



Small-disturbance input-to-state stability of perturbed gradient flows: Applications to LQR problem

Leilei Cui^{a,*}, Zhong-Ping Jiang^a, Eduardo D. Sontag^b

^a Department of Electrical and Computer Engineering, New York University, Brooklyn, NY, USA

^b Department of Electrical and Computer Engineering and Department of BioEngineering, Northeastern University, Boston, MA, USA

ARTICLE INFO

Keywords:

Input-to-state stability (ISS)
Learning-based control
Policy optimization
Gradient systems
Linear quadratic regulator (LQR)

ABSTRACT

This paper studies the effect of perturbations on the gradient flow of a general nonlinear programming problem, where the perturbation may arise from inaccurate gradient estimation in the setting of data-driven optimization. Under suitable conditions on the objective function, the perturbed gradient flow is shown to be small-disturbance input-to-state stable (ISS), which implies that, in the presence of a small-enough perturbation, the trajectories of the perturbed gradient flow must eventually enter a small neighborhood of the optimum. This work was motivated by the question of robustness of direct methods for the linear quadratic regulator problem, and specifically the analysis of the effect of perturbations caused by gradient estimation or round-off errors in policy optimization. We show small-disturbance ISS for three of the most common optimization algorithms: standard gradient flow, natural gradient flow, and Newton gradient flow.

1. Introduction

Gradient-based optimization of loss functions constitutes a key tool in contemporary machine learning. Thus, the theoretical analysis of convergence to the minima of loss functions, in gradient-like iterations and/or in their continuous analogue, gradient-like flows (viewed as the limit of discrete-time gradient descent algorithms with infinitesimally small step size) have attracted considerable attention from both academic and industrial researchers. Besides convergence under ideal no-noise situations, a useful optimization algorithm should be capable of finding a near-optimal solution while degrading elegantly in the face of perturbations that might arise from noisy measurements of experimental data, arithmetic rounding errors due to numerical computation, numerically approximating the gradient from data through two-point estimates, discretization error when solving ordinary differential equations, or even early stopping when estimating gradients in a hierarchical learning setup [1–5]. Especially in the setting of data-driven optimization, the analytical form of the gradient is typically unknown, and consequently the gradient has to be numerically approximated through sampling and experiments, which unavoidably introduces perturbations to the gradient iteration or flow. In computer science foundations of optimization theory, similarly, noisy or error-prone operations such as inexact or stochastic gradient computations have led to the introduction of the concept of “reproducibility in optimization” which is concerned conceptually with the same issues [6]. One could also view adversarial attacks on neural network training as

affecting gradient computations in “backpropagation” algorithms, and the effect of disturbances in that context has been the subject of recent work [7]. In summary, both the convergence and robustness properties of gradient descent should be theoretically analyzed in the presence of perturbations. Mathematically, gradient flows are more amenable to mathematical analysis than discrete iterations, so they are the main object of study in this paper.

In order to formulate precisely the effect of perturbations on gradient flows, we employ as in [5] the formalism of input-to-state stability (ISS) introduced originally in [8] (see for example [9] for an exposition).

The key to proving ISS for perturbed gradient flows is to verify a *Polyak-Lojasiewicz (PL)* type of condition [10–12] on the loss function to be optimized, meaning that the gradient of the loss function should not be “too small” compared to the loss. Roughly (precise definitions to be given) if we wish to minimize a continuously differentiable (but not necessarily convex) function $J(z)$ on a domain \mathcal{Z} and if a global minimizer z^* exists, then we would like that, for some continuous function $\kappa : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ which satisfies that $\kappa(0) = 0$ and $\kappa(r) > 0$ for all $r > 0$ (a “positive definite” function), there should hold an estimate of the form $\|\nabla J(z)\| \geq \kappa(J(z) - J(z^*))$ valid for all z in \mathcal{Z} . The classical PL condition is often stated in a slightly different “semi-global” form, by requiring the existence, for each r , of a constant c_r such that $\|\nabla J(z)\| \geq c_r(J(z) - J(z^*))$ for every z in the sublevel set $\{z \mid J(z) \leq r\}$.

* Corresponding author.

E-mail address: l.cui@nyu.edu (L. Cui).

For a coercive function \mathcal{J} , this is equivalent to the above estimate using positive definite functions.

Sometimes, the PL condition is stated globally, that is, with a constant c which is independent of r (in other words, κ can be picked as a linear function), but such a global condition is too strong for many applications, including the one to be pursued here. On the other hand, one could think of stronger forms of the PL condition, weaker than the existence of a linear function but stronger than merely requiring κ to be positive definite. One particularly useful strengthening is to ask that κ be a function of class \mathcal{K} , that is, that it be a strictly increasing function. For example, one could take $\kappa(r) = \frac{ar}{b+r}$: note that this function saturates, in the sense that it approaches a finite limit as $r \rightarrow +\infty$. This is stronger than asking that κ be only positive definite, as illustrated by $\kappa(r) = \frac{ar}{(b+r)^2}$ which is positive definite but is not of class \mathcal{K} . An even further strengthening would be to ask that κ be of class \mathcal{K}_∞ , meaning that κ does not saturate, $\kappa(r) \rightarrow \infty$ as $r \rightarrow \infty$, as for example when κ is linear. Estimates with κ of class \mathcal{K}_∞ lead to ISS estimates for perturbed gradient flows, as discussed in [5], and a similar proof to that in [5] can be used to show that estimates with $\kappa(r)$ only positive definite lead to the weaker property of “integral ISS (iISS)” [13,14] for perturbed gradient flows. We may call the intermediate type of PL estimate, in which κ is required to be of class \mathcal{K} (a stronger property than positive definiteness, but not as strong as class \mathcal{K}_∞) a *CJS-PL* (“comparison just saturated”) estimate. It turns out that CJS-PL estimates are exactly what is required in order to establish “small-disturbance” ISS as studied in [15,16]. To be precise, we will show that when the objective function is coercive (the value of the objective function blows up when the decision variable approaches the boundary of \mathcal{Z}) and the CJS-PL condition holds, the perturbed gradient flow is small-disturbance ISS. This implies that the trajectories of the perturbed gradient flow will eventually enter a small neighborhood of the optimal solution, as long as the perturbation is sufficiently small. In addition, the size of the neighborhood is (in a nonlinear manner) proportional to the magnitude of the perturbation. In the application that motivated this work, the linear regulator problem (see below), CJS-PL is the correct notion to use, and we believe that this notion might be of more general applicability in optimization problems as well.

We should remark that generalizations of the PL condition, and relations to ISS types of properties, can be found in other recent work. This includes [17], which studies the gradient minimization of a function J_q on Euclidean space, where the parameter q represents time-varying uncertainty. In that paper, an ISS property is established with respect to the rate of change of the parameter q , essentially showing differential ISS (DISS) [18]. Extremum-seeking controllers based on gradient flows and an ISS property with respect to disturbances, specifically for an integrator and a kinematic unicycle, are designed and analyzed in [2]; in that paper the domain is a closed submanifold of an Euclidean space. In [1] one finds results on gradient flows that are ISS with respect to additive errors, but assuming a “convex-concave” property for the loss function, and in [3] the authors solve an output regulation problem for switched linear systems, and show an ISS property for gradient flows with respect to unknown disturbances acting on the plant.

We now turn to the main motivation for this work. Reinforcement learning (RL) is an active research field in which gradient-based optimization plays a pivotal role [19, Chapter 13]. In the setting of RL, an agent interacts continuously with an unknown environment, and iteratively optimizes a performance index by collecting data from the environment. By adopting gradient-based optimization methods, various policy optimization (PO) algorithms have been developed, such as actor-critic methods [20], deep deterministic policy gradient [21], and trust region policy optimization [22]. The critical strategy of the policy optimization methods is to parameterize the policy by universal approximations and update the parameters of the policy along the gradient descent direction of the performance index.

Starting in the early 1960s with the work of Kalman, the linear quadratic regulator (LQR) problem was shown to be theoretically

tractable, and has become a widely utilized tool for optimal control and feedback design in engineering applications.

In the classical approach, the (infinite-horizon) LQR problem relies upon the solution of a Riccati equation. In 1970, Athans and Levine [23] introduced the idea of a direct gradient descent computation of optimal feedback gains, a procedure which can be interpreted as a form of RL. Thus, the LQR problem offers an ideal benchmark for better understanding policy optimization methods in the RL field, as one can compare solutions to the known optimal solution, and analysis of gradient methods can take advantage of theory developed for LQR. For policy optimization in the LQR problem, the objective function is a cumulative quadratic function of the state and control inputs, the control policy is parameterized as a linear function (feedback) of the state, and the admissible set, consisting of all the stabilizing control gains, is an open subset of an Euclidean space. As investigated for example in [4,24,25], the gradient of the objective function can be computed by using a Lyapunov equation that depends on the system matrices. Nevertheless, if precise system knowledge is unavailable, as in the setting of model-free RL, the gradient has to be numerically approximated through sampling and experiments. For example, by utilizing the approximate dynamic programming technique [26,27], the Lyapunov equation was solved by data-driven methods in [28–30]. In [4,31,32], the gradient is directly calculated by the finite differences method [33, Section 7.1], based on the change in function values in response to small perturbations near a given point. For these data-driven methods, a gradient estimation error is inevitable due to noisy data and insufficient samples. Therefore, the robustness analysis of the policy optimization algorithm in the presence of perturbations is critical for efficient learning, and lays the foundations for better understanding RL algorithms.

Our main result will be that, for the LQR problem, the loss function is coercive and satisfies the CJS-PL property, and therefore, by the results in the first part of the paper, we conclude that the perturbed standard gradient flow is small-disturbance ISS. We also show that two variants of gradient flows, natural gradient flows and Newton gradient flows, are small-disturbance ISS. The new contribution is to establish the CJS-PL property for the LQR problem. This considerably extends previous work [4,24] that only showed a semiglobal estimate (and thus would imply merely iISS). In [5], it was mistakenly stated that the magnitude of the gradient is lower bounded by a \mathcal{K}_∞ -function, which is a stronger property. This is incorrect. Indeed, take a one-dimensional linear system with scalar inputs and assume that all constants in the system and cost function are equal to one. Then the loss function is $J(z) = \frac{z^2+1}{2(z-1)}$, so that its gradient is $J'(z) = \frac{z^2-2z-1}{2(z-1)^2}$. The domain of J is the open set $(1, \infty)$. We claim that there is no function κ of class \mathcal{K}_∞ such that $|J'(z)| \geq \kappa(J(z) - J(z^*))$, where z^* is the global minimizer of J . Indeed, as $z \rightarrow \infty$ we would have that $J(z) \rightarrow \infty$, so, as κ is of class \mathcal{K}_∞ , also $\kappa(J(z) - J(z^*)) \rightarrow \infty$. However, the left-hand side $|J'(z)|$ is bounded, and in fact converges to $1/2$, showing a contradiction.

