Lyapunov-Based Certified Region-of-Attraction Estimation for Constrained LQR-Controlled Systems

Ece Nur Demirhan Coskun*ᵁ Kemal Caglar Coskunᵁ Christoph Lüth*ᵁ Dieter Hutter*ᵁ
*Cyber-Physical Systems, DFKI GmbH, 28359 Bremen, Germany
ᵁInstitute of Computer Science, University of Bremen, 28359 Bremen, Germany
{christoph.lueth, dieter.hutter}@dfki.de {ecenur, kcoskun}@uni-bremen.de

***Abstract*—Many cyber-physical systems operate in safety-critical settings where formal stability guarantees are required. In this context, methods for *Region-of-Attraction* (ROA) are of particular interest, as they provide certified characterizations of the initial states from which system executions remain stable despite disturbances and constraints.**

We present a Lyapunov-based method for computing certified ROAs for constrained nonlinear systems controlled by *Linear Quadratic Regulators* (LQRs). The approach constructs a composite Lyapunov function that combines a problem-informed energy-like term with the LQR cost-to-go, and leverages a targeted sampling strategy to efficiently identify invariant sublevel sets while preserving formal guarantees. Rather than proposing a universal ROA solver, this work presents a practical and sound framework that can be instantiated in a system-specific manner. The approach is evaluated on a torque-limited inverted pendulum benchmark and produces certified ROAs that are less conservative than those obtained using standard quadratic Lyapunov functions and analytic baselines.

***Index Terms*—cyber-physical systems, region of attraction, Lyapunov methods, LQR, nonlinear control, safety verification**

I. Introduction

Cyber-Physical Systems (CPSs) operating under safety constraints require formal guarantees to ensure stable behavior under disturbances and actuator limits [1]. In such settings, estimating the *Region-of-Attraction* (ROA) provides a principled way to certify the set of initial conditions from which closed-loop trajectories remain stable and remain within the safe state space, despite disturbances or uncertainties [2]. This capability is particularly important in safety-critical applications such as robot control, where it is essential to ensure that all system trajectories remain stable and converge under expected operational conditions [3].

Traditional methods for estimating the ROA, such as optimization or sampling, can be computationally demanding and pose challenges for formal verification in safety-critical settings. In contrast, methods based on *Lyapunov functions* (LFs) and energy concepts are computationally tractable and provide reliable estimations [3]–[6]. However, these methods do not capture the entire ROA, but rather find a conservative underapproximation. In time-critical and safety-critical applications, this can be an acceptable compromise depending on the severity of the underapproximation. Therefore, it is important to construct LFs that yield less conservative certified ROA estimates.

In this paper, we present a Lyapunov-based ROA estimation methodology for computing certified ROAs for constrained nonlinear systems under *Linear Quadratic Regulator* (LQR) control, and illustrate the approach on an inverted pendulum benchmark. We propose a function that integrates various energy functions and controller-specific quadratic functions. This function is then combined with a targeted-sampling-based ROA calculation process to estimate the ROA under various parameters. Through experiments, we demonstrate the applicability and effectiveness of our method.

Section II summarizes a brief background on ROA estimation, including existing methods and open problems. Section III presents the necessary background information relevant to this work and introduces the notation used throughout the paper. Section IV introduces our Lyapunov-based method, which enables efficient computation of certified ROA estimates for LQR-controlled inverted pendulums. In Section V we present the results of our experiments. The paper concludes with Section VI.

II. Related Work

Recent advancements in CPSs have prioritized safety under dynamic uncertainties by estimating the Region of Attraction (ROA), crucial for operational stability.

Reachability-analysis-based methods are based on either forward or backward reachability. Forward-reachability methods identify all states accessible from a specified initial set, considering disturbance limits [7]. In ROA analysis, this approach is utilized by partitioning the state-space into multiple discrete regions and assessing whether these regions are part of the ROA. Conversely, backward-reachability methods compute ROA boundaries by applying a backward-reachability algorithm [8]. This algorithm identifies all initial states from which system trajectories converge to a target set. Therefore, when the target set is defined as a very small area around the equilibrium point, the algorithm returns the ROA.

Learning-based methods employ historical data and machine learning to predict system safety under actual conditions,

This research has been supported in part by the German Ministry for Research, Technology and Space (BMFTR) with projects VeryHuman (grant no. 01IW20004) and ExaVerse (grant no. 01IW25003).

deriving ROA from these predictions. Neural-network-based Lyapunov approaches have also been proposed for ROA estimation, offering increased expressiveness through learning-based constructions of LFs [9].

