

Learning Certified Control Using Contraction Metric

Dawei Sun

SRI International and
University of Illinois at Urbana-Champaign, US
daweis2@illinois.edu

Susmit Jha

Computer Science Laboratory
SRI International, US
susmit.jha@sri.com

Chuchu Fan

Massachusetts Institute of Technology
United States
chuchu@mit.edu

Abstract: In this paper, we solve the problem of finding a certified control policy that drives a robot from any given initial state and under any bounded disturbance to the desired reference trajectory, with guarantees on the convergence or bounds on the tracking error. Such a controller is crucial in safe motion planning. We leverage the advanced theory in Control Contraction Metric and design a learning framework based on neural networks to co-synthesize the contraction metric and the controller for control-affine systems. We further provide methods to validate the convergence and bounded error guarantees. We demonstrate the performance of our method using a suite of challenging robotic models, including models with learned dynamics as neural networks. We compare our approach with leading methods using sum-of-squares programming, reinforcement learning, and model predictive control. Results show that our methods indeed can handle a broader class of systems with less tracking error and faster execution speed. Code is available at <https://github.com/sundw2014/C3M>.

Keywords: Learning Tracking Controller, Control Contraction Metric, Convergence Guarantee, Neural Networks

1 Introduction

Designing safe motion planning controllers for nonholonomic robotic systems is a critical yet extremely challenging problem. In general, motion planning for robots is notoriously difficult. For example, even planning for a robot (composed of n polyhedral bodies) in a 3D world that contains a fixed number of polyhedral obstacles is PSPACE-hard [1]. Simultaneously providing formal guarantees on safety and robustness when the robots are facing uncertainties and disturbances makes the problem even more challenging. Sample-based planning techniques [2, 1] can plan safe trajectories by exploring the environment but cannot handle uncertainty and disturbances. A natural thought that has been extensively explored is to exploit a separation of concerns that exists in the problem: (A) how to design a reference (also called expected or nominal) trajectory $x^*(t)$ to drive a robot to its goal safely without considering any uncertainty or disturbance, and (B) how to design a tracking controller to make sure the actual trajectories $x(t)$ of the system under disturbances can converge to $x^*(t)$ with guaranteed bounds for the tracking error between $x(t)$ and $x^*(t)$ [3, 4, 5, 6]. Combining controllers from solving (A) and (B) can make sure that the actual behaviors of the robot are all safe.

While the reference trajectories can be found using efficient planning techniques [2, 1, 7] for a very broad class of systems, the synthesis of a guaranteed tracking controller, also called trajectory stabilization [8] is not an automatic process, especially for nonholonomic systems. Control theoretic techniques exist to guide the dual synthesis [6] of control policies and certificates at the same time, where the certificates can make sure certain required properties are provably satisfied when the corresponding controller is run in (closed-loop) composition with the plant. For example, control Lyapunov Function (CLF) [9, 10] ensures the existence of a controller so the controlled system is

Lyapunov stable, and Control Barrier Function (CBF) [11, 12, 13] ensures that a controller exists so the controlled system always stays in certain safety invariant sets defined by the corresponding barrier function. Other methods pre-compute the tracking error through reachability analysis [4], Funnels [14, 5], and Hamilton-Jacobi analysis [3, 15]. Despite the benefit brought by these certificates (e.g. Lyapunov Function, Barrier Function), finding the correct function representation for the certificates is non-trivial. Various methods have been studied to learn the certificate as polynomials [16, 10], support vectors [17], Gaussian processes [18], temporal logic formulae [19, 20], and neural networks (NN) [9, 13, 21, 22, 23]. Unlike reinforcement learning (RL) [24, 25, 26, 27] which focuses on learning a control policy that maximizes an accumulated reward and usually lacks formal safety guarantees, certificate-guided controller learning focuses on learning a sufficient condition for the desired property. In this paper, we follow this idea and learn a tracking controller by learning a convergence certificate simultaneously as guidance.

In this paper, we leverage recent advances in Control Contraction Metric (CCM) [28, 29] theory that extends the contraction theory to control-affine systems to prove the existence of a tracking controller so the closed-loop system is contracting. In the CCM theory [28], it has been shown that a valid CCM implies the existence of a contracting tracking controller for *any* reference trajectories. In our framework, we model both the controller and the metric function as neural networks and optimize them jointly with the cost functions inspired by the CCM theory. Due to the constraints imposed during training, the tracking error of the learned controller is guaranteed to be bounded even with external disturbances.