To summarize, the contributions of the paper are as follows. First, we provide a Lyapunov-like necessary and sufficient condition for small-disturbance ISS. Second, under assumptions of coercivity and the CJS-PL property, we use the Lyapunov characterization to show that the perturbed gradient flow for a general constrained nonlinear programming problem is small-disturbance ISS. Finally, we show the CJS-PL property for the LQR loss function, which in turn then implies that the standard gradient flow, natural gradient flow, and Newton gradient flow, are all small-disturbance ISS. The remaining contents of the paper are organized as follows. The notations and preliminaries are introduced in Section 2. In Section 3, the concept of small-disturbance ISS is reviewed, followed by a necessary and sufficient condition. Section 4 introduces the perturbed gradient flow for a general constrained nonlinear programming problem over an open admissible set, and it is shown that the perturbed gradient flow is small-disturbance ISS under appropriate conditions on the loss function. In Section 5, we study the CSJ-PL property for the LQR problem, and three different kinds of the perturbed gradient flows for LQR are shown to be small-disturbance ISS. Some concluding remarks are given in Section 6.

2. Notations and preliminaries

In this paper, \mathbb{R} (\mathbb{R}_+) denotes the set of (nonnegative) real numbers. \mathbb{P}^n denotes the set of n -dimensional real symmetric and positive definite matrices. $\lambda_{\min}(\cdot)$ and $\lambda_{\max}(\cdot)$ denote the minimal and maximal eigenvalues of a real symmetric matrix, respectively. $\text{Tr}(\cdot)$ denotes the trace of a square matrix. $\|\cdot\|$ denotes the spectral norm of a matrix or Euclidean norm of a vector, and $\|\cdot\|_F$ denotes the Frobenius norm of a matrix. \mathcal{L}^∞ ($\mathcal{L}^\infty_{\text{loc}}$) denotes the set of measurable and locally essentially bounded functions $w : \mathbb{R}_+ \rightarrow \mathbb{R}^n$ ($K : \mathbb{R}_+ \rightarrow \mathbb{R}^{m \times n}$), endowed with the essential supremum norm $\|w\|_\infty = \text{ess sup}_{s \in \mathbb{R}_+} \|w(s)\|$ ($\|K\|_\infty = \text{ess sup}_{s \in \mathbb{R}_+} \|K(s)\|_F$). w_t denotes the truncation of w at t , that is, $w_t(s) = w(s)$ if $s \leq t$, and $w_t(s) = 0$ if $s > t$. I_n denotes the n -dimensional identity matrix. Id denotes the identity function. For any $K_1, K_2 \in \mathbb{R}^{m \times n}$ and $Y \in \mathbb{P}^n$, define the inner product $\langle K_1, K_2 \rangle_Y = \text{Tr}(K_1^T K_2 Y)$. In addition, to simplify the notation, we denote $\langle K_1, K_2 \rangle = \langle K_1, K_2 \rangle_{I_n}$. Recall that for any $K \in \mathbb{R}^{m \times n}$, $\|K\|_F^2 = \langle K, K \rangle$. For any two real symmetric matrices A and B , $A > B$ means that $A - B$ is positive definite, and $A \geq B$ means that $A - B$ is positive semidefinite.

In the remainder of this section we gather definitions and technical results which we shall use throughout the rest of the paper. A number of these results appear across the literature.

Definition 2.1 (Definitions 2.5 and 2.4.2 in [34]). A function $\alpha : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a \mathcal{K} -function if it is continuous, strictly increasing, and vanishes at zero. For any $d > 0$, a function $\alpha : [0, d) \rightarrow \mathbb{R}_+$ is a $\mathcal{K}_{[0, d)}$ -function if it is continuous, strictly increasing, and vanishes at zero. A function $\alpha : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a \mathcal{K}_∞ -function if it is a \mathcal{K} -function and also satisfies $\alpha(r) \rightarrow \infty$ as $r \rightarrow \infty$. A function $\beta : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a \mathcal{KL} -function if for any fixed $t \geq 0$, $\beta(\cdot, t)$ is a \mathcal{K} -function, and for any fixed $r \geq 0$, $\beta(r, \cdot)$ is decreasing and $\beta(r, t) \rightarrow 0$ as $t \rightarrow \infty$.

Lemma 2.1 (The Cyclic Property of the Trace, Equation (16) in [35]). For any $X, Y, Z \in \mathbb{R}^{n \times n}$, $\text{Tr}(XYZ) = \text{Tr}(ZXY) = \text{Tr}(YZX)$.

Lemma 2.2 (Trace Inequality [36]). Let $S \in \mathbb{R}^{n \times n}$ be real symmetric and $P \in \mathbb{R}^{n \times n}$ be real symmetric and positive semidefinite. Then,

$$\lambda_{\min}(S) \text{Tr}(P) \leq \text{Tr}(SP) \leq \lambda_{\max}(S) \text{Tr}(P).$$

Lemma 2.3 (Cauchy–Schwarz Inequality). For any $K_1, K_2 \in \mathbb{R}^{m \times n}$, $R \in \mathbb{P}^n$, and $Y \in \mathbb{P}^n$, we have

$$\langle K_1, RK_2 \rangle_Y \leq \sqrt{\langle K_1, RK_1 \rangle_Y} \sqrt{\langle K_2, RK_2 \rangle_Y}. \quad (1)$$

Proof. Inequality (1) is obviously true if $K_2 = 0$. If $K_2 \neq 0$, define K_3 as

$$K_3 = K_1 - \frac{\langle K_1, RK_2 \rangle_Y}{\langle K_2, RK_2 \rangle_Y} K_2. \quad (2)$$

It is clear that $\langle K_3, RK_2 \rangle_Y = 0$. Therefore, by plugging (2) into $\langle K_1, RK_1 \rangle_Y$, we can obtain

$$\langle K_1, RK_1 \rangle_Y = \langle K_3, RK_3 \rangle_Y + \frac{\langle K_1, RK_2 \rangle_Y^2}{\langle K_2, RK_2 \rangle_Y^2} \langle K_2, RK_2 \rangle_Y \geq \frac{\langle K_1, RK_2 \rangle_Y^2}{\langle K_2, RK_2 \rangle_Y}. \quad (3)$$

Hence, (1) readily follows from (3). \square

Lemma 2.4. The map $h(v) = \frac{1}{1+\|v\|} v : \mathbb{R}^m \rightarrow \mathcal{W} := \{w \in \mathbb{R}^m \mid \|w\| < 1\}$ is a homeomorphism.

Proof. For any $w \in \mathcal{W}$, let $g(w) = \frac{1}{1-\|w\|} w$. Clearly, for any $v \in \mathbb{R}^m$, $g(h(v)) = v$ and for any $w \in \mathcal{W}$, $h(g(w)) = w$. Hence, g is the inverse function of h , i.e. $g = h^{-1}$. Since both h and g are continuous, h is a homeomorphism. \square

Lemma 2.5. Suppose $\omega_1, \omega_2 : \mathbb{R}^n \rightarrow \mathbb{R}$ are continuous, positive definite with respect to χ^* , and radially unbounded. Then, there exist \mathcal{K}_∞ -functions ρ_1 and ρ_2 such that

$$\rho_1(\omega_2(\chi)) \leq \omega_1(\chi) \leq \rho_2(\omega_2(\chi)), \quad \forall \chi \in \mathbb{R}^n.$$

Proof. The proof follows from [5, Proposition 2.6] by considering the open subset as \mathbb{R}^n and the compact set as $\{\chi^*\}$. \square

Lemma 2.6 (Weak triangle inequality in [37]). For any \mathcal{K} -function α , any \mathcal{K}_∞ -function ρ , and any nonnegative real numbers a and b , we have

$$\alpha(a + b) \leq \alpha \circ (\text{Id} + \rho)(a) + \alpha \circ (\text{Id} + \rho^{-1})(b).$$

Lemma 2.7 (Theorem 18 in [38]). If $A \in \mathbb{R}^{n \times n}$ is Hurwitz, then the Lyapunov equation

$$A^T P + PA + Q = 0$$

has a unique solution for any $Q \in \mathbb{R}^{n \times n}$, and the solution can be expressed as

$$P = \int_0^\infty e^{A^T t} Q e^{At} dt.$$

3. Small-disturbance input-to-state stability

Let S denote an open subset of \mathbb{R}^n , which will be called the admissible set of states. Consider the following nonlinear system

$$\dot{\chi}(t) = f(\chi(t), w(t)), \quad (4)$$

where $f : S \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ is a smooth function, and inputs $w : \mathbb{R}_+ \rightarrow \mathbb{R}^m$ are measurable and locally essentially bounded functions. Assume the unforced system has an equilibrium χ^* , i.e. $f(\chi^*, 0) = 0$.

Definition 3.1 (Definition 2.1 in [5]). A function $\mathcal{V} : S \rightarrow \mathbb{R}_+$ is a size function for (S, χ^*) if \mathcal{V} is

1. continuous;
2. positive definite with respect to χ^* , i.e. $\mathcal{V}(\chi^*) = 0$ and $\mathcal{V}(\chi) > 0$ for all $\chi \neq \chi^*$, $\chi \in S$;
3. coercive, i.e. for any sequence $\{\chi_k\}_{k=0}^\infty$, $\chi_k \rightarrow \partial S$ or $\|\chi_k\| \rightarrow \infty$, it holds that $\mathcal{V}(\chi_k) \rightarrow \infty$, as $k \rightarrow \infty$.

Definition 3.2 ([15, 16]). System (4) is small-disturbance input-to-state stable (ISS) if there exist a size function \mathcal{V} , a constant $d > 0$ (possibly ∞), a \mathcal{KL} -function β , and a $\mathcal{K}_{[0, d)}$ -function γ , such that for all inputs w essentially bounded by d (i.e. $\|w\|_\infty < d$), and all initial states $\chi(0) \in S$, $\chi(t)$ remains in S and satisfies

$$\mathcal{V}(\chi(t)) \leq \beta(\mathcal{V}(\chi(0)), t) + \gamma(\|w\|_\infty), \quad \forall t \geq 0. \quad (5)$$

As shown in [39], by causality, the same definition would result if one would replace $\|w\|_\infty$ by $\|w_t\|_\infty$ in (5).