Robust and optimal control designs ensure system stability and safety under uncertainties and environmental changes, often through worst-case scenario analysis to guarantee stability across all conditions within the uncertainty bounds [10].

Lyapunov-based methods are foundational techniques that use classical approaches and *Sum of Squares* (SOS) programming to affirm system stability. The classical approaches involve constructing a LF, a scalar function that decreases along the trajectories of the system, signifying stability. Using SOS programming and semidefinite programming allows for the estimation of the ROA, defined as the largest set where the LF decreases [11].

Various forms of LFs have been employed to estimate the ROA for nonlinear systems. The approach in [5] expands ROA estimation by defining it with a variable shape factor. In another work [6], the LF space is used to apply Ellipsoidal ROA to systems with polynomial dynamics when restricted to quadratic functions.

In [12], a fast method for estimating the *Domain-of-Attraction* (DOA), where ROA $\subseteq$ DOA, is proposed. This method, based on random samples in the state-space, does not concentrate on the selection of LFs. In another work [3], the authors present an analytic approach for estimating the ROA for a torque-limited simple pendulum. Although this approach is faster compared to the quadratic-based LF method described in [12], it faces challenges when extended to higher-dimensional systems and provides no formal guarantees.

In summary, while non-Lyapunov-based methods often incur higher computational cost or lack formal guarantees, commonly used Lyapunov-based approaches relying on quadratic LFs may yield conservative ROA estimates. Other techniques, such as numerical optimization, analytic methods, and statistical sampling, typically involve trade-offs between computational efficiency, conservatism, and the availability of formal guarantees, highlighting the need for approaches that balance practicality and certification.

III. Preliminaries

In this section, we introduce LQR-controlled system dynamics, the ROA, and LF-based ROA estimation, using an inverted pendulum as a running example and case study. A torque-limited inverted pendulum is a simple yet representative example of an underactuated CPS.

A. LQR-Controlled System Dynamics

Consider a closed-loop dynamics analysis of an inverted pendulum controlled by an LQR to stabilize the pendulum in its upright position, as seen in Figure 1. It consists of a rigid arm of length *l* with an actuator at its axis and a mass *m* attached at the end. The arm itself is weightless, and the actuator is located at its pivot point. The dynamics of the pendulum are influenced by gravitational forces, characterized

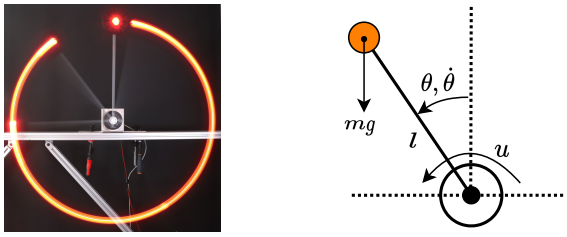

(a) (b)

Fig. 1: An Inverted Pendulum **(a)** and Its Schematic **(b)**

by the gravitational constant g, damping forces with a damping constant b, and the actuating torque u. To address the real-world constraint of limited actuator output, the actuator torque is bounded by $|u| \leq \dot{u}$, making the system under-actuated.

The *Equation of Motion* (EOM) for the pendulum, which describes the angular acceleration $\ddot{\theta}$, is given by:

$$\ddot{\theta} = \frac{1}{ml^2}(mgl\sin(\theta) - b\dot{\theta} + u) \tag{1}$$

where θ represents the angle from the vertical and $\dot{\theta}$ (or w) is the angular velocity. This equation accounts for the forces acting on the pendulum due to gravity, damping, and actuation.

To design the LQR controller, the system's behavior is linearized around the unstable equilibrium point at $\theta = 0$ and $\dot{\theta} = 0$. Furthermore, the state vector is defined as $x = (\theta, \dot{\theta})^T$ to describe the linearized system in the form of a state-space model $\dot{x} = Ax + Bu$. The resulting set of first-order *ordinary differential equations* (ODEs) are given as:

$$\dot{x} = \begin{pmatrix} 0 & 1 \\ \frac{g}{l} & -\frac{b}{ml^2} \end{pmatrix} x + \begin{pmatrix} 0 \\ \frac{1}{ml^2} \end{pmatrix} u \tag{2}$$

Here, matrix *A* contains the parameters governing the effects of gravity and damping, and matrix B accounts for the input due to the actuating torque.

For such a continuous-time linear system, a quadratic cost function J is defined as follows:

$$J = \int_0^\infty \left(x^T Q x + u^T R u\right) dt \tag{3}$$

where Q is a matrix that weights the state vector, and R is a matrix that weights the control input in the cost function. The objective is to determine the control input u that minimizes this cost.