Synthesis of CCM certificate has been formulated as solving a Linear Matrix Inequality (LMI) problem using Sum-of-Squares (SoS) programming [30] or RKHS theory [8] even when the model dynamics is unknown. However, such LMI-based techniques have to rely on an assumption on the special control input structure of a class of underactuated systems (see Sec. 3.1 for details) and therefore cannot be applied to general robotic systems. Moreover, the above methods only learn the CCM. The controller needs to be found separately by computing geodesics of the CCM, which cannot be solved exactly and has high computational complexity. To use SoS, the system dynamics need to be polynomial equations or can be approximated as polynomial equations. The degree and template of the polynomials in SoS play a crucial role in determining whether a solution exists and need to be tuned for each system. In [31], the authors proposed a synthesis framework using recurrent NNs to model the contraction metric and this framework works for nonlinear systems which are a convex combination of multiple state-dependent coefficients (i.e. $f(x, t)$ is written as $A(x, t)x$). Again, the controller in [31] is constructed from the learned metric. In contrast, our approach can simultaneously synthesize the controller and the CCM certificate for control-affine systems without any additional structural assumptions.

We provide two methods to prove convergence guarantees of the learned controller and CCM (Section 3.2). The first method provides deterministic guarantees by leveraging the Lipschitz continuity of the CCM’s condition. Furthermore, we observed that even if the CCM’s condition does not hold globally for every state in the state space, the resulting trajectories still often converge to the reference trajectories in our experience. This motivates our second approach based on conformal regression to give probabilistic guarantees on the convergence of the tracking error. Both methods can provide upper bounds on the tracking error, using which one can explore trajectory planning methods such as sampling-based methods (e.g. RRT [32], PRM [2]), model predictive control (MPC) [7], and satisfiability-based methods [33] to find safe reference trajectories to accomplish the safe motion planning missions. We compare our approach Certified Control using Contraction Metric (C3M) with the SoS programming [30], RL [34], and MPC [7, 35, 36] on several representative robotic systems, including the ones whose dynamics are learned from physical data using neural networks. We show that C3M outperforms other methods by being the only approach that can find converging tracking controllers for all the benchmarks. We provide two metrics for evaluating the tracking performance and show that controllers achieved through C3M have a much smaller tracking error. In addition, controllers found by C3M can be executed in sub-milliseconds at each step and is much faster than methods that only learn the metric (e.g. SoS programming) and online control methods (e.g. MPC).

2 Preliminaries and notations

We denote by \mathbb{R} and $\mathbb{R}^{\geq 0}$ the set of real and non-negative real numbers respectively. For a symmetric matrix $A \in \mathbb{R}^{n \times n}$, the notation $A \succ 0$ ($A \prec 0$) means A is positive (negative) definite. The set of positive definite $n \times n$ matrices is denoted by $S_n^{>0}$. For a matrix-valued function $M(x) : \mathbb{R}^n \mapsto \mathbb{R}^{n \times n}$, its element-wise Lie derivative along a vector $v \in \mathbb{R}^n$ is $\partial_v M := \sum_i v^i \frac{\partial M}{\partial x^i}$. Unless otherwise stated, x^i denotes the i -th element of vector x . For $A \in \mathbb{R}^{n \times n}$, we denote $A + A^\top$ by \hat{A} .