Definition 3.3. A continuously differentiable function $\mathcal{V} : S \rightarrow \mathbb{R}$ is a small-disturbance ISS-Lyapunov function for system (4) if

1. \mathcal{V} is a size function for (S, χ^*) ;
2. there exist a \mathcal{K} -function α_1 and a continuous and positive definite function α_2 such that if $\|\mu\| \leq \alpha_1(\mathcal{V}(\chi))$,

$$\nabla \mathcal{V}(\chi)^T f(\chi, \mu) \leq -\alpha_2(\mathcal{V}(\chi)). \quad (6)$$

Theorem 3.1. System (4) is small-disturbance ISS if and only if it admits a small-disturbance ISS-Lyapunov function.

Proof. Sufficiency: This is an adaptation of the proof of the analogous result for the ISS property [8]. Let $d = \sup_{r \in \mathbb{R}_+} \alpha_1(r)$. For any inputs w with $\|w\|_\infty < d$, define the sublevel set $S_c = \{\chi \in S \mid \mathcal{V}(\chi) \leq c\}$,

where $c = \alpha_1^{-1}(\|w\|_\infty)$. Borrowing techniques similar to those in the proof of [8, Theorem 1], we can show that S_c is forward invariant, i.e. if $\chi(t_0) \in S_c$ for some $t_0 \geq 0$, then $\chi(t) \in S_c$ for all $t \geq t_0$.

Now, let $t_1 = \inf\{t \in \mathbb{R}_+ | \chi(t) \in S_c\} \leq \infty$. Therefore, for any $t \geq t_1$, we have

$$\mathcal{V}(\chi(t)) \leq \alpha_1^{-1}(\|w\|_\infty). \quad (7)$$

For $t < t_1$, $\alpha_1(\mathcal{V}(\chi(t))) \geq \|w\|_\infty$, which implies that

$$\frac{d\mathcal{V}(\chi(t))}{dt} \leq -\alpha_2(\mathcal{V}(\chi(t))), \quad \forall t < t_1.$$

Hence, $\mathcal{V}(\chi(t)) \leq \mathcal{V}(\chi(0))$, $\forall t < t_1$. By the comparison principle [40, Lemma 4.4], there exists a \mathcal{KL} -function β such that

$$\mathcal{V}(\chi(t)) \leq \beta(\mathcal{V}(\chi(0)), t), \quad \forall t < t_1. \quad (8)$$

Combining (7) and (8), the small-disturbance ISS property (5) follows readily with $\gamma = \alpha_1^{-1}$.

Necessity: We first prove the case when $S = \mathbb{R}^n$. Reparameterize the input as

$$w(t) = \frac{d}{d + \|v(t)\|} v(t) =: h(v(t)).$$

It is shown in Lemma 2.4 that h is a homeomorphism from \mathbb{R}^m to $\mathcal{W} := \{w \in \mathbb{R}^m | \|w\| < d\}$ with $v(t) = h^{-1}(w(t)) = \frac{d}{d - \|w(t)\|} w(t)$. With the input change, we have

$$\dot{\chi}(t) = f(\chi(t), h(v(t))) =: f_1(\chi(t), v(t)). \quad (9)$$

Since γ is a $\mathcal{K}_{[0,d]}$ -function and $\gamma_1(r) = \frac{dr}{d+r}$ is a \mathcal{K} -function with the range $[0, d)$, $\gamma_2 = \gamma \circ \gamma_1$ is a \mathcal{K} -function. According to (5), it holds

$$\mathcal{V}(\chi(t)) \leq \beta(\mathcal{V}(\chi(0)), t) + \gamma_2(\|v\|_\infty). \quad (10)$$

Since \mathcal{V} is a size function for (\mathbb{R}^n, χ^*) , according to Lemma 2.5, there exist \mathcal{K}_∞ -functions ρ_1 and ρ_2 , such that

$$\rho_1(\|\chi - \chi^*\|) \leq \mathcal{V}(\chi) \leq \rho_2(\|\chi - \chi^*\|), \quad \forall \chi \in \mathbb{R}^n. \quad (11)$$

Plugging (11) into (10) yields

$$\|\chi(t) - \chi^*\| \leq \rho_1^{-1}[\beta(\rho_2(\|\chi(0) - \chi^*\|), t) + \gamma_2(\|v\|_\infty)].$$

By Lemma 2.6, there exist a \mathcal{KL} -function β_1 and a \mathcal{K} -function γ_3 , such that

$$\|\chi(t) - \chi^*\| \leq \beta_1(\|\chi(0) - \chi^*\|, t) + \gamma_3(\|v\|_\infty).$$

Hence, system (9) is ISS with respect to v . According to [39, Theorem 1], there exists an ISS-Lyapunov function \mathcal{V}_1 , such that

$$\rho_3(\|\chi - \chi^*\|) \leq \mathcal{V}_1(\chi) \leq \rho_4(\|\chi - \chi^*\|), \quad \forall \chi \in \mathbb{R}^n,$$

and

$$\nabla \mathcal{V}_1(\chi)^T f_1(\chi, v) \leq -\rho_6(\|\chi - \chi^*\|), \quad (12)$$

for any $\chi \in \mathbb{R}^n$ and any $v \in \mathbb{R}^m$ satisfying $\|v\| \leq \rho_5(\|\chi - \chi^*\|)$, where ρ_i ($i = 3, 4, 5$) are \mathcal{K}_∞ -functions, and ρ_6 is a \mathcal{K} -function. This, in turn, implies that if $\|w\| \leq \gamma_1 \circ \rho_5 \circ \rho_4^{-1}(\mathcal{V}_1(\chi))$, Eq. (12) holds. Since $\gamma_1 \circ \rho_5 \circ \rho_4^{-1}$ is a \mathcal{K} -function with the range $[0, d)$, we obtain that \mathcal{V}_1 is a small-disturbance ISS-Lyapunov function.

Next, we prove the general case when S is an open subset of \mathbb{R}^n . By assumption, for the unforced system $\dot{\chi}(t) = f(\chi(t), 0)$, χ^* is an asymptotically stable point and the domain of stability is S . By [41, Theorem 2.2] and the Brown–Stallings Theorem [42], S is diffeomorphic to \mathbb{R}^n . We denote by $\varphi : S \rightarrow \mathbb{R}^n$ the diffeomorphism, $\zeta = \varphi(\chi)$, and $\zeta^* = \varphi(\chi^*)$. As a consequence, we have

$$\dot{\zeta}(t) = J_\varphi(\varphi^{-1}(\zeta(t)))f(\varphi^{-1}(\zeta(t)), w(t)) =: f_2(\zeta(t), w(t)), \quad (13)$$

where $J_\varphi(\chi)$ is the Jacobian matrix of $\varphi(\chi)$. Since system (4) is small-disturbance ISS over $S \times \mathbb{R}^m$, system (13) is small-disturbance ISS over $\mathbb{R}^n \times \mathbb{R}^m$. Then, according to the conclusion from the case of $S = \mathbb{R}^n$, there exists a small-disturbance ISS-Lyapunov function $\mathcal{V}_1 : \mathbb{R}^n \rightarrow \mathbb{R}_+$

for system (13) with respect to ζ^* . It readily follows that $\mathcal{V}_1 \circ \varphi : S \rightarrow \mathbb{R}_+$ is a small-disturbance ISS-Lyapunov function for system (4). Therefore, the necessity holds. \square

Remark 3.1. It should be mentioned that small-disturbance ISS is not equivalent to the notion of integral ISS [13]. Consider a scalar nonlinear system $\dot{\chi}(t) = -\frac{\chi(t)}{1+\chi(t)^2} + w(t)$. Differentiating the Lyapunov function $\mathcal{V}(\chi) = \log(1 + \chi^2)$ with respect to time yields $\dot{\mathcal{V}} \leq \frac{-2\chi^2}{(1+\chi^2)^2} + |w|$, which implies that the system is integral ISS [14]. However, for any arbitrarily small input $0 < \bar{w} < 0.5$, the trajectories of the system diverge whenever the initial conditions $\chi(0) > \frac{1+\sqrt{1-4\bar{w}^2}}{2\bar{w}}$. Hence, the system is not small-disturbance ISS.

4. Robustness analysis of perturbed gradient flows

Consider the following constrained nonlinear programming problem

$$\min_{z \in \mathcal{Z}} \mathcal{J}(z) \quad (14)$$

where \mathcal{Z} is an open subset of \mathbb{R}^n , which is called an admissible set of variables; $\mathcal{J} : \mathcal{Z} \rightarrow \mathbb{R}$ is an objective function with a global minimizer z^* .

Definition 4.1. A function $\mathcal{J} : \mathcal{Z} \rightarrow \mathbb{R}$ is a proper objective function if

1. $\mathcal{J}(z)$ is a smooth function;
2. $\mathcal{J}(z) - \mathcal{J}(z^*)$ is a size function for (\mathcal{Z}, z^*) ;
3. there exists a \mathcal{K} -function α_3 , such that $\|\nabla \mathcal{J}(z)\| \geq \alpha_3(\mathcal{J}(z) - \mathcal{J}(z^*))$ (CJS-PL estimate).

The perturbed gradient flow for (14) is

$$\dot{z}(t) = -\eta \nabla \mathcal{J}(z(t)) + e(t), \quad (15)$$

where $\eta > 0$ is a constant, and $e \in \mathcal{L}_\infty^n$ denotes the perturbation to the gradient flow. The perturbation $e(t)$ may arise from inaccurate gradient estimation for data-driven optimization or arithmetic rounding errors of numerical computation.

Theorem 4.1. If \mathcal{J} is a proper objective function, then system (15) is small-disturbance ISS.