The optimal control input u can be determined using the feedback control law $u = -Kx$, where the gain matrix $K = [K_0, K_1]$ is calculated as:

$$K = R^{-1}B^T S \tag{4}$$

Here, S is the solution to the continuous-time Algebraic Riccati Equation (ARE):

$$A^T S + SA - SBR^{-1}B^T S + Q = 0 \tag{5}$$

which is a fundamental equation in control theory for determining the optimal state feedback.

Integrating the LQR controller into Eq. (1), we get the closed-loop EOM as:

$$\ddot{\theta} = -\frac{1}{ml^2}\left(mg\ell\sin(\theta) - K_0\theta - (b + K_1)\dot{\theta}\right) \tag{6}$$

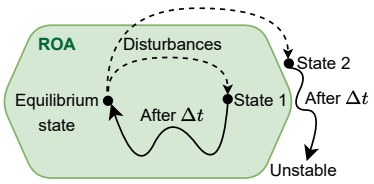

Fig. 2: Illustration of Region-of-Attraction

The actuator torque is bounded by $|u| \leq \bar{u}$. We can rearrange the EOM under these constraints as follows:

$$\ddot{\theta} = \begin{cases} \frac{1}{ml^2}(mg\ell\sin(\theta) - b\dot{\theta} - \bar{u}) & u(x) < -\bar{u} \\ \frac{1}{ml^2}(mg\ell\sin(\theta) - b\dot{\theta} + \bar{u}) & u(x) > \bar{u} \\ \frac{1}{ml^2}(mg\ell\sin(\theta) - K_0\theta - (b+K_1)\dot{\theta}) & \text{else} \end{cases} \tag{7}$$

B. Region-of-Attraction

In this section, we formally define the ROA and introduce relevant notation. The ROA includes all states from which a dynamic system can return to a stable equilibrium after disturbances, as illustrated in Figure 2. If a system is perturbed but remains within the ROA, it is capable of naturally reestablishing equilibrium [13]. A larger ROA indicates a greater ability of the system to maintain stability after disturbances, reflecting enhanced transient stability.

The equilibrium point is stable overall for the system given in Eq. (6), but only locally stable for the nonlinear, torque-limited system given in Eq. (7). Therefore, we define the ROA $\mathscr{S}$:

$$\mathscr{S} = \left\{ y \,\middle|\, \lim_{t\to\infty} x_y(t) \to 0 \right\}. \tag{8}$$

where all initial states will eventually return to the origin as time goes to infinity.

C. Lyapunov-Based ROA Estimation

Lyapunov functions (LFs) are scalar functions fundamental to stability theory in dynamic systems and control theory, as they can prove the stability of an equilibrium in ODEs. Specifically, the existence of an LF is both a necessary and sufficient condition for ensuring stability.

Furthermore, LFs are closely associated with invariant sets, defined as sets that, once entered, are never exited. Specifically, any sublevel set of an LF is also an invariant set. This relation allows for the use of LF sublevel sets to approximate the ROA for nonlinear systems [2]. The ROA for a locally attractive fixed point x^* is defined as the largest subset $D \subseteq \mathbb{R}^n$ such that if $x(0) \in D$, then $\lim_{t\to\infty} x(t) = x^*$.

To analyze the stability of the system $\dot{x} = f(x)$ around the origin, we can introduce a scalar LF, $V(x)$, where

$$\begin{aligned} x \neq 0 &\longrightarrow V(x) > 0 \\ V(0) &\geq 0 \end{aligned} \tag{9}$$

A bounded sublevel set $\mathcal{G}(\rho)$ of $V(x)$ can be defined as:

$$\mathcal{G}(\rho) = \{x \,|\, V(x) \leq \rho\} \tag{10}$$

where ρ is a positive scalar. Within this $\mathcal{G}(\rho)$, if the derivative of V with respect to time, denoted $\dot{V}(x)$, satisfies the condition

$$x \in \mathcal{G}(\rho) \longrightarrow \dot{V}(x) \leq 0$$

then $\mathcal{G}(\rho)$ is an invariant set under the system dynamics. Furthermore, if

$$\begin{aligned} x \in \mathcal{G}(\rho) \setminus \{0\} &\longrightarrow \dot{V}(x) < 0 \\ x = 0 &\longrightarrow \dot{V}(x) = 0 \end{aligned} \tag{11}$$

then the origin is asymptotically stable, and $\mathcal{G}$ is a subset of the ROA.

IV. Methodology

In this section, we introduce our methodology to estimate the ROA.

A. Proposed Function

We define our LF for LQR-controlled systems.