Dynamical system We consider *control-affine* systems of the form

$$\dot{x}(t) = f(x(t)) + B(x(t))u(t) + d(t), \quad (1)$$

where $x(t) \in \mathcal{X} \subseteq \mathbb{R}^n$, $u(t) \in \mathcal{U} \subseteq \mathbb{R}^m$, and $d(t) \in \mathcal{D} \subseteq \mathbb{R}^n$ for all $t \in \mathbb{R}^{\geq 0}$ are states, inputs, and disturbances respectively. Here, \mathcal{X}, \mathcal{U} and \mathcal{D} are compact sets that represent the state, input, and disturbance space respectively. We assume that $f : \mathbb{R}^n \mapsto \mathbb{R}^n$ and $B : \mathbb{R}^n \mapsto \mathbb{R}^{n \times m}$ are smooth functions, the control input $u : \mathbb{R}^{\geq 0} \mapsto \mathcal{U}$ is a piece-wise continuous function, and the right-side of Eq. (1) holds at discontinuities. Furthermore, we assume that the disturbance is bounded, i.e. $\sup_t \|d(t)\|_2 < \infty$. Given an input signal $u : \mathbb{R}^{\geq 0} \mapsto \mathcal{U}$, a disturbance signal $d : \mathbb{R}^{\geq 0} \mapsto \mathcal{D}$, and an initial state $x_0 \in \mathcal{X}$, a (*state*) *trajectory* of the system is a function $x : \mathbb{R}^{\geq 0} \mapsto \mathbb{R}^n$ such that $x(t), u(t), d(t)$ satisfy Eq. (1) and $x(0) = x_0$. The goal of this paper is to design a tracking controller $u(\cdot)$ such that the controlled trajectory $x(t)$ can track *any* target trajectory $x^*(t)$ generated by a reference control $u^*(t)$ when $x(0)$ is in a neighborhood of $x^*(0)$ and there is no disturbance (i.e. $d(t) = 0$). Furthermore, when $d(t) \neq 0$, the tracking error can also be bounded (Sec. 3.3). Fig. 1 (a) illustrates the reference trajectory and the actual trajectory controlled by the proposed method.

Control contraction metric theory Contraction theory [37] analyzes the incremental stability of a system by considering the evolution of the distance between any pair of arbitrarily close neighboring trajectories. Let us first consider a time-invariant autonomous system of the form $\dot{x} = f(x)$. Given a pair of neighboring trajectories, denote the infinitesimal displacement between them by δ_x , which is also called a *virtual displacement*. The evolution of δ_x is dominated by a linear time varying (LTV) system: $\dot{\delta}_x = \frac{\partial f}{\partial x}(x)\delta_x$. Thus, the dynamics of the squared distance $\delta_x^\top \delta_x$ is given by $\frac{d}{dt}(\delta_x^\top \delta_x) = 2\delta_x^\top \dot{\delta}_x = 2\delta_x^\top \frac{\partial f}{\partial x} \delta_x$. If the symmetrical part of the Jacobian $\frac{\partial f}{\partial x}$ is uniformly negative definite, i.e. there exists a constant $\lambda > 0$ such that for all x , $\frac{1}{2} \widehat{\frac{\partial f}{\partial x}} \preceq -\lambda \mathbf{I}$, then $\delta_x^\top \delta_x$ converges to zero exponentially at rate 2λ . Hence, all trajectories of this system will converge to a common trajectory [37]. Such a system is referred to be *contracting*.

The above analysis can be generalized by introducing a *contraction metric* $M : \mathbb{R}^n \mapsto S_n^{>0}$, which is a smooth function. Then, $\delta_x^\top M(x)\delta_x$ can be interpreted as a Riemannian squared length. Since $M(x) \succ 0$ for all x , if $\delta_x^\top M(x)\delta_x$ converges to 0 exponentially, then the system is contracting. The converse is also true. As shown in [37], if a system is contracting, then there exists a contraction metric $M(x)$ and a constant $\lambda > 0$ such that $\frac{d}{dt}(\delta_x^\top M(x)\delta_x) < -\lambda \delta_x^\top M(x)\delta_x$ for all x and δ_x .

Contraction theory can be further extended to control-affine systems. First, let us ignore the disturbances in system (1), then the dynamics of the corresponding virtual system is given by $\dot{\delta}_x = A(x, u)\delta_x + B(x)\delta_u$, where $A(x, u) := \frac{\partial f}{\partial x} + \sum_{i=1}^m u^i \frac{\partial b_i}{\partial x}$, b_i is the i -th column of B , u^i is the i -th element of u , and δ_u is the infinitesimal difference of u . A fundamental theorem in Control Contraction Metric (CCM) theory [28] says if there exists a metric $M(x)$ such that the following conditions hold for all x and some $\lambda > 0$,

$$B_\perp^\top \left(-\partial_f W(x) + \widehat{\frac{\partial f(x)}{\partial x} W(x)} + 2\lambda W(x) \right) B_\perp \prec 0, \quad (2)$$

$$B_\perp^\top \left(\partial_{b_j} W(x) - \widehat{\frac{\partial b_j(x)}{\partial x} W(x)} \right) B_\perp = 0, \quad j = 1, \dots, m, \quad (3)$$

where $B_\perp(x)$ is an annihilator matrix of $B(x)$ satisfying $B_\perp^\top B = 0$, and $W(x) = M(x)^{-1}$ is the dual metric, then there exists a tracking controller $k(x, x^*, u^*)$ such that the closed-loop system controlled by $u = k(x, x^*, u^*)$ is contracting with rate λ under metric $M(x)$ [28], which means $x(t)$ converges to $x^*(t)$ exponentially. However, as shown in [28], either finding a valid metric or constructing a controller from a given metric is not straightforward. Hence, we proposed to jointly learn the metric and controller using neural networks.