Proof. We will prove that $\mathcal{V}_2(z) = \mathcal{J}(z) - \mathcal{J}(z^*)$ is a small-disturbance ISS-Lyapunov function. Firstly, notice that \mathcal{V}_2 is a size function for (\mathcal{Z}, z^*) . Then, by the Cauchy–Schwarz inequality and Young’s inequality, it holds

$$\begin{aligned} \nabla \mathcal{V}_2(z)^T (-\eta \nabla \mathcal{J}(z) + e) &= -\eta \|\nabla \mathcal{J}(z)\|^2 + \nabla \mathcal{J}(z)^T e \\ &\leq -\frac{\eta}{2} \|\nabla \mathcal{J}(z)\|^2 + \frac{1}{2\eta} \|e\|^2 \leq -\frac{\eta}{2} \alpha_3^2(\mathcal{V}_2(z)) + \frac{1}{2\eta} \|e\|^2, \end{aligned} \quad (16)$$

where the last inequality is a direct consequence of the CJS-PL property. Hence, if $\|e\| \leq \frac{\eta}{\sqrt{2}} \alpha_3(\mathcal{V}_2(z))$, it follows from (16) that

$$\nabla \mathcal{V}_2(z)^T (-\eta \nabla \mathcal{J}(z) + e) \leq -\frac{\eta}{4} \alpha_3^2(\mathcal{V}_2(z)).$$

Since α_3 is a \mathcal{K} -function, $\frac{\eta}{\sqrt{2}} \alpha_3$ is also a \mathcal{K} -function. Therefore, by Definition 3.3, $\mathcal{V}_2(z) = \mathcal{J}(z) - \mathcal{J}(z^*)$ is a small-disturbance ISS-Lyapunov function. According to Theorem 3.1, we conclude that system (15) is small-disturbance ISS. \square

Remark 4.1. Suppose that one has the classical PL inequality, namely for each $r > 0$, there is a $c_r > 0$ such that $\|\nabla \mathcal{J}(z)\|^2 \geq c_r(\mathcal{J}(z) - \mathcal{J}(z^*))$ for all z in the sublevel set $\{z \in \mathcal{Z} | \mathcal{J}(z) - \mathcal{J}(z^*) \leq r\}$. Then there is some positive definite function α_3 such that $\|\nabla \mathcal{J}(z)\| \geq \alpha_3(\mathcal{J}(z) - \mathcal{J}(z^*))$ for all $z \in \mathcal{Z}$. By following [14] and (16), it shows that the classical PL condition gives integral ISS.

Remark 4.2. As discussed in the Introduction, the gradient dominance condition given by the CJS-PL estimate is weaker than its counterpart in [5, Definition 4.1], where α_3 is required to be a \mathcal{K}_∞ -function. In addition, the CJS-PL estimate implies that the perturbed gradient flow (15) is integral ISS.

5. Application to LQR problem

5.1. Preliminaries of LQR

Consider the following linear time-invariant system

$$\dot{x}(t) = Ax(t) + Bu(t), \quad x(0) = x_0,$$

where $x(t) \in \mathbb{R}^n$ is the state; x_0 is the initial state; $u(t) \in \mathbb{R}^m$ is the control input; A and B are constant matrices with compatible dimensions. The continuous-time LQR aims at finding a state-feedback controller by solving the following optimal control problem

$$\min_{u \in \mathcal{L}^m_\infty} J_1(x_0, u) = \int_0^\infty x^T(t)Qx(t) + u^T(t)Ru(t)dt, \quad (17)$$

with $Q = Q^T \geq 0$, and $R = R^T > 0$. Under the assumption that (A, B) is stabilizable and (A, \sqrt{Q}) is observable, as shown in [38, Section 8.4], the optimal controller is

$$u^*(x(t)) = -K^*x(t), \quad K^* = R^{-1}B^T P^*, \quad (18)$$

where $P^* = (P^*)^T$ is the unique positive definite solution of the following algebraic Riccati equation (ARE)

$$A^T P^* + P^* A + Q - P^* B R^{-1} B^T P^* = 0. \quad (19)$$

Let $\mathcal{G} = \{K \in \mathbb{R}^{m \times n} \mid A - BK \text{ is Hurwitz}\}$ denote the admissible set of all stabilizing control gains. For any stabilizing controller $u(t) = -Kx(t)$, where $K \in \mathcal{G}$, and any nonzero initial state $x_0 \in \mathbb{R}^n$, the corresponding cost is

$$J_1(x_0, K) = \int_0^\infty x_0^T e^{(A-BK)t} (Q + K^T R K) e^{(A-BK)t} x_0 dt = x_0^T P_K x_0,$$

where $P_K = P_K^T$ is the unique positive definite solution of the following Lyapunov equation

$$(A - BK)^T P_K + P_K (A - BK) + Q + K^T R K = 0. \quad (20)$$

Since K^* is the optimal control gain and $P^* = P_{K^*}$ is the minimal cost matrix, by [38, page 382], it holds

$$J_1(x_0, K) = x_0^T P_K x_0 \geq x_0^T P^* x_0 = J_1(x_0, K^*), \quad \forall x_0 \in \mathbb{R}^n. \quad (21)$$

This implies that $P_K \geq P^*$ for all $K \in \mathcal{G}$.

Since the objective function J_1 of the LQR problem depends on the initial condition, we are motivated to study an equivalent optimization problem ($\min_{K \in \mathcal{G}} J_2(K)$), which is independent of the initial condition. For any initial state, an upper bound for $J_1(x_0, K)$ is

$$J_1(x_0, K) \leq \|x_0\|^2 \text{Tr}(P_K) = \|x_0\|^2 J_2(K),$$

where

$$J_2(K) := \text{Tr}(P_K), \quad (22)$$

which is independent of x_0 and is an analytic function [24, Proposition 3.2]. Since $P_K \geq P^*$ for any $K \in \mathcal{G}$, $J_2(K) \geq J_2(K^*)$. In addition, since $\text{Tr}(P_K) = \text{Tr}(P^*)$ implies $P_K = P^*$ and $K = K^*$, $J_2(K)$ has a unique minimum at K^* . Thus, the optimal control gain K^* can be obtained by the following policy optimization problem

$$\min_{K \in \mathcal{G}} J_2(K).$$

Before calculating $\nabla J_2(K)$, which denotes the gradient of $J_2(K)$ over the Euclidean space, let us define the matrix $Y_K \in \mathbb{P}^n$ as the solution of

$$(A - BK)Y_K + Y_K(A - BK)^T + I_n = 0. \quad (23)$$

It is noticed that according to [43, Lemma 3.18], $Y_K > 0$ for any $K \in \mathcal{G}$. In addition, Y^* is defined as the solution of (23) with K replaced by K^* . Since $A - BK$ is Hurwitz, by Lemma 2.7, Y_K can be expressed as

$$Y_K = \int_0^\infty e^{(A-BK)t} e^{(A-BK)^T t} dt. \quad (24)$$

Lemma 5.1. For any $K \in \mathcal{G}$, when it is perturbed by E with $K + E \in \mathcal{G}$ (recall that \mathcal{G} is an open set), the second-order Taylor series approximation of $J_2(K + E)$ is

$$J_2(K + E) = J_2(K) + 2\text{Tr}(E^T(RK - B^T P_K)Y_K) + \text{Tr}(E^T R E Y_K) + 2\text{Tr}(E^T(RK - B^T P_K)\Delta Y_K) + O(\|E\|_F^3), \quad (25)$$

where

$$\Delta Y_K := - \int_0^\infty e^{(A-BK)t} (B E Y_K + Y_K E^T B^T) e^{(A-BK)^T t} dt, \quad (26)$$

and $O(\|E\|_F^3)$ is the remainder of the approximation.

Proof. Firstly, we calculate Y_{K+E} when K is perturbed to $K + E$. Using (23), we have

$$(A - BK - BE)Y_{K+E} + Y_{K+E}(A - BK - BE)^T + I_n = 0. \quad (27)$$

Subtracting (23) from (27), we have

$$(A - BK - BE)(Y_{K+E} - Y_K) + (Y_{K+E} - Y_K)(A - BK - BE)^T - B E Y_K - Y_K E^T B^T = 0.$$

Let ΔY_K denote the first-order term in the Taylor expansion of Y_{K+E} , i.e. $Y_{K+E} = Y_K + \Delta Y_K + O(\|E\|_F^2)$. Then, ΔY_K satisfies

$$(A - BK)\Delta Y_K + \Delta Y_K(A - BK)^T - B E Y_K - Y_K E^T B^T = 0,$$

which, in turn, implies (26).

Then, we will calculate P_{K+E} for the perturbed control gain $K + E$. By (20), we have

$$(A - BK - BE)^T P_{K+E} + P_{K+E}(A - BK - BE) + Q + (K + E)^T R (K + E) = 0. \quad (28)$$

Subtracting (20) from (28) yields

$$(A - BK - BE)^T (P_{K+E} - P_K) + (P_{K+E} - P_K)(A - BK - BE) + E^T (RK - B^T P_K) + (RK - B^T P_K)^T E + E^T R E = 0,$$

which is equivalent to

$$P_{K+E} - P_K = \int_0^\infty e^{(A-BK-BE)t} [E^T (RK - B^T P_K) + (RK - B^T P_K)^T E + E^T R E] e^{(A-BK-BE)^T t} dt. \quad (29)$$

Taking the trace of (29), considering $J_2(K+E) - J_2(K) = \text{Tr}(P_{K+E} - P_K)$, and using the cyclic property of the trace in Lemma 2.1, we can obtain

$$J_2(K + E) - J_2(K) = 2\text{Tr}(E^T (RK - B^T P_K)Y_{K+E}) + \text{Tr}(E^T R E Y_{K+E}). \quad (30)$$

Since the first-order Taylor series approximation of Y_{K+E} is $Y_K + \Delta Y_K$, plugging it into (30), we can obtain (25). \square

In (25), the first-order term is $2\text{Tr}(E^T (RK - B^T P_K)Y_K) = \langle E, 2(RK - B^T P_K)Y_K \rangle$, which should be equal to $\langle E, \nabla J_2(K) \rangle$ by Taylor expansion. In other words, $\langle E, 2(RK - B^T P_K)Y_K \rangle = \langle E, \nabla J_2(K) \rangle$ for any small perturbation $E \in \mathbb{R}^{m \times n}$. Hence, the gradient of the objective function $J_2(K)$ is

$$\nabla J_2(K) = 2(RK - B^T P_K)Y_K. \quad (31)$$

We claim that $J_2(K)$ is a proper objective function (Definition 4.1). To this end, we will first prove several intermediate lemmas providing bounds on Y_K and P_K . The following lemma gives a lower bound of $\lambda_{\min}(Y_K)$.

Lemma 5.2. For any $K \in \mathcal{G}$, we have

$$\lambda_{\min}(Y_K) \geq \frac{1}{2\|A - BK^*\|_F + 2\|B\|\|K - K^*\|_F},$$

where K^* is the optimal control gain in (18).