Existing Lyapunov-based ROA approaches often rely on purely quadratic LFs or optimization-based constructions, which can yield conservative estimates or require expensive computations. In contrast, our approach combines system-informed energy terms with the LQR cost-to-go to construct a Lyapunov candidate that better captures nonlinear dynamics while remaining computationally tractable.

Since the LF is supposed to represent a distance to the target equilibrium point (Eq. (9)) and should decrease over time (Eq. (11)), it is beneficial to include the required energy change necessary for the system to reach the equilibrium point. The total required energy change (E_{req}) is a function of the required potential energy gain $E_{\text{P,req}}(\theta)$ and the kinetic energy $E_{\text{K}}(\dot{\theta})$.

The required potential energy gain for an inverted pendulum is found from the potential energy function $E_{\text{P}}(\theta) = mgl\cos(\theta)$ as:

$$\begin{aligned} E_{\text{P,req}}(\theta) &= mgl\cos(0) - mgl\cos(\theta) \\ &= mgl(1-\cos(\theta)) \end{aligned} \tag{12}$$

and the kinetic energy is given as:

$$E_{\text{K}}(\dot{\theta}) = \frac{1}{2}ml^2\dot{\theta}^2 \tag{13}$$

The total required energy change is found by $E_{\text{P,req}}(\theta) + E_{\text{K}}(\dot{\theta})$ if θ and $\dot{\theta}$ have the same sign, and by $|E_{\text{P,req}}(\theta) - E_{\text{K}}(\dot{\theta})|$ otherwise. However, a function based on such an if-else condition is not differentiable. To achieve a differentiable function with a similar effect, we start by representing the signs of the energies with the sign function sgn:

$$E_{\text{req}}(\theta,\dot{\theta}) = |\,\text{sgn}(\theta)\,E_{\text{P,req}}(\theta) + \text{sgn}(\dot{\theta})\,E_{\text{K}}(\dot{\theta})|$$

where sgn is defined as:

$$\text{sgn}(x) = \begin{cases} 1 & x > 0 \\ 0 & x = 0 \\ -1 & x < 0 \end{cases}$$

Since the absolute value function is also not differentiable at zero, we approximate it by taking the square root of the two energies and squaring the result:

$$\tilde{E}_{\text{req}}(\theta,\dot{\theta}) = \left(\text{sgn}(\theta)\sqrt{E_{\text{P,req}}(\theta)} + \text{sgn}(\dot{\theta})\sqrt{E_{\text{K}}(\dot{\theta})}\right)^2$$

To eliminate the sgn functions, we derive from Eq. (13) that

$$\text{sgn}(\dot{\theta})\sqrt{E_{\text{K}}(\dot{\theta})} = \sqrt{\frac{1}{2}m}\,l\dot{\theta}$$

and, using the identity $1-\cos(\theta) = 2\sin^2(\theta/2)$, we derive from Eq. (12) that

$$\operatorname{sgn}(\theta)\sqrt{E_{\mathrm{P,req}}(\theta)} = \sqrt{2mgl}\ \sin\Big(\frac{\theta}{2}\Big) \quad \text{if} \quad -2\pi \le \theta \le 2\pi$$

With these, we can find the approximated required energy change $\tilde{E}_{\mathrm{req}}$ as:

$$\tilde{E}_{\mathrm{req}}(\theta,\dot{\theta}) = \Big(\sqrt{2mgl}\ \sin\Big(\frac{\theta}{2}\Big) + \sqrt{\frac{1}{2}m}\ l\dot{\theta}\Big)^2$$

When used as an LF, this function will satisfy the positiveness condition in Eq. (9), but not the negativeness condition in Eq. (11), since the LQR controller is not designed to reduce the required energy change. To address this, a quadratic function based on the LQR controller is added to $\tilde{E}_{\mathrm{req}}$ for our LF, V_{ours}, as follows:

$$V_{\mathrm{ours}}(\theta,\dot{\theta}) = \tilde{E}_{\mathrm{req}}(\theta,\dot{\theta}) + k_q\, x^T S x \tag{14}$$

Here, $x^T S x$ is the quadratic cost-to-go function of the LQR controller. The scaling factor k_q is added for further flexibility. The quadratic cost-to-go function is a viable candidate LF on its own [14], since it is a positive-definite quadratic form and locally satisfies the condition given in Eq. (11) around the origin. Therefore, the addition of this term to $\tilde{E}_{\mathrm{req}}$ will not affect the positiveness condition, and will create a local region that satisfies the condition given in Eq. (11) around the origin.

B. Targeted-Sampling-Based ROA Estimation

We established in Section III-C that any sublevel set $\mathcal{G}(\rho)$ of V_{ours} (Eq. (14)) that satisfies the condition in Eq. (11) is a subset of the actual ROA ($\mathscr{S}$). Since we want the estimation to be as close as possible to $\mathscr{S}$, we need to find the largest possible $\mathcal{G}(\rho)$, $\mathscr{S}_{\mathrm{ours}}$.