Figure 1: Components of the proposed method. (a) Trajectories of the Dubins car model, which is a classical nonholonomic system. Here, x^* is the reference trajectory, and x is the actual trajectory controlled by our learned controller. (b) Joint training for the CCM and the controller. (c) Illustration on the online execution.

3 Learning tracking controllers using CCM with formal guarantees

In this paper, we utilize machine learning methods to jointly learn a Contraction Metric and a tracking controller for System (1) with *known* dynamics. The overall learning system is illustrated in Fig. 1 (b). Different from RL-based methods, since all the conditions used for learning (Eq. (2), (3), and (4)) are defined for a single state at a specific time instant instead of the whole trajectory, running of the system is not required for training. Hence, the training is fully offline and fast. The data set for training is of the form $\{(x_i, x_i^*, u_i^*) \in \mathcal{X} \times \mathcal{X} \times \mathcal{U}\}_{i=1}^N$, and samples are drawn independently. The contraction metric and the tracking controller are parameterized using neural networks. In what follows we first ignore the disturbance $d(t)$ and then show in Sec. 3.3 that with $d(t) \neq 0$, the tracking error of the learned controller can still be bounded.

3.1 Controller and metric learning

The controller $u(x, x^*, u^*; \theta_u)$ is a neural network with parameters θ_u , and the dual metric $W(x; \theta_w)$ is also a neural network with parameters θ_w . By design, the controller satisfies that $x = x^* \Rightarrow u(x, x^*, u^*; \theta_u) = u^*$ for all θ_u , and $W(x; \theta_w)$ is a symmetric matrix satisfying $W(x; \theta_w) \succeq \underline{w}\mathbf{I}$ for all x and all θ_w , where \underline{w} is a hyper-parameter. More details about the design can be found in Sec. 4. Plugging $u(x, x^*, u^*; \theta_u)$ into System (1) with $d(t) = 0$, the dynamics of the generalized squared length of the virtual displacement under metric $M(x; \theta_w) = W(x; \theta_w)^{-1}$ is given by $\frac{d}{dt}(\delta_x^T M(x; \theta_w) \delta_x) = \delta_x^T (\dot{M} + \overline{M(A + BK)}) \delta_x$, where $K = \frac{\partial u}{\partial x}$, and $\dot{M} = \sum_{i=1}^n \frac{\partial M}{\partial x^i} \dot{x}^i$, or using the Lie derivative notation $\dot{M} = \partial_{f(x)+B(x)u} M$. A sufficient condition from [28] for the closed-loop system being contracting is that there exists $\lambda > 0$ such that for all $(x, x^*, u^*) \in \mathcal{X} \times \mathcal{X} \times \mathcal{U}$,

$$\dot{M} + \overline{M(A + BK)} + 2\lambda M \prec 0. \quad (4)$$

Since $x = x^* \Rightarrow u(x, x^*, u^*; \theta_u) = u^*$, the reference $(x^*(t), u^*(t))$ is a valid trajectory of the closed-loop system. If the system is furthermore contracting, then starting from any initial state, the trajectory converges to a common trajectory which is indeed $x^*(t)$ as formally stated below.

Proposition 1. *If condition (4) holds and $\bar{m} \geq \underline{m} > 0$ satisfying $\underline{m}\mathbf{I} \preceq M(x) \preceq \bar{m}\mathbf{I}$ for all x , then the initial difference between the reference and the actual trajectory is exponentially forgotten, i.e. $\|x(t) - x^*(t)\|_2 \leq \sqrt{\bar{m}/\underline{m}} e^{-\lambda t} \|x(0) - x^*(0)\|_2$.*

Denoting the LHS of Eq.(4) by $C_u(x, x^*, u^*; \theta_w, \theta_u)$ and the uniform distribution over the data space $\mathcal{S} := \mathcal{X} \times \mathcal{X} \times \mathcal{U}$ by $\rho(\mathcal{S})$, then the *contraction risk* of the system is defined as

$$\mathcal{L}_u(\theta_w, \theta_u) = \mathbb{E}_{(x, x^*, u^*) \sim \rho(\mathcal{S})} L_{PD}(-C_u(x, x^*, u^*; \theta_w, \theta_u)), \quad (5)$$

where $L_{PD}(\cdot) \geq 0$ is for penalizing non-positive definiteness, and $L_{PD}(A) = 0$ iff. $A \succeq 0$. Obviously, $\mathcal{L}_u(\theta_w^*, \theta_u^*) = 0$ implies that $u(x, x^*, u^*; \theta_u^*)$ and $W(x; \theta_w^*)$ satisfy inequality (4) exactly.