Proof. Let $q \in \mathbb{R}^n$ denote a unit eigenvector of Y_K associated with the eigenvalue $\lambda_{\min}(Y_K)$. Pre and post multiplying (23) by q^T and q , respectively, we have

$$\lambda_{\min}(Y_K) q^T [(A - BK) + (A - BK)^T] q = -1. \quad (32)$$

Since $Y_K > 0$, it follows from (32) that $q^T [(A - BK) + (A - BK)^T] q < 0$. Consequently, we have

$$\begin{aligned} \lambda_{\min}((A - BK) + (A - BK)^T) &= \min_{\|v\|=1} v^T [(A - BK) + (A - BK)^T] v \\ &\leq q^T [(A - BK) + (A - BK)^T] q < 0. \end{aligned} \quad (33)$$

where the first equality is obtained by the Rayleigh principle [44, Theorem 4.2.2]. By (32) and (33), it holds

$$\begin{aligned} \lambda_{\min}(Y_K) &= \frac{1}{-q^T [(A - BK) + (A - BK)^T] q} \\ &\geq \frac{1}{-\lambda_{\min}((A - BK) + (A - BK)^T)} \\ &\geq \frac{1}{2\|A - BK^*\|_F + 2\|B\|\|K - K^*\|_F}, \end{aligned}$$

where the last line uses the relation $\|B(K^* - K)\| \leq \|B\|\|(K^* - K)\| \leq \|B\|\|(K^* - K)\|_F$. Thus, the proof is completed. \square

The following lemma gives the bounds of $\text{Tr}(Y_K)$.

Lemma 5.3. Given $K \in \mathcal{G}$ and $Q > 0$, we have

$$\frac{\text{Tr}(P_K - P^*)}{\|R\|\|K - K^*\|_F^2} \leq \text{Tr}(Y_K) \leq \frac{\text{Tr}(P_K)}{\lambda_{\min}(Q)} \quad (34)$$

Proof. Considering $K^* = R^{-1}B^T P^*$, we can rewrite (19) as

$$(A - BK^*)^T P^* + P^*(A - BK^*) + Q + (K^*)^T R K^* = 0. \quad (35)$$

It follows from (18) and (35) that

$$\begin{aligned} (A - BK)^T P^* + P^*(A - BK) + Q + (K^*)^T R K^* \\ + (K - K^*)^T R K^* + (K^*)^T R (K - K^*) = 0. \end{aligned} \quad (36)$$

Subtracting (36) from (20) and completing the squares yield

$$(A - BK)^T (P_K - P^*) + (P_K - P^*)(A - BK) + (K - K^*)^T R (K - K^*) = 0.$$

Since $A - BK$ is Hurwitz, by Lemma 2.7, we have

$$P_K - P^* = \int_0^\infty e^{(A-BK)t} (K - K^*)^T R (K - K^*) e^{(A-BK)t} dt. \quad (37)$$

Taking the trace of (37) and using the cyclic property of the trace in Lemma 2.1 and Eq. (24), we obtain

$$\begin{aligned} \text{Tr}(P_K - P^*) &= \text{Tr} \left(\int_0^\infty e^{(A-BK)t} e^{(A-BK)^T t} dt (K - K^*)^T R (K - K^*) \right) \\ &= \text{Tr}(Y_K (K - K^*)^T R (K - K^*)). \end{aligned} \quad (38)$$

By the trace inequality in Lemma 2.2 and considering the following relation

$$\|(K - K^*)^T R (K - K^*)\| \leq \text{Tr}((K - K^*)^T R (K - K^*)) \leq \|R\|\|K - K^*\|_F^2,$$

we have

$$\text{Tr}(P_K - P^*) \leq \|R\|\|K - K^*\|_F^2 \text{Tr}(Y_K).$$

Hence, the lower bound of $\text{Tr}(Y_K)$ in (34) is obtained.

Since $A - BK$ is Hurwitz, it follows from Lemma 2.7 and (20) that

$$P_K = \int_0^\infty e^{(A-BK)t} (Q + K^T R K) e^{(A-BK)t} dt. \quad (39)$$

Taking the trace of (39), and again using the cyclic property of the trace in Lemma 2.1 and the trace inequality in Lemma 2.2, we have

$$\begin{aligned} \text{Tr}(P_K) &\geq \text{Tr} \left(\int_0^\infty e^{(A-BK)^T t} Q e^{(A-BK)t} dt \right) \\ &= \text{Tr} \left(Q \int_0^\infty e^{(A-BK)t} e^{(A-BK)^T t} dt \right) \geq \lambda_{\min}(Q) \text{Tr}(Y_K). \end{aligned}$$

Hence, the right inequality in (34) follows readily. \square

Lemma 5.4. For any $K \in \mathcal{G}$, $\text{Tr}(P_K - P^*) \geq \alpha_4(\|K - K^*\|_F)$, where α_4 is a \mathcal{K}_∞ -function defined as

$$\alpha_4(r) := \frac{\lambda_{\min}(R) r^2}{2\|A - BK^*\|_F + 2\|B\|r}, \quad \forall r \geq 0.$$

Proof. Taking the trace of (37), and using the cyclic property of the trace in Lemma 2.1, we have

$$\begin{aligned} \text{Tr}(P_K - P^*) &= \text{Tr}((K - K^*)^T R (K - K^*) Y_K) \\ &\geq \lambda_{\min}(Y_K) \lambda_{\min}(R) \|K - K^*\|_F^2. \end{aligned}$$

Considering the lower bound of $\lambda_{\min}(Y_K)$ in Lemma 5.2, we can obtain

$$\text{Tr}(P_K - P^*) \geq \frac{\lambda_{\min}(R) \|K - K^*\|_F^2}{2\|A - BK^*\|_F + 2\|B\|\|K - K^*\|_F} = \alpha_4(\|K - K^*\|_F). \quad \square$$

Lemma 5.5 (Lemma 3.3 in [24]). The objective function $\mathcal{J}_2(K)$ is coercive, i.e. for any sequence $\{K_k\}_{k=0}^\infty$, $K_k \rightarrow \partial\mathcal{G}$ or $\|K_k\|_F \rightarrow \infty$, it holds $\mathcal{J}_2(K_k) \rightarrow \infty$, as $k \rightarrow \infty$.

Lemma 5.6. For any $K \in \mathcal{G}$, let $K' := R^{-1}B^T P_K$ and $M_K := (K - K')^T R (K - K')$. Then,

$$\text{Tr}(M_K) \geq a\|K - K^*\|_F^2 + a' \text{Tr}(P_K - P^*), \quad (40)$$

where a and a' are constants defined as

$$a := \frac{\lambda_{\min}(R) \lambda_{\min}(Y^*)}{2\lambda_{\min}(Y^*) + 2\lambda_{\max}(Y^*)}, \quad a' := \frac{1}{\lambda_{\min}(Y^*) + \lambda_{\max}(Y^*)}. \quad (41)$$

Proof. We can rewrite (20) as

$$\begin{aligned} (A - BK^*)^T P_K + P_K (A - BK^*) + Q + K^T R K \\ + (K^* - K)^T B^T P_K + P_K B (K^* - K) = 0. \end{aligned}$$

Considering $K' = R^{-1}B^T P_K$ and completing the squares, we have

$$\begin{aligned} (A - BK^*)^T P_K + P_K (A - BK^*) + Q + (K^*)^T R K^* \\ + (K - K')^T R (K - K') - (K' - K^*)^T R (K' - K^*) = 0. \end{aligned} \quad (42)$$

Subtracting (35) from (42) yields

$$\begin{aligned} (A - BK^*)^T (P_K - P^*) + (P_K - P^*)(A - BK^*) \\ + (K - K')^T R (K - K') - (K' - K^*)^T R (K' - K^*) = 0. \end{aligned} \quad (43)$$

Since $A - BK^*$ is Hurwitz, according to Lemma 2.7, we have

$$P_K - P^* = \int_0^\infty e^{(A-BK^*)t} [M_K - (K' - K^*)^T R (K' - K^*)] e^{(A-BK^*)t} dt. \quad (44)$$

Taking the trace of (44), and using the cyclic property of the trace in Lemma 2.1 yield

$$\text{Tr}(P_K - P^*) = \text{Tr}(Y^* M_K) - \text{Tr}(Y^* (K' - K^*)^T R (K' - K^*)).$$

By the trace inequality in Lemma 2.2, we have

$$\begin{aligned} \lambda_{\max}(Y^*) \text{Tr}(M_K) &\geq \lambda_{\min}(Y^*) \langle (K' - K^*)^T R (K' - K^*) \rangle \\ &\quad + \text{Tr}(P_K - P^*). \end{aligned} \quad (45)$$

Using Lemma 2.3 and Young's inequality, we can obtain

$$\begin{aligned}
& \langle K - K^*, R(K - K^*) \rangle \\
&= \langle (K - K') + (K' - K^*), R[(K - K') + (K' - K^*)] \rangle \\
&= \langle K - K', R(K - K') \rangle + \langle K' - K^*, R(K' - K^*) \rangle \\
&+ 2\langle K - K', R(K' - K^*) \rangle \\
&\leq 2\langle K - K', R(K - K') \rangle + 2\langle K' - K^*, R(K' - K^*) \rangle.
\end{aligned} \tag{46}$$

Noticing that $\text{Tr}(M_K) = \langle K - K', R(K - K') \rangle$. Plugging (45) into (46) and using the trace inequality in Lemma 2.2 yield

$$\begin{aligned}
\lambda_{\min}(R) \|K - K^*\|_F^2 &\leq \left(2 + 2\frac{\lambda_{\max}(Y^*)}{\lambda_{\min}(Y^*)}\right) \text{Tr}(M_K) \\
&- \frac{2}{\lambda_{\min}(Y^*)} \text{Tr}(P_K - P^*).
\end{aligned} \tag{47}$$

Hence, (40) follows from (47). \square

5.2. Perturbed standard gradient flow

In this subsection, we apply small-disturbance ISS to analyze the robustness property of the perturbed standard gradient flow for the LQR problem. At any position $K \in \mathcal{G}$, the steepest-descent direction of $J_2(K)$ is the solution to the problem [33]

$$\min_E \langle E, \nabla J_2(K) \rangle, \quad \text{subject to } \langle E, E \rangle = 1.$$

Therefore, by the method of Lagrange multipliers, the steepest-descent direction of $J_2(K)$ is

$$-\nabla J_2(K) / \|\nabla J_2(K)\|_F.$$

Consequently, along the steepest-descent direction and considering the expression of $\nabla J_2(K)$ in (31), the standard gradient flow is [4,24]

$$\frac{dK(s)}{ds} = -\eta \nabla J_2(K(s)) = -2\eta(RK(s) - B^T P(s))Y(s),$$

where $\eta > 0$ is a constant, $P(s) := P_{K(s)}$, and $Y(s) := Y_{K(s)}$. It is seen from (31) that the calculation of the gradient $\nabla J_2(K(s))$ relies on the system matrices, which are unknown in the setting of model-free RL. Many data-driven methods, e.g. approximate dynamic programming [30] and finite-difference algorithms [31], are proposed to numerically approximate the gradient. Therefore, in practice, $\widehat{\nabla J_2}(K(s))$, instead of $\nabla J_2(K(s))$, is utilized to optimize the control gain, and the perturbation $W(s) = \eta[\nabla J_2(K(s)) - \widehat{\nabla J_2}(K(s))]$ is unavoidable for the gradient flow. Hence, with the perturbation $W \in \mathcal{L}_{\infty}^{m \times n}$, the perturbed gradient flow of the LQR problem is

$$\frac{dK(s)}{ds} = -2\eta(RK(s) - B^T P(s))Y(s) + W(s). \tag{48}$$

The following lemma shows that the gradient of $J_2(K)$ is lower bounded by a \mathcal{K} -function of the deviation from the optimal value (gradient dominance condition), which is an important property of a proper objective function in Definition 4.1.