For this purpose, we use a variant of the sampling-based memoryless algorithm presented in [12]. In short, the algorithm starts with a large ρ (as defined in Eq. (10)) and samples random points in the state-space. For every sample point x_s, if $\dot{V}_{\mathrm{ours}}(x_s) \ge 0$, we reduce ρ to $V_{\mathrm{ours}}(x_s)$ if $\rho > V_{\mathrm{ours}}(x_s)$.

The sampling can be made uniformly in the space spanned by all possible initial states (broad sampling), but the convergence of the algorithm can be improved by not sampling in areas known to be outside the ROA (targeted sampling). It was shown in [12] that this approach improves the convergence, but the presented targeted-sampling method in the paper is based on polar coordinates and is only applicable for ellipsoidal ROAs.

For our approach, we use the inverse transform sampling method. This method uses a *Cumulative Distribution Function* (CDF), and an inverse of the CDF. Shortly, it generates a random number u between 0 and 1, and calculates the point x where the CDF is equal to u. Fortunately, our LF can be used as a CDF, since it is positive everywhere, is zero at the origin, and is decreasing towards the origin. Since we want to sample everywhere inside $\mathcal{G}(\rho)$, we generate the number u between 0 and ρ.

There are two challenges with this approach. First, V_{ours} is not dependent on one, but two variables. Second, the inverse of V_{ours} is not analytically solvable. To address the first challenge, we use a second random number ϕ between $-\pi/2$ and $\pi/2$, such that

$$\dot{\theta} = k\theta \quad \text{where} \quad k = \tan(\phi)$$

which we can substitute into V_{ours} to get

$$V_{\mathrm{ours}}(\theta,k) = p_1k^2\theta^2 + p_2k\theta^2 + p_3k\theta\sin\Big(\frac{\theta}{2}\Big) + p_4\theta^2 + p_5(1-\cos(\theta))$$

where the constants p_1 to p_5, which depend on the system parameters and S, are introduced for simplicity, and p_3 and p_5 are always positive.

To address the second challenge, we notice that it is permissible to sample slightly outside $\mathcal{G}(\rho)$, at the expense of some efficiency. This happens when we use a lower approximation of V_{ours} that still evaluates to zero at the origin.

First, we set the $p_5(1-\cos(\theta))$ term, which is always non-negative, to zero:

$$V_{\mathrm{ours}}(\theta,k) \ge p_1k^2\theta^2 + p_2k\theta^2 + p_3k\theta\sin\left(\frac{\theta}{2}\right) + p_4\theta^2$$

Next, we derive from the identity

$$\frac{\theta^2}{2} - \frac{\theta^4}{48} \le \theta\sin\left(\frac{\theta}{2}\right) \le \frac{\theta^2}{2}$$

that

$$k\theta\sin\left(\frac{\theta}{2}\right) > \begin{cases} k\left(\frac{\theta^2}{2} - \frac{\theta^4}{48}\right) & \text{if} \quad k > 0 \\ k\frac{\theta^2}{2} & \text{if} \quad k \le 0 \end{cases}$$

With this, we can define a lower bound of V_{ours} separately for $k > 0$ and $k \le 0$:

$$V_{\mathrm{lower,pos}}(\theta,k) = p_1k^2\theta^2 + p_2k\theta^2 + p_3k\Big(\frac{\theta^2}{2} - \frac{\theta^4}{48}\Big) + p_4\theta^2,\ k > 0$$

$$V_{\mathrm{lower,neg}}(\theta,k) = p_1k^2\theta^2 + p_2k\theta^2 + p_3k\frac{\theta^2}{2} + p_4\theta^2,\qquad k \le 0$$

and, denoting the result of the function as v, solve for θ to get the inverse transform sampling function:

$$\theta(v,k) = \begin{cases} V^{-1}_{\mathrm{lower,pos}}(v,k) & \text{if} \quad k > 0 \\ V^{-1}_{\mathrm{lower,neg}}(v,k) & \text{if} \quad k \le 0 \end{cases} \tag{15}$$

We use a symbolic solver to invert the functions. The inverted functions return both a positive and negative value for θ. We randomly choose one and calculate the corresponding $\dot{\theta}$ with $k\theta$, to get our sample point.

If the inverse sampling procedure produces samples outside the true invariant sublevel set, soundness is preserved, since such samples can only decrease the estimated level ρ, yielding a more conservative certified ROA.