As shown in Proposition 1, the overshoot of the tracking error is affected by the condition number \bar{m}/\underline{m} of the metric. Since the smallest eigenvalue of the dual metric is lower bounded by \underline{w} by design, penalizing large condition numbers is equivalent to penalizing the largest eigenvalues, and thus the following risk function is used

$$\mathcal{L}_c(\theta_w) = \mathbb{E}_{(x, x^*, u^*) \sim \rho(\mathcal{S})} L_{PD}(\underline{w}\mathbf{I} - W(x; \theta_w)), \quad (6)$$

where \bar{w} is a hyper-parameter.

However, jointly learning a metric and a controller by minimizing \mathcal{L}_u solely is hard and leads to poor results. Inspired by the CCM theory, we add two auxiliary loss terms for the dual metric to mitigate the difficulty of minimizing \mathcal{L}_u . As shown in Sec. 2, conditions (2) and (3) are sufficient for a metric to be a valid CCM. Intuitively, imposing these constraints on the dual metric $W(x; \theta_w)$ provides more guidance for optimization. More discussion can be found in Appendix. Denoting the LHS of Eq. (2) and (3) by $C_1(x; \theta_w)$ and $\{C_2^j(x; \theta_w)\}_{j=1}^m$, the following risk functions are used

$$\mathcal{L}_{w1}(\theta_w) = \mathbb{E}_{(x, x^*, u^*) \sim \rho(\mathcal{S})} L_{PD}(-C_1(x; \theta_w)); \mathcal{L}_{w2}(\theta_w) = \sum_{j=1}^m \mathbb{E}_{(x, x^*, u^*) \sim \rho(\mathcal{S})} \|C_2^j(x; \theta_w)\|_F, \quad (7)$$

where $\|\cdot\|_F$ is the Frobenius norm.

Please note that in some cases, condition (3) can be automatically satisfied for all θ_w by designing $W(\cdot)$ appropriately. As stated in [8], if the matrix $B(x)$ is sparse and of the form $B(x) = \begin{bmatrix} 0_{(n-m) \times m} \\ b(x) \end{bmatrix}$, where $b(x)$ is an invertible matrix, condition (3) can be automatically satisfied for all x and all θ_w by making the upper-left $(n-m) \times (n-m)$ block of $W(x; \theta_w)$ not a function of the last $(n-m)$ elements of x . In [8, 30], such sparsity assumptions are necessary for their approach to work. For dynamical models satisfying this assumption, we make use of this property by designing W to satisfy such a structure and eliminate \mathcal{L}_{w2} since it will always be 0. For models not satisfying this assumption, we will show in Sec. 4 the impact of the loss term \mathcal{L}_{w2} .

In order to train the neural network using sampled data, the following *empirical risk* function is used

$$\mathcal{L}(\theta_u, \theta_w) = \frac{1}{N} \sum_{i=1}^N \left[L_{PD}(-C_u(x_i, x_i^*, u_i^*; \theta_u, \theta_w)) + L_{PD}(-C_1(x_i; \theta_w)) \right. \\ \left. + \sum_{j=1}^m \|C_2^j(x_i; \theta_w)\|_F + L_{PD}(\bar{w}\mathbf{I} - W(x_i; \theta_w)) \right], \quad (8)$$

where $\{(x_i, x_i^*, u_i^*)\}_{i=1}^N$ are drawn independently from $\rho(\mathcal{S})$, and L_{PD} is implemented as follows. Given a matrix $A \in \mathbb{R}^{n \times n}$, we randomly sample K points $\{p_i \in \mathbb{R}^n \mid \|p_i\|_2 = 1\}_{i=1}^K$. Then, the loss function is calculated as $L_{PD}(A) = \frac{1}{K} \sum_{i=1}^K \min\{0, -p_i^T A p_i\}$.