Lemma 5.7. *There exists a \mathcal{K} -function ξ_1 , such that for any $K \in \mathcal{G}$,*

$$\|\nabla J_2(K)\|_F \geq \xi_1(J_2(K) - J_2(K^*)).$$

Proof. Considering the expression of $\nabla J_2(K)$ in (31), the expression of K' in Lemma 5.6, and the cyclic property of the trace in Lemma 2.1, we have

$$\text{Tr}(\nabla J_2^T(K) \nabla J_2(K)) = 4\text{Tr}(Y_K^2(K - K')^T R^2(K - K')).$$

By the trace inequality in Lemma 2.2, it holds

$$\begin{aligned}
\text{Tr}(\nabla J_2^T(K) \nabla J_2(K)) &\geq 4\lambda_{\min}(Y_K)^2 \text{Tr}\left((K - K')^T \sqrt{RR} \sqrt{R}(K - K')\right) \\
&= 4\lambda_{\min}(Y_K)^2 \text{Tr}\left(\sqrt{R}(K - K')(K - K')^T \sqrt{RR}\right) \\
&\geq 4\lambda_{\min}(Y_K)^2 \lambda_{\min}(R) \text{Tr}\left(\sqrt{R}(K - K')(K - K')^T \sqrt{R}\right) \\
&= 4\lambda_{\min}(Y_K)^2 \lambda_{\min}(R) \text{Tr}(M_K).
\end{aligned}$$

By Lemmas 5.2 and 5.6, it follows that

$$\text{Tr}(\nabla J_2^T(K) \nabla J_2(K)) \geq \lambda_{\min}(R) \frac{a\|K - K^*\|_F^2 + a'\text{Tr}(P_K - P^*)}{(\|A - BK^*\|_F + \|B\|\|K - K^*\|_F)^2}. \tag{49}$$

To simplify notations, let

$$a_1 := \lambda_{\min}(R)a, \quad a_2 := \lambda_{\min}(R)a', \quad a_3 := \|A - BK^*\|_F, \quad a_4 := \|B\|, \tag{50}$$

and

$$\sigma(r, p) := \frac{a_1 r^2 + a_2 p}{(a_3 + a_4 r)^2}.$$

It is clear that $\sigma(r, p) \geq \sigma(r, 0)$ for $p \geq 0$. Taking the derivative of $\sigma(r, p)$ with respect to r , we have

$$\frac{\partial \sigma(r, p)}{\partial r} = \frac{2a_1 a_3 r - 2a_2 a_4 p}{(a_3 + a_4 r)^3}.$$

When $p \geq \frac{a_1 a_3}{a_2 a_4} r$, $\sigma(r, p)$ is strictly decreasing in r . When $p \leq \frac{a_1 a_3}{a_2 a_4} r$, $\sigma(r, p)$ is strictly increasing in r . Hence, for each p , $r = \frac{a_2 a_4}{a_1 a_3} p$ is the minimum point of $\sigma(r, p)$, and

$$\sigma(r, p) \geq \sigma\left(\frac{a_2 a_4}{a_1 a_3} p, p\right) \geq \sigma\left(\frac{a_2 a_4}{a_1 a_3} p, 0\right) =: \xi_1^2(p), \tag{51}$$

where

$$\xi_1(p) := \frac{a_5 p}{a_3 + a_6 p},$$

and

$$a_5 := \frac{a_2 a_4}{\sqrt{a_1 a_3}}, \quad a_6 := \frac{a_2 a_4^2}{a_1 a_3}. \tag{52}$$

Since $\frac{d\xi_1(p)}{dp} = \frac{a_3 a_5}{(a_3 + a_6 p)^2} > 0$, ξ_1 is a \mathcal{K} -function with the range $[0, \frac{\sqrt{a_1}}{a_4}]$.

Plugging (51) into (49) yields

$$\|\nabla J_2(K)\|_F = \text{Tr}(\nabla J_2^T(K) \nabla J_2(K))^{\frac{1}{2}} \geq \xi_1\left(\text{Tr}(P_K - P^*)\right). \quad \square$$

Remark 5.1. As remarked in the Introduction, the gradient dominance condition in Lemma 5.7, the CJS-PL condition, can be considered as a generalization of the well-known Polyak-Lojasiewicz (PL) condition which only holds on a compact set of stabilizing control gains [4, 24]. The gradient dominance condition in Lemma 5.7 removes the restriction to compact sets.

We next revisit the one-dimensional system mentioned in the Introduction, in order to illustrate Lemma 5.7. Suppose that $m = n = 1$ and $A = B = Q = R = 1$. In this case, the admissible set is $\mathcal{G} = \{K | K > 1\}$, and one obtains

$$P^* = 1 + \sqrt{2}, \quad K^* = 1 + \sqrt{2}, \quad Y_K = \frac{1}{2(K-1)}, \quad Y^* = \frac{\sqrt{2}}{4}, \tag{53}$$

$$J_2(K) = \frac{1 + K^2}{2(K-1)}, \quad J_2(K) - J_2(K^*) = \frac{(K - K^*)^2}{2(K-1)}.$$

The constants in (41), (50) and (52) can be computed as

$$\begin{aligned}
a &= \frac{1}{4}, \quad a' = \sqrt{2}, \quad a_1 = \frac{1}{4}, \quad a_2 = \sqrt{2}, \\
a_3 &= \sqrt{2}, \quad a_4 = 1, \quad a_5 = 2, \quad a_6 = 4.
\end{aligned}$$

Consequently, we have

$$\xi_1(p) = \frac{2p}{\sqrt{2} + 4p}, \quad \xi_1(J_2(K) - J_2(K^*)) = \frac{(K - K^*)^2}{\sqrt{2}(K-1) + 2(K - K^*)^2}.$$

The gradient of $J_2(K)$ is

$$\nabla J_2(K) = \frac{K^2 - 2K - 1}{2(K-1)^2} = \frac{(K - K^*)^2 + 2\sqrt{2}(K - K^*)}{2(K-1)^2}.$$

When $K \geq K^*$, $\nabla J_2(K) \geq 0$, and

$$\begin{aligned} & \frac{\nabla J_2(K)}{\xi_1 (J_2(K) - J_2(K^*))} \\ &= \frac{2(K - K^*)^3 + 5\sqrt{2}(K - K^*)^2 + 6(K - K^*) + 4\sqrt{2}}{2(K - K^*)^3 + 4\sqrt{2}(K - K^*)^2 + 4(K - K^*)} \geq 1. \end{aligned}$$

When $K^* \geq K > 1$, $\nabla J_2(K) \leq 0$, and

$$\begin{aligned} & -\nabla J_2(K) - \xi_1 (J_2(K) - J_2(K^*)) \\ &= \frac{-4(K - 1)^4 + 7\sqrt{2}(K - 1)^3 - 4(K - 1)^2 - 6\sqrt{2}(K - 1) + 8}{2\sqrt{2}(K - 1)^3 + 4(K - 1)^2(K - K^*)^2} \geq 0. \end{aligned}$$

Therefore, $\|\nabla J_2(K)\| \geq \xi_1 (J_2(K) - J_2(K^*))$, which is consistent with Lemma 5.7.

Based on Lemma 5.7, we are ready to state the main result on the small-disturbance ISS property of the perturbed standard gradient flow in (48).

Theorem 5.1. *System (48) is small-disturbance ISS with respect to W .*

Proof. Define

$$\mathcal{V}_3(K) := J_2(K) - J_2(K^*) \quad (54)$$

It follows from (22) that $\mathcal{V}_3(K) = \text{Tr}(P_K - P)^*$. Clearly, $\mathcal{V}_3(K)$ is smooth in K [24, Proposition 3.2]. Since K^* is the unique minimum of $J_2(K)$, $\mathcal{V}_3(K)$ is a positive definite function with respect to K^* . The coercivity of $\mathcal{V}_3(K)$ can be obtained by Lemma 5.5. Therefore, by Definition 3.1, $\mathcal{V}_3(K)$ is a size function for (\mathcal{G}, K^*) . Hence, by Lemma 5.7 and Definition 4.1, J_2 is a proper objective function. According to Theorem 4.1, the proof of Theorem 5.1 is completed. \square

As a direct consequence of Theorem 5.1, an estimate of $\|K(s) - K^*\|_F$ is provided in the following corollary.

Corollary 5.1. *There exist a constant $d_1 > 0$, a \mathcal{KL} -function $\beta_2(\cdot, \cdot)$, and a $\mathcal{K}_{[0, d_1]}$ -function γ_4 , such that for all perturbations W essentially bounded by d_1 (i.e. $\|W\|_\infty < d_1$), and all initial conditions $K(0) \in \mathcal{G}$, $K(s)$ satisfies*

$$\|K(s) - K^*\|_F \leq \beta_2(\mathcal{V}_3(K(0)), s) + \gamma_4(\|W\|_\infty), \quad \forall s \geq 0, \quad (55)$$

where \mathcal{V}_3 is defined in (54).

