none"> 1. Sample x uniformly at random from $[0, 1]$ 2. Query $a \leftarrow \varepsilon\text{-Eval}(f, x, c)$ 3. If $a \in [\bar{f}_1, \bar{f}_2]$, break and return x,

Figure 2: Algorithms for 1 dimensional search used as subroutines for Step 4 in the discount annealing algorithm.

where we use $d_x/\alpha \geq 1$. Thus, for

$$\|\mathbf{x}_0\| \leq \frac{d_x}{64c^2\beta_{\text{nl}}(1 + \|K\|)},$$

the third term in the minimum is at most c , so that

$$J_{\text{nl}}^{(H)}(K \mid \gamma, \mathbf{x}_0) \geq \frac{\alpha\|\mathbf{x}_0\|^2}{8d_x^2} \min\{J_{\text{lin}}(K \mid \gamma), c\}.$$

Lastly, using [Lemma C.3](#),

$$\begin{aligned} J_{\text{nl}}^{(H)}(K \mid \gamma, r) &= \frac{d_x}{r^2} \cdot \mathbb{E}_{\mathbf{x}_0 \sim r \cdot \mathcal{S}^{d_x-1}} J_{\text{nl}}^{(H)}(K \mid \gamma, \mathbf{x}_0) \\ &\geq \mathbb{P}[x_0 \in \mathcal{X}_1] \cdot \frac{1}{8d_x^2} \min\{J_{\text{lin}}(K \mid \gamma), c\} \\ &\geq \frac{1}{80d_x^2} \min\{J_{\text{lin}}(K \mid \gamma), c\}. \end{aligned}$$

□

C.4 Search analysis

Lemma C.6. *Let $f : [0, 1] \rightarrow \mathbb{R} \cup \{\infty\}$ be a nondecreasing function over the unit interval. Then, given \bar{f}_1, \bar{f}_2 such that $[\bar{f}_1, \bar{f}_2] \subseteq [f_1, f_2]$ and for which there exist $[x_1, x_2] \subseteq [0, 1]$ such that for all $x' \in [x_1, x_2]$, $f(x') \in [\bar{f}_1, \bar{f}_2]$, binary search as defined in [Figure 2](#) returns a value $x_\star \in [0, 1]$ such that $f(x_\star) \in [\bar{f}_1, \bar{f}_2]$ in at most $\lceil \log_2(1/\Delta) \rceil$ many iterations where $\Delta = x_2 - x_1$.*

Lemma C.7. *Let $f : [0, 1] \rightarrow \mathbb{R} \cup \{\infty\}$ be a function over the unit interval. Then, given given \bar{f}_1, \bar{f}_2 such that $[\bar{f}_1, \bar{f}_2] \subseteq [f_1, f_2]$ and for which there exist $[x_1, x_2] \subseteq [0, 1]$ such that for all $x' \in [x_1, x_2]$, $f(x') \in [\bar{f}_1, \bar{f}_2]$, with probability $1 - \delta$, noisy random search as defined in [Figure 2](#) returns a value $x_\star \in [0, 1]$ such that $f(x_\star) \in [f_1, f_2]$ in at most $1/\Delta \log(1/\delta)$ many iterations where $\Delta = x_2 - x_1$.*

The analysis of the correctness and runtime of this classical algorithms for search problems in 1 dimension are standard. We omit the proofs for the sake of concision.

D Additional Experiments

D.1 Cart-pole dynamics

The state of the cart-pole is given by $\mathbf{x} = (x, \theta, \dot{x}, \dot{\theta})$, where x and \dot{x} denote the horizontal position and velocity of the cart, and θ and $\dot{\theta}$ denote the angular position and velocity of the pole, with $\theta = 0$ corresponding to the upright equilibrium. The control input $u \in \mathbb{R}$ corresponds to the horizontal force applied to the cart. The continuous time dynamics are given by

$$\begin{bmatrix} m_p + m_c & -m_p l \cos(\theta) \\ -m_p l \cos(\theta) & m_p l^2 \end{bmatrix} \begin{bmatrix} \ddot{x} \\ \ddot{\theta} \end{bmatrix} = \begin{bmatrix} u - m_p l \sin(\theta) \dot{\theta}^2 \\ m_p g l \sin(\theta) \end{bmatrix}$$

where m_p denotes the mass of the pole, m_c the mass of the cart, l the length of the pendulum, and g acceleration due to gravity. For our experiments, we set all parameters to unity, i.e. $m_p = m_c = l = g = 1$.

D.2 \mathcal{H}_∞ synthesis

Consider the following linear system

$$\mathbf{x}_{t+1} = A\mathbf{x}_t + B\mathbf{u}_t + \mathbf{w}_t, \quad \mathbf{z}_t = \begin{bmatrix} Q^{1/2}\mathbf{x}_t \\ R^{1/2}\mathbf{u}_t \end{bmatrix}, \quad (\text{D.1})$$

where \mathbf{w} denotes an additive disturbance to the state transition, and \mathbf{z} denotes the so-called performance output. Notice that $\|\mathbf{z}_t\|_2^2 = \mathbf{x}_t^\top Q \mathbf{x}_t + \mathbf{u}_t^\top R \mathbf{u}_t$. The \mathcal{H}_∞ optimal controller minimizes the \mathcal{H}_∞ -norm of the closed-loop system from input \mathbf{w} to output \mathbf{z} , i.e. the smallest η such that

$$\sum_{t=0}^T \|\mathbf{z}_t\|_2^2 \leq \eta \sum_{t=0}^T \|\mathbf{w}_t\|_2^2$$

holds for all $\mathbf{w}_{0:T}$, all T , and $\mathbf{x}_0 = 0$. In essence, the \mathcal{H}_∞ optimal controller minimizes the effect of the worst-case disturbance on the cost. In this setting, additive disturbances \mathbf{w} serve as a crude (and unstructured) proxy for modeling error. We synthesize the \mathcal{H}_∞ optimal controller using Matlab's `hinfsyn` function. For the system (D.1) with perfect state observation, we require only static state feedback $\mathbf{u}_t = K\mathbf{x}_t$ to implement the controller.

D.3 Difference between LQR and discount annealing as a function of discount

In Figure 3 we plot the error $\|K_{\text{pg}}^*(\gamma) - K_{\text{lin}}^*(\gamma)\|_F$ between the policy $K_{\text{pg}}^*(\gamma)$ returned by policy gradient and the optimal policy $K_{\text{lin}}^*(\gamma)$ for the (damped) linearized system, as a function of the discount γ during the discount annealing process. Observe that for small radius of the ball of initial conditions ($r = 0.05$), the optimal controller from policy gradient remains very close to the exact optimal controller for the linearized system; however, for larger radius ($r = 0.7$) the controller from policy gradient differs significantly.




Figure 3: Error between the policy returned by policy gradient and optimal LQR policy for the (damped) linearized system during discount annealing, for two different radius values, $r = 0.05$ and $r = 0.7$, of the ball of initial conditions. Five independent trials are plotted for each radius value. The values across trials highly overlap.