3.2 Theoretical convergence guarantees

As discussed in Sec. 3.1, satisfaction of inequality (4) for all $(x, x^*, u^*) \in \mathcal{S}$ is a strong guarantee to show the validity of $M(x)$ as a contraction metric and therefore provide convergence guarantees. SMT solvers [9] are valid tools for verifying the satisfaction of (4) on the uncountable set \mathcal{S} . However, SMT solvers for NN have poor scalability and cannot handle the neural networks used in our experiments. Moreover, we observed in experiments that even if inequality (4) does not hold for all $(x, x^*, u^*) \in \mathcal{S}$, the learned controller still has perfect performance. Therefore, we introduce two approaches to provide guarantees on the satisfaction of inequality (4) or on the tracking error directly, which in practice are easier to validate and still give the needed level of assurance.

Deterministic guarantees using Lipschitz continuity. For a Lipschitz continuous function $f : \mathcal{X} \mapsto \mathbb{R}$ with Lipschitz constant L_f , discretizing the domain \mathcal{X} such that the distance between any grid point and its nearest neighbor is less than τ , if $f(x_i) < -L_f \tau$ holds for all grid point x_i , then $f(x) < 0$ holds for all $x \in \mathcal{X}$. We show in the following proposition that the largest eigenvalue of the LHS of inequality (4) indeed has a Lipschitz constant if both the dynamics and the learned controller and metric have Lipschitz constants.

Proposition 2. *Let A, B, K , and M be functions of x, x^* , and u^* . If \dot{M}, M, A, B , and K all have Lipschitz constants $L_{\dot{M}}, L_M, L_A, L_B$, and L_K respectively, and 2-norms of the last four are all bounded by S_M, S_A, S_B , and S_K respectively, then the largest eigenvalue function $\lambda_{\max}(\dot{M} + \overline{M(A+BK)} + 2\lambda M)$ has a Lipschitz constant $L_{\dot{M}} + 2(S_M L_A + S_A L_M + S_M S_B L_K + S_B S_K L_M + S_M S_K L_B + \lambda L_M)$.*

The proof of Proposition 2 is provided in Appendix. Combining discrete samples and the Lipschitz constant from Proposition 2 can guarantee the strict satisfaction of inequality (4) on the uncountable set \mathcal{S} . The Lipschitz constants for the metric and controller NNs can be computed using various existing tools such as the methods in [38]. An example demonstrating the verification process is given in Appendix. Since the estimate of the Lipschitz constant is too conservative, it usually requires a huge number of samples to verify the learned controller. Fortunately, this process only requires forward computation of the NN and thus can be done relatively fast.

Probabilistic guarantees. We observed that even if inequality (4) does not hold for all points in \mathcal{S} , the learned controller can still drive all simulated trajectories to converge to reference trajectories. This motivates us to derive probabilistic guarantees on the convergence as an alternative in the trajectory space. The probabilistic guarantee is derived based on a simple result from conformal prediction [39]. Let us consider the process of evaluating a tracking controller using a *quality metric function*. Given an *evaluation configuration* including the initial state $x(0)$, and the initial reference state $x^*(0)$ and reference control signal $u^*(t)$, we can get the actual trajectory $x(t)$ controlled by the tracking controller. The quality metric function summarizes the tracking error curve $\|x^*(t) - x(t)\|_2$ into a scalar which can be viewed as a score for the tracking. Now, it is clear that the quality metric function is a mapping from the evaluation configuration to a score. The following proposition gives a probabilistic guarantee on the distribution of the quality score based on empirical observations.

Proposition 3. *Given a set of n i.i.d. evaluation configurations, let $\{m_i\}_{i=1}^n$ be the quality scores for each configuration. Then, for a new i.i.d. evaluation configuration and the corresponding quality score m_{n+1} , we have $\Pr(m_{n+1} \geq q_{1-\alpha}) \leq 1 - \alpha$, where $q_{1-\alpha}$ is the $(1 - \alpha)$ -th quantile of $\{m_i\}_{i=1}^n$.*

Note that Proposition 3 asserts guarantees on the marginal distribution of m_{n+1} , which should be distinguished from the conditional distribution. Such probabilistic guarantees work on the trajectory space and therefore can be applied to any method that finds tracking controllers. In Sec. 4, we evaluate all the methods in comparison by computing the probabilistic guarantees using either the average tracking error or the convergence rate as the quality metric.