Proof. Since system (48) is small-disturbance ISS, when $\|W\|_\infty < d_1$, it holds

$$\mathcal{V}_3(K(s)) \leq \beta_2'(\mathcal{V}_3(K(0)), s) + \gamma_4'(\|W\|_\infty), \quad \forall s \geq 0,$$

where β_2' is a \mathcal{KL} -function and γ_4' is a $\mathcal{K}_{[0, d_1]}$ -function. According to Lemma 5.4 and recalling that $\mathcal{V}_3(K(s)) = \text{Tr}(P(s) - P)^*$, we have

$$\|K(s) - K^*\|_F \leq \alpha_4^{-1}(\beta_2'(\mathcal{V}_3(K(0)), s) + \gamma_4'(\|W\|_\infty)).$$

Using Lemma 2.6, we can obtain (55). \square

5.3. Perturbed natural gradient flow

In this subsection, we will show that the perturbed natural gradient flow is small-disturbance ISS. It follows from (38) that

$$J_2(K) - J_2(K^*) = \langle K - K^*, R(K - K^*) \rangle_{Y_K}.$$

Hence, the objective function can be viewed as a quadratic function over the Riemannian manifold $(\mathcal{G}, \langle \cdot, \cdot \rangle_{Y_K})$. As seen in the expression of $\nabla J_2(K)$, the magnitude of the gradient depends on Y_K , and Y_K may tend to infinity when $K \rightarrow \partial\mathcal{G}$, and tend to zero when $\|K\|_F \rightarrow \infty$ (see the illustrative one-dimensional system in (53)). The non-isotropic property of the magnitude of the gradient may slow down the convergence of the gradient flow. To handle the non-isotropic property, in [45,46], the natural gradient was proposed as a way to modify

the standard gradient search direction according to the Riemannian structure of the parameter space.

Over the Riemannian manifold $(\mathcal{G}, \langle \cdot, \cdot \rangle_{Y_K})$, the steepest-descent direction can be obtained by solving

$$\min_E \langle E, \nabla J_2(K) \rangle, \quad \text{subject to } \langle E, E \rangle_{Y_K} = 1. \quad (56)$$

By the method of Lagrange multipliers, the solution of (56) is

$$E = -\nabla J_2(K) Y_K^{-1} / \langle \nabla J_2(K) Y_K^{-1}, \nabla J_2(K) Y_K^{-1} \rangle_{Y_K}^{1/2}.$$

The natural gradient of $J_2(K)$ over the Riemannian manifold $(\mathcal{G}, \langle \cdot, \cdot \rangle_{Y_K})$, denoted by $\text{grad}(J_2(K))$, is

$$\text{grad}(J_2(K)) := \nabla J_2(K) Y_K^{-1} = 2(RK - B^T P_K).$$

Considering the perturbation, the natural gradient flow is

$$\frac{dK(s)}{ds} = -2\eta(RK(s) - B^T P(s)) + W(s). \quad (57)$$

The following lemma is introduced to pave the foundation for the proof of small-disturbance ISS property of system (57).

Lemma 5.8. *For any $K \in \mathcal{G}$, we have*

$$2\langle K - K^*, R(K - K') \rangle_{Y^*} = \text{Tr}(P_K - P^*) + \langle K - K^*, R(K - K^*) \rangle_{Y^*}. \quad (58)$$

Proof. By completing the squares, we have

$$\begin{aligned} & (K - K')^T R(K - K') - (K' - K^*)^T R(K' - K^*) \\ &= (K - K')^T R(K - K^*) + (K - K^*)^T R(K - K') - (K - K^*)^T R(K - K^*). \end{aligned}$$

Therefore, we can rewrite (43) as

$$\begin{aligned} & (A - BK^*)^T (P_K - P^*) + (P_K - P^*) (A - BK^*) + (K - K')^T R(K - K^*) \\ &+ (K - K^*)^T R(K - K') - (K - K^*)^T R(K - K^*) = 0. \end{aligned}$$

Since $A - BK^*$ is Hurwitz, by Lemma 2.7 we have

$$\begin{aligned} P_K - P^* &= \int_0^\infty e^{(A - BK^*)^T t} [(K - K')^T R(K - K^*) \\ &+ (K - K^*)^T R(K - K') - (K - K^*)^T R(K - K^*)] e^{(A - BK^*) t} dt. \end{aligned} \quad (59)$$

Taking the trace of (59) and using the cyclic property of trace in Lemma 2.1, we have (58). \square

Next, we will prove the small-disturbance ISS property of the perturbed natural gradient flow.

Theorem 5.2. *Given $Q > 0$, system (57) is small-disturbance ISS with respect to W .*

Proof. Let $\mathcal{V}_4(K) := \frac{1}{2} \langle K - K^*, K - K^* \rangle_{Y^*}$. Clearly, by Lemma 2.2, $\mathcal{V}_4(K)$ is bounded by

$$\frac{1}{2} \lambda_{\min}(Y^*) \|K - K^*\|_F^2 \leq \mathcal{V}_4(K) \leq \frac{1}{2} \|Y\| \|K - K^*\|_F^2.$$

Then, it holds:

$$\begin{aligned} & \langle \nabla \mathcal{V}_4(K), -2\eta(RK - B^T P_K) + W \rangle \\ &= -2\eta \langle K - K^*, R(K - K') \rangle_{Y^*} + \langle K - K^*, W \rangle_{Y^*}. \end{aligned}$$

Recall that $K' = R^{-1} B^T P_K$. According to Lemma 5.8, we obtain

$$\begin{aligned} & \langle \nabla \mathcal{V}_4(K), -2\eta(RK - B^T P_K) + W \rangle \\ &= -\eta \text{Tr}(P_K - P^*) - \eta \langle K - K^*, R(K - K^*) \rangle_{Y^*} + \langle K - K^*, W \rangle_{Y^*} \\ &\leq -\eta \text{Tr}(P_K - P^*) - \eta \langle K - K^*, R(K - K^*) \rangle_{Y^*} \\ &+ \langle K - K^*, K - K^* \rangle_{Y^*}^{1/2} \langle W, W \rangle_{Y^*}^{1/2} \\ &\leq -\eta \text{Tr}(P_K - P^*) - \frac{\eta \lambda_{\min}(R)}{2} \langle K - K^*, K - K^* \rangle_{Y^*} + \frac{\|Y^*\|}{2\eta \lambda_{\min}(R)} \|W\|_F^2, \end{aligned} \quad (60)$$

where the first inequality follows from Lemma 2.3, and the last inequality is according to Young's inequality and the trace inequality in Lemma 2.2.

$$\begin{aligned} \text{Differentiating } \mathcal{V}_3(K) = \mathcal{J}_2(K) - \mathcal{J}_2(K^*) \text{ and considering (31) yield} \\ \langle \nabla \mathcal{V}_3(K), -2\eta(RK - B^T P_K) + W \rangle &= -4\eta \langle RK - B^T P_K, RK - B^T P_K \rangle_{Y_K} \\ &+ 2 \langle RK - B^T P_K, W \rangle_{Y_K} \\ &\leq -3\eta \langle RK - B^T P_K, RK - B^T P_K \rangle_{Y_K} + \frac{1}{\eta} \langle W, W \rangle_{Y_K}, \end{aligned}$$

where the last line is obtained by Lemma 2.3 and Young's inequality. Using Lemma 5.3 and the trace inequality in Lemma 2.2, we have

$$\begin{aligned} \langle \nabla \mathcal{V}_3(K), -2\eta(RK - B^T P_K) + W \rangle &\leq -3\eta \langle RK - B^T P_K, RK - B^T P_K \rangle_{Y_K} \\ &+ \frac{\mathcal{V}_3(K) + \text{Tr}(P)^*}{\eta \lambda_{\min}(Q)} \|W\|_F^2. \end{aligned} \quad (61)$$

Let $\mathcal{V}_5(K) := \mathcal{V}_3(K) + \mathcal{V}_4(K)$. Since \mathcal{V}_3 is a size function and \mathcal{V}_4 is positive definite with respect to K^* , \mathcal{V}_5 is also a size function. It follows from (60) and (61) that

$$\begin{aligned} \langle \nabla \mathcal{V}_5(K), -2\eta(RK - B^T P_K) + W \rangle &\leq -\eta \mathcal{V}_5(K) - \eta \lambda_{\min}(R) \mathcal{V}_4(K) \\ &+ \frac{\mathcal{V}_3(K) + b_1}{\eta b_2} \|W\|_F^2, \end{aligned} \quad (62)$$

where

$$b_1 := \frac{\|Y^*\| \lambda_{\min}(Q)}{2\lambda_{\min}(R)} + \text{Tr}(P^*), \quad b_2 := \lambda_{\min}(Q), \quad b_3 := \eta^2 \lambda_{\min}(R) \lambda_{\min}(Q).$$

Without losing generality, assume that $\lambda_{\min}(R) \leq 1$. Then, it follows from (62) that

$$\langle \nabla \mathcal{V}_5(K), -2\eta(RK - B^T P_K) + W \rangle \leq -\eta \lambda_{\min}(R) \mathcal{V}_5(K) + \frac{\mathcal{V}_5(K) + b_1}{\eta b_2} \|W\|_F^2.$$

Thus, if

$$\|W\|_F \leq \left(\frac{b_3 \mathcal{V}_5(K)}{2\mathcal{V}_5(K) + 2b_1} \right)^{\frac{1}{2}} =: \xi_2(\mathcal{V}_5(K)),$$

it is guaranteed that

$$\langle \nabla \mathcal{V}_5(K), -2\eta(RK - B^T P_K) + W \rangle \leq -\frac{\eta \lambda_{\min}(R)}{2} \mathcal{V}_5(K).$$

Since $\frac{d\xi_2(r)}{dr} = \frac{1}{2} \xi^{-1}(r) \frac{2b_1 b_3}{(2r+2b_1)^2} > 0$, $\forall r > 0$, ξ_2 is a \mathcal{K} -function, and its range is $[0, \sqrt{\frac{b_3}{2}})$. Consequently, \mathcal{V}_5 is a small-disturbance ISS-Lyapunov function. According to Theorem 3.1, we conclude that system (57) is small-disturbance ISS. \square

5.4. Perturbed Newton gradient flow

By applying the results of Section 3, we will show that the perturbed Newton gradient flow is small-disturbance ISS. The Newton gradient descent method was first adopted in [47] for solving the LQR problem, and it converges to the optimum at a quadratic convergence rate. The Newton direction is derived from the second-order Taylor series approximation of $\mathcal{J}_2(K + E)$, which, according to Lemma 5.1, can be expressed as

$$\mathcal{J}_2(K + E) = \mathcal{J}_2(K) + \langle E, 2(RK - B^T P_K) \rangle_{Y_K} + \langle E, RE \rangle_{Y_K} + O(\|E\|_F^2).$$

By minimizing the second-order Taylor series approximation of $\mathcal{J}_2(K)$ over E , the Newton direction is obtained as $-(K - R^{-1} B^T P_K)$. Considering the perturbation, the Newton gradient flow is

$$\frac{dK(s)}{ds} = -\eta(K(s) - R^{-1} B^T P(s)) + W(s). \quad (63)$$

Theorem 5.3. *Given $Q > 0$, system (63) is small-disturbance ISS with respect to W .*

Proof. The proof follows from the proof of Theorem 5.2 by defining $\mathcal{V}_6(K) = \mathcal{V}_3(K) + \frac{1}{2} \langle K - K^*, R(K - K^*) \rangle_{Y^*}$. \square

6. Conclusions

In this paper, we studied the small-disturbance ISS property of continuous-time gradient flows on an open subset of certain Euclidean space. In the framework of small-disturbance ISS, the transient behavior, the convergence speed, and the robustness to the perturbations of gradient flows can be well quantified. As a by-product, a Lyapunov characterization of small-disturbance ISS is given. Upon specification to the policy optimization of the LQR problem, three kinds of perturbed gradient flows, including standard gradient flow, natural gradient flow, and Newton gradient flow, were studied in greater details. In particular, they are all small-disturbance ISS.