V. Experimental Evaluation

We demonstrate our method's applicability and efficiency on a torque-limited inverted pendulum system and compare it to other approaches.

A. Pendulum Parameters

We selected three parameter sets to represent different system configurations for the pendulum, as seen in Table I. The 'normal' set contains the measured parameters of the physical system. With these parameters, the maximum torque that gravity can exert on the system, specifically at a 90° angle,

TABLE I: Inverted Pendulum Parameters and k_q Values

	mass (kg)	length (m)	inertia (kg m^2)	$\bar{u}=\tilde{u}/2$	$\bar{u}=\tilde{u}/4$	$\bar{u}=\tilde{u}/8$
normal	0.676	0.45	0.137	0.04	0.04	0.08
long	0.174	1.744	0.531	0.04	0.04	0.04
short	1.744	0.174	0.0531	0.2	0.2	0.2

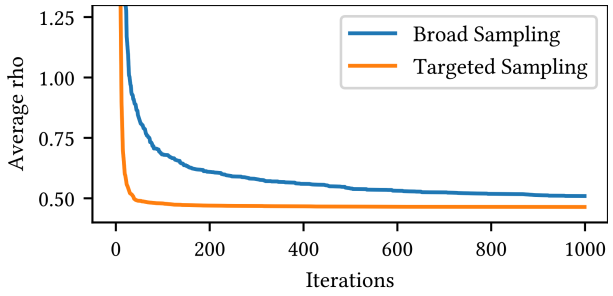

Fig. 3: Convergence Comparison

is calculated as $\tilde{u} := mgl \approx 2.98$ Nm.

For the other two sets, we modified the pendulum length while keeping the maximum gravity torque $\tilde{u}$ constant, thus maintaining actuator torques within a practical range. We examined a 'long' system with $\ell/m = 0.1$ and a 'short' system with $\ell/m = 10$. Additionally, we simulated the underactuation by limiting the controllers torque output to three separate values: $\tilde{u}/2$, $\tilde{u}/4$, and $\tilde{u}/8$ for a total of 9 distinct scenarios.

In the experiments, the cost parameters of the LQR controller are chosen as $Q = \text{diag}(1, 1)$ and $R = 1$. We assumed the damping coefficient to be approximately $b \approx 0.1$ Nm s/rad. We did not account for Coulomb friction in our calculations.

Furthermore, we set the LF scaling factor k_q separately for each scenario to ensure a balanced contribution between the required energy change and the quadratic cost-to-go function. This parameter allows balancing the contribution of the energy-based and quadratic components; in this work it is selected empirically to illustrate the effectiveness of the composite Lyapunov construction.

B. Evaluation of Targeted-Sampling-Based ROA Estimation

In this section, we demonstrate the effectiveness of the targeted-sampling method, by comparing it with the broad-sampling method. For this purpose, we used the normal pendulum system with the torque limitation set to $\tilde{u}/2$. We then used V_{ours} as our LF and estimated the ROA using both the targeted-sampling and broad-sampling methods, performing 1000 iterations for each method.

The targeted-sampling method focuses samples on more relevant regions of the state space than broad sampling. The efficiency difference is shown in Figure 3. For a fair comparison, we ran both methods a hundred times and averaged the ρ values at each iteration. As shown, targeted sampling converges much faster than broad sampling.

C. Evaluation of Methodology and Comparison

We evaluate our methodology by first describing the experimental setup, followed by explaining our approach to approximating the ROAs, and finally, presenting and interpreting the results. To evaluate the performance and validity of our ROA estimation method, we compared it with the quadratic-based LF $\mathscr{S}_{\text{quadratic}}$ [12] and an analytic approach $\mathscr{S}_{\text{analytic}}$ [3].

To compare our method with $\mathscr{S}_{\text{quadratic}}$, we consider a candidate LF, specifically the cost-to-go function, which is given by $x^T S x$. We define the largest set that guarantees stability by identifying the boundary ρ through state-space sampling and checking the Lyapunov conditions, as detailed in [12].

For the comparison with $\mathscr{S}_{\text{analytic}}$, we use the equation from [3] as:

$$\begin{aligned}\mathscr{S}_{\text{analytic}} = \quad & \{\gamma \mid mgl\,|\sin(\theta) - \theta| \le \bar{u}\} \\ \cap\, & \{\gamma \mid |u_{\text{lin},\gamma}(0)| \le \bar{u}\} \\ \cap\, & \{\gamma \mid t^*_\gamma > 0 \Rightarrow |u_{\text{lin},\gamma}(t^*_\gamma)| \le \bar{u}\}\end{aligned} \tag{16}$$

where the unknown variables and functions are calculated as in [3] and γ represents the initial conditions.