3.3 Robustness of the learned controller

When the closed-loop system run with disturbances, the tracking error is still bounded as follows.

Theorem 4. *Given $M(x)$ and $u(x, x^*, u^*)$ satisfying inequality (4), since $M(x) \in S_n^{>0}$, there exist $\bar{m} \geq \underline{m} > 0$ such that $\underline{m}\mathbf{I} \preceq M(x) \preceq \bar{m}\mathbf{I}$ for all x . Assume that the disturbance is uniformly bounded as $\|d(t)\|_2 \leq \epsilon$. Now, for the same reference, considering the trajectories $x_1(t)$ and $x_2(t)$ of the unperturbed and the perturbed closed-loop system respectively, the distance between these two trajectories can be bounded as $\|x_1(t) - x_2(t)\|_2 \leq \frac{R_0}{\sqrt{\underline{m}}} e^{-\lambda t} + \sqrt{\frac{\bar{m}}{\underline{m}}} \cdot \frac{\epsilon}{\lambda} (1 - e^{-\lambda t})$, where $R_0 = \int_{x_1(0)}^{x_2(0)} \sqrt{\delta_x^\top M(x) \delta_x}$ is the geodesic distance between $x_1(0)$ and $x_2(0)$ under metric $M(x)$.*

The proof of Theorem 4 is inspired by Theorem 1 in [31] and is provided in Appendix. In practice \bar{m} and \underline{m} can be computed using samples and the Lipschitz constant of the eigenvalues of $M(x)$ as discussed in Sec. 3.2.

4 Evaluation of performance

We evaluate the C3M approach on 9 representative case studies (5 of them are reported in Appendix), including high-dimensional (up to 9 state variables and 3 control variables) models and a system with learned dynamics represented by a neural network. All our experiments were conducted on a Linux workstation with two Xeon Silver 4110 CPUs, 32 GB RAM, and a Quadro P5000 GPU. Please note that the execution time reported in Tab. 1 were all evaluated on CPU.

Comparison methods. We compare the performance of C3M with 3 different leading approaches for synthesizing tracking controllers, including both model-based and model-free methods. To be specific, the methods include (1.) **SoS**: the SoS-based method proposed in [30], using the official implementation¹. (2.) **MPC**: an open-source implementation of MPC on PyTorch [7], which solves

¹<https://github.com/StanfordASL/RobustMP>

nonlinear implicit MPC problems using the first-order approximation of the dynamics [40]. (3.) **RL**: the Proximal Policy Optimization (PPO) RL algorithm [34] implemented in Tianshou². Note that although RL is often used in a model-free setting, the advances in deep neural networks have made RL an outstanding tool for controller learning.

Studied systems. We study 4 representative system models adopted from classical benchmarks: (1). **PVTOL** models a planar vertical-takeoff-vertical-landing (PVTOL) system for drones and is adopted from [30, 8]. (2). **Quadrotor** models a physical quadrotor and is adopted from [30]. (3). **Neural lander** models a drone flying close to the ground so that ground effect is prominent and therefore could not be ignored [41]. Note that in this model the ground effect is learned from empirical data of a physical drone using a 4-layer neural network [41]. Due to the neural network function in the dynamics, the SoS-based method as in [30] cannot handle such a model since neural networks are hard to be approximated by polynomials. (4). **SEGWAY** models a real-world Segway robot and is adopted from [13]. Since this model does not satisfy the sparsity assumption for matrix B in Sec. 3.1, SoS-based method cannot handle it. With this model, we also study the impact of the regularization term \mathcal{L}_{w_2} in Equation (7) for learning a tracking controller.

We report the performance of C3M on 5 other examples in Appendix, including a 10D quadrotor model, a cart-pole model, a pendulum model, a two-link planar robot arm system, and the Dubin’s vehicle model in Fig. 1. The detailed dynamics of the above 4 representative system models are also reported in Appendix.