CRedit authorship contribution statement

Leilei Cui: Writing – review & editing, Writing – original draft, Formal analysis, Conceptualization. **Zhong-Ping Jiang:** Writing – review & editing, Project administration, Funding acquisition, Formal analysis, Conceptualization. **Eduardo D. Sontag:** Writing – review & editing, Project administration, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Acknowledgments

This work has been supported in part by the National Science Foundation grants CNS-2227153 and ECCS-2210320 (L.C. and Z.P.J.), AFOSR grant FA9550-21-1-0289 and ONR, United States grant N00014-21-1-2431 (E.D.S.).

References

- [1] A. Cherukuri, E. Mallada, S. Low, J. Cortés, The role of convexity in saddle-point dynamics: Lyapunov function and robustness, *IEEE Trans. Automat. Control* 63 (8) (2018) 2449–2464.
- [2] R. Suttner, S. Dashkovskiy, Robustness and averaging properties of a large-amplitude, high-frequency extremum seeking control scheme, *Automatica* 136 (2022) 110020.
- [3] G. Bianchin, J.I. Poveda, E. Dall'Anese, Online optimization of switched LTI systems using continuous-time and hybrid accelerated gradient flows, *Automatica* 146 (2022) 110579.
- [4] H. Mohammadi, A. Zare, M. Soltanolkotabi, M.R. Jovanović, Convergence and sample complexity of gradient methods for the model-free linear–quadratic regulator problem, *IEEE Trans. Automat. Control* 67 (5) (2022) 2435–2450.
- [5] E.D. Sontag, Remarks on input to state stability of perturbed gradient flows, motivated by model-free feedback control learning, *Syst. Control Lett.* 161 (2022) 105138.
- [6] K. Ahn, P. Jain, Z. Ji, S. Kale, P. Netrapalli, G.I. Shamir, Reproducibility in optimization: Theoretical framework and limits, in: S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, A. Oh (Eds.), in: *Advances in Neural Information Processing Systems*, vol. 35, Curran Associates, Inc., 2022, pp. 18022–18033.
- [7] A.C.B. de Olivera, M. Siami, E.D. Sontag, Dynamics and perturbations of overparameterized linear neural networks, in: 2023 62nd IEEE Conference on Decision and Control, (CDC), 2023, pp. 7356–7361.
- [8] E.D. Sontag, Smooth stabilization implies coprime factorization, *IEEE Trans. Automat. Control* 34 (4) (1989) 435–443.
- [9] E. Sontag, Input to State Stability: Basic Concepts and Results, in: *Lecture Notes in Mathematics*, Springer Verlag, Germany, 2008, pp. 163–220.
- [10] B.T. Polyak, Gradient methods for the minimisation of functionals, *USSR Comput. Math. Math. Phys.* 3 (4) (1963) 864–878.
- [11] S. Łojasiewicz, A topological property of real analytic subsets, in: *Colloques internationaux du C.N.R.S. 117*, in: *Les Équations aux Dérivées Partielles*, 1963, pp. 87–89, in French.

- [12] H. Karimi, J. Nutini, M. Schmidt, Linear convergence of gradient and proximal-gradient methods under the Polyak-Lojasiewicz condition, in: P. Frasconi, N. Landwehr, G. Manco, J. Vreeken (Eds.), *Machine Learning and Knowledge Discovery in Databases*, Springer International Publishing, Cham, 2016, pp. 795–811.
- [13] E.D. Sontag, Comments on integral variants of ISS, *Systems Control Lett.* 34 (1–2) (1998) 93–100.
- [14] D. Angeli, E.D. Sontag, Y. Wang, A characterization of integral input-to-state stability, *IEEE Trans. Autom. Control* 45 (6) (2000) 1082–1097.
- [15] B. Pang, Z.P. Jiang, Robust reinforcement learning: A case study in linear quadratic regulation, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, (10) 2021, pp. 9303–9311.
- [16] B. Pang, T. Bian, Z.P. Jiang, Robust policy iteration for continuous-time linear quadratic regulation, *IEEE Trans. Automat. Control* 67 (1) (2022) 504–511.
- [17] J.I. Poveda, M. Krstić, Nonsmooth extremum seeking control with user-prescribed fixed-time convergence, *IEEE Trans. Automat. Control* 66 (12) (2021) 6156–6163.
- [18] D. Angeli, E.D. Sontag, Y. Wang, Input-to-state stability with respect to inputs and their derivatives, *Internat. J. Robust Nonlinear Control* 13 (11) (2003) 1035–1056.
- [19] R.S. Sutton, A.G. Barto, *Reinforcement Learning: An Introduction*, 2nd Edition, MIT Press, Cambridge, MA, 2018.
- [20] V. Konda, J. Tsitsiklis, Actor-critic algorithms, in: *Advances in Neural Information Processing Systems*, vol. 12, 1999, pp. 1008–1014.
- [21] T.P. Lillicrap, J.J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, D. Wierstra, *Continuous control with deep reinforcement learning*, 2015, arXiv preprint arXiv:1509.02971.
- [22] J. Schulman, S. Levine, P. Abbeel, M. Jordan, P. Moritz, Trust region policy optimization, in: F. Bach, D. Blei (Eds.), in: *Proceedings of the 32nd International Conference on Machine Learning*, vol. 37, Lille, France, 2015, pp. 1889–1897.
- [23] W.S. Levine, M. Athans, On the determination of the optimal constant output feedback gains for linear multivariable systems, *IEEE Trans. Automat. Control* 15 (1) (1970) 44–48.
- [24] J. Bu, A. Mesbahi, M. Mesbahi, Policy gradient-based algorithms for continuous-time linear quadratic control, 2020, arXiv preprint arXiv:2006.09178.
- [25] B. Hu, K. Zhang, N. Li, M. Mesbahi, M. Fazel, T. Başar, Toward a theoretical foundation of policy optimization for learning control policies, *Annu. Rev. Control, Robotics, Auton. Syst.* 6 (1) (2023) 123–158.
- [26] D.P. Bertsekas, J.N. Tsitsiklis, *Neuro-Dynamic Programming*, Athena Scientific, New Hampshire, 1996.
- [27] W.B. Powell, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, John Wiley & Sons, Hoboken, New Jersey, 2007.
- [28] Y. Jiang, Z.P. Jiang, *Robust Adaptive Dynamic Programming*, Wiley-IEEE Press, Hoboken, New Jersey, 2017.
- [29] F.L. Lewis, D. Liu, *Reinforcement Learning and Approximate Dynamic Programming for Feedback Control*, Wiley-IEEE Press, New Jersey, 2013.
- [30] Z.P. Jiang, T. Bian, W. Gao, Learning-based control: A tutorial and some recent results, *Found. Trends® Syst. Control* 8 (3) (2020) 176–284.
- [31] M. Fazel, R. Ge, S. Kakade, M. Mesbahi, Global convergence of policy gradient methods for the linear quadratic regulator, in: *International Conference on Machine Learning*, 2018, pp. 1467–1476.
- [32] Y. Li, Y. Tang, R. Zhang, N. Li, Distributed reinforcement learning for decentralized linear quadratic control: A derivative-free policy optimization approach, *IEEE Trans. Automat. Control* 67 (12) (2022) 6429–6444.
- [33] J. Nocedal, S.J. Wright, *Numerical Optimization*, Springer, New York, 2006.
- [34] W. Hahn, *Stability of Motion*, Springer, Berlin, Heidelberg, 1967.
- [35] K.B. Petersen, M.S. Pedersen, *The Matrix Cookbook*, version 20081110, 2008.
- [36] S.-D. Wang, T.-S. Kuo, C.-F. Hsu, Trace bounds on the solution of the algebraic matrix Riccati and Lyapunov equation, *IEEE Trans. Automat. Control* 31 (7) (1986) 654–656.
- [37] Z.P. Jiang, A.R. Teel, L. Praly, Small-gain theorem for ISS systems and applications, *Math. Control Signals Systems* 7 (1994) 95–120.
- [38] E.D. Sontag, *Mathematical Control Theory: Deterministic Finite Dimensional Systems*, vol. 6, Springer Science & Business Media, New York, 1998.
- [39] E.D. Sontag, Y. Wang, On characterizations of the input-to-state stability property, *Systems Control Lett.* 24 (5) (1995) 351–359.
- [40] Y. Lin, E.D. Sontag, Y. Wang, A smooth converse Lyapunov theorem for robust stability, *SIAM J. Control Optim.* 34 (1) (1996) 124–160.
- [41] F.W. Wilson, The structure of the level surfaces of a Lyapunov function, *J. Differential Equations* 3 (3) (1967) 323–329.
- [42] J. Milnor, Differential topology, in: T.L. Saaty (Ed.), 2nd, in: *Lectures in Modern Mathematics*, John Wiley & Sons, New York, US, 1964, pp. 165–183.
- [43] K. Zhou, J.C. Doyle, K. Glover, *Robust and Optimal Control*, Prentice Hall, New Jersey, 1996.
- [44] R.A. Horn, C.R. Johnson, *Matrix Analysis*, 2nd Edition, Cambridge University Press, New York, US, 2013.
- [45] S.-I. Amari, S.C. Douglas, Why natural gradient? in: *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP'98 (Cat. No. 98CH36181)*, vol. 2, 1998, pp. 1213–1216.
- [46] S.-I. Amari, Natural gradient works efficiently in learning, *Neural Comput.* 10 (2) (1998) 251–276.
- [47] D.L. Kleinman, On an iterative technique for Riccati equation computations, *IEEE Trans. Automat. Control* 13 (1) (1968) 114–115.