To approximate $\mathscr{S}_{actual}$, since it cannot be computed directly, we conducted 100,000 simulations for each parameter set. Initial values were randomly selected within the ranges of $-\pi < \theta < \pi$ and $-10\,\text{rad/s} < \dot{\theta} < 10\,\text{rad/s}$. We numerically integrated the torque-limited EOM given in Eq. (7), using a Runge-Kutta algorithm, with a step size of 0.1 seconds and a total duration of 10 seconds. Initial values were considered within the set $\mathscr{S}_{actual}$ if the solution converged to the origin with a precision of 5 digits.

To evaluate $\mathscr{S}_{analytic}$, $\mathscr{S}_{quadratic}$, and $\mathscr{S}_{ours}$, we used the same initial conditions as for $\mathscr{S}_{actual}$. We calculated the oracle functions for each method that determine whether a given initial condition falls within the respective ROA. For $\mathscr{S}_{analytic}$, the oracle function is a combination of all the conditions given in the set builders in Eq. (16). For $\mathscr{S}_{quadratic}$ and $\mathscr{S}_{ours}$, the oracle functions are derived from the respective LFs and calculated ρ values, such that the initial condition is within the ROA if the LF is less than ρ.

We compare our method with $\mathscr{S}_{actual}$ and the other approaches, [3] and [12], for estimating the ROA shown in Figure 4, across nine distinct scenarios, using a range of parameters and torque limits. In the figure, our method $\mathscr{S}_{ours}$ is represented in green, the quadratic-based method $\mathscr{S}_{quadratic}$ in yellow, the analytic approach $\mathscr{S}_{analytic}$ in blue, and the actual ROA $\mathscr{S}_{actual}$ in grey. Notably, the controller stabilized all states within the estimated ROA for all methods.

Overall, the results show that the proposed method yields less conservative certified ROA estimates than the considered analytic and quadratic Lyapunov-based baselines. We also report the computation times required to configure the oracle functions for each method using new parameter sets, excluding the time required to generate Figure 4 from oracle evaluations over multiple initial conditions. While the analytic approach is the fastest (0.00084 s), it produces more conservative ROA estimates and does not provide formal guarantees. The quadratic-based method is slightly faster (0.126 s) than the proposed approach (0.143 s), but yields significantly more conservative certified ROAs.

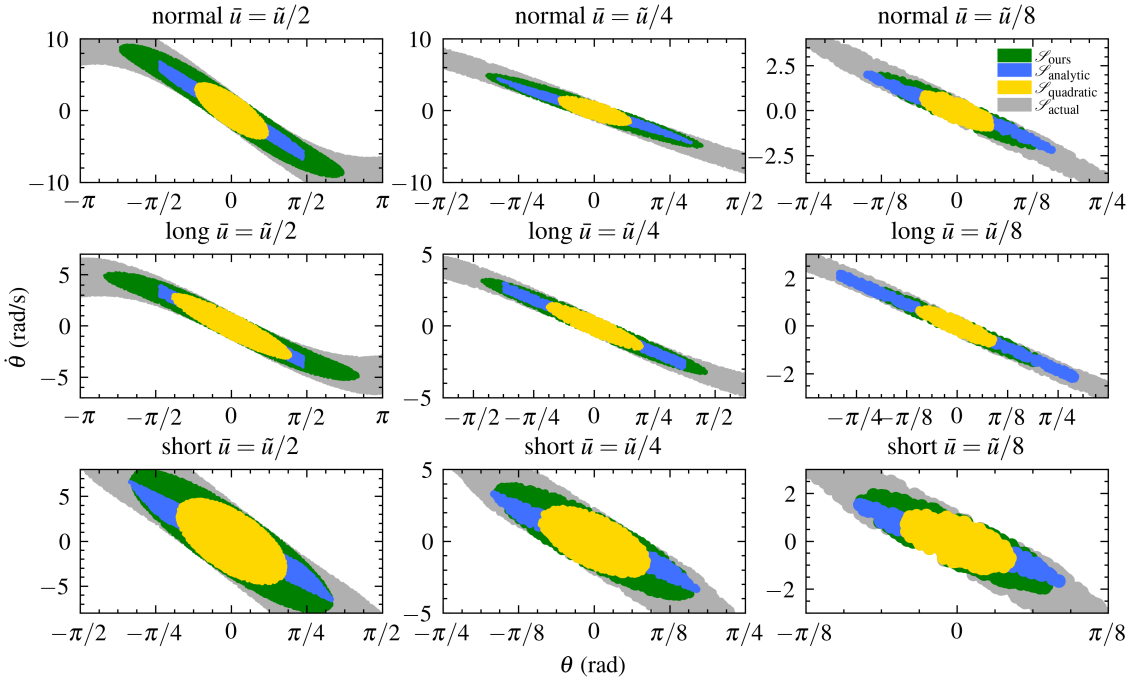

Fig. 4: Comparison of ROA Estimation for Torque-Limited Inverted Pendulums Across Nine Distinct Scenarios

VI. Conclusion

In this paper, we presented a Lyapunov-based method for computing certified regions of attraction for constrained nonlinear systems under LQR control. The proposed approach combines a problem-informed energy-like term with the LQR cost-to-go to construct a composite Lyapunov function and employs a targeted sampling strategy to estimate invariant sublevel sets with formal guarantees. The method was evaluated on a torque-limited inverted pendulum benchmark, where it produced certified ROAs that are less conservative than those obtained using standard quadratic Lyapunov functions and analytic baselines. While demonstrated on a torque-limited inverted pendulum, the proposed framework can be extended to other nonlinear systems by incorporating problem-specific energy-like terms reflecting the underlying system dynamics. These results demonstrate that the proposed approach provides a practical and sound framework for certified ROA estimation within the considered scope, making it suitable for safety-critical cyber-physical systems requiring certified stability guarantees.

References

[1] V. Bolbot, G. Theotokatos, L. M. Bujorianu, E. Boulougouris, and D. Vassalos, "Vulnerabilities and safety assurance methods in Cyber-Physical Systems: A comprehensive review," *Reliability Engineering & System Safety*, vol. 182, pp. 179–193, Feb. 2019.

[2] H. Khalil, *Nonlinear Systems (3rd Ed.)*. Upper Saddle River NJ: Prentice-Hall, 2002.

[3] L. Gross, L. Maywald, S. Kumar, F. Kirchner, and C. Lüth, "Analytic Estimation of Region of Attraction of an LQR Controller for Torque Limited Simple Pendulum," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. Cancun, Mexico: IEEE, Dec. 2022, pp. 2695–2701.

[4] R. Iervolino, D. Tangredi, and F. Vasca, "Lyapunov stability for piecewise affine systems via cone-copositivity," *Automatica*, vol. 81, pp. 22–29, Jul. 2017.

[5] Z. W. Jarvis-Wloszek, *Lyapunov Based Analysis and Controller Synthesis for Polynomial Systems Using Sum-of-Squares Optimization*. Berkeley: University of California, 2003.

[6] G. Valmórbida, S. Tarbouriech, and G. Garcia, "Region of attraction estimates for polynomial systems," in *Proceedings of the 48h IEEE Conference on Decision and Control (CDC) Held Jointly with 2009 28th Chinese Control Conference*. Shanghai, China: IEEE, Dec. 2009, pp. 5947–5952.

[7] A. El-Guindy, D. Han, and M. Althoff, "Estimating the region of attraction via forward reachable sets," in *2017 American Control Conference (ACC)*. Seattle, WA, USA: IEEE, May 2017, pp. 1263–1270.

[8] L. Jin, R. Kumar, and N. Elia, "Reachability analysis based transient stability design in power systems," *International Journal of Electrical Power & Energy Systems*, vol. 32, no. 7, pp. 782–787, Sep. 2010.

[9] S. Chen, M. Fazlyab, M. Morari, G. J. Pappas, and V. M. Preciado, "Learning Region of Attraction for Nonlinear Systems," in *2021 60th IEEE Conference on Decision and Control (CDC)*. Austin, TX, USA: IEEE, Dec. 2021, pp. 6477–6484.

[10] A. I. Zečević and D. D. Šiljak, "Estimating the region of attraction for large-scale systems with uncertainties," *Automatica*, vol. 46, no. 2, pp. 445–451, Feb. 2010.

[11] U. Topcu, A. Packard, and P. Seiler, "Local stability analysis using simulations and sum-of-squares programming," *Automatica*, vol. 44, no. 10, pp. 2669–2675, Oct. 2008.

[12] E. Najafi, R. Babuška, and G. A. D. Lopes, "A fast sampling method for estimating the domain of attraction," *Nonlinear Dynamics*, vol. 86, no. 2, pp. 823–834, Oct. 2016.

[13] F. Bellizio, J. L. Cremer, and G. Strbac, "Transient Stable Corrective Control Using Neural Lyapunov Learning," *IEEE Transactions on Power Systems*, vol. 38, no. 4, pp. 3245–3253, Jul. 2023.

[14] B. D. O. Anderson and J. B. Moore, *Optimal Control: Linear Quadratic Methods*. New Jersey, USA: Prentice-Hall, 1989.