Implementation details. For all systems studied here, we model the dual metric as $W(x) = C(x; \theta_w)^\top C(x; \theta_w) + \underline{w}\mathbf{I}$ where $C(x; \theta_w) \in \mathbb{R}^{n \times n}$ is a 2-layer neural network, of which the hidden layer contains 128 neurons. We use a relative complex structure for the controller. First, two weight matrices of the controller $w_1 = w_1(x, x^*; \theta_{u1})$ and $w_2 = w_2(x, x^*; \theta_{u2})$ are modeled using two 2-layer neural networks with 128 neurons in the hidden layer, where θ_{u1} and θ_{u2} are the parameters. Then the controller is given by $u(x, x^*, u^*; \theta_u) = w_2 \cdot \tanh(w_1 \cdot (x - x^*)) + u^*$, where $\theta_u = \{\theta_{u1}, \theta_{u2}\}$ and $\tanh(\cdot)$ is the hyperbolic tangent function. It is easy to verify the assumptions made in Sec. 3.1 are satisfied: For the controller, $x = x^* \Rightarrow u(x, x^*, u^*; \theta_u) = u^*$ for all θ_u . For the metric, $W(x; \theta_w)$ is a symmetric matrix satisfying $W(x; \theta_w) \succeq \underline{w}\mathbf{I}$ for all x and all θ_w . A training set with 130K samples is used. We train the NN for 20 epochs with the Adam [42] optimizer.

The process of generating random reference trajectories x^* is critical in the evaluation of tracking controllers and also a fundamental part in RL training. We pre-define a group of sinusoidal signals with some fixed frequencies and randomly sampled a weight for each frequency component, then the reference control inputs $u^*(t)$ are calculated as the linear combination of the sinusoidal signals. The initial states of the reference trajectories $x^*(0)$ are uniformly randomly sampled from a compact set. Using $x^*(0)$ and $u^*(t)$, we can get the reference trajectories $x^*(t)$ following the system dynamics (1) with $d(t) = 0$. For evaluation, the initial errors $x_e(0)$ are uniformly randomly sampled from a bounded set, and the initial states $x(0)$ are computed as $x(0) = x^*(0) + x_e(0)$. Trajectories of the closed-loop systems are then simulated on a bounded time horizon $[0, T]$.

To make use of the model-free PPO RL library, we formulate some key concepts as follows. The state of the environment at time t is the concatenation of $x(t)$ and $x^*(t)$. At the beginning of each episode, the environment randomly samples two initial points as $x(0)$ and $x^*(0)$ respectively, and a reference control input $u^*(t)$ using the aforementioned sampling method. At each transition, the environment takes the action $u(t)$ from the agent, and returns the next state as $x(t + \Delta t) = x(t) + \Delta t(f(x(t)) + B(x(t))(u(t) + u^*(t)))$ and $x^*(t + \Delta t) = x^*(t) + \Delta t(f(x^*(t)) + B(x^*(t))u^*(t))$, and the reward as $r_t = 1/(1 + \|w \circ (x(t + \Delta t) - x^*(t + \Delta t))\|_2)$, where w are predefined weights for each state. For a fair comparison, RL and our method share the same architecture for the controller.

For the MPC-PyTorch library, we observe that the time horizon, also called the receding time window for computing the control sequence, plays a crucial role. With smaller time window, the resulting controller cannot always produce trajectories that converge to the reference trajectory. The larger the time window is, the smaller the tracking errors are. However, a large time window can also cause the computational time to increase. As a trade-off, we set time window to be 50 time units in all cases.

Results and discussion Results are reported in Tab. 1 and Fig. 2. For each benchmark, a group of evaluation configurations was randomly sampled using the aforementioned sampling method,

²<https://github.com/thu-ml/tianshou>

Table 1: Comparison results of C3M vs. other methods. n and m are dimensions of the state and input space.

Model	Dim		Execution time per step (ms)				Tracking error (AUC)				Convergence rate (λ/C)			
	n	m	C3M	SoS	MPC	RL	C3M	SoS	MPC	RL	C3M	SoS	MPC	RL
PVTOL	6	2	0.41	3.4	1968	0.41	0.659	0.975	0.892	0.735	1.153/3.95	0.864/4.91	0.489/4.95	0.799/3.84
Quadrotor	9	3	0.40	12.6	3535	0.40	0.772	1.103	0.977	1.416	1.078/3.63	0.821/3.30	0.401/2.23	0.187/1.37
Neural lander	6	3	0.36	-	3385	0.36	0.588	-	0.713	0.793	1.724/2.89	-	0.822/1.34	0.606/1.30
Segway	4	1	0.29	-	-	0.29	0.704	-	-	1.408	0.446/3.11	-	-	0.168/8.12




Figure 2: Normalized tracking error on the benchmarks using different approaches. The y axes are in log-scale. The tubes are tracking errors between mean plus and minus one standard deviation over 100 trajectories.

and then each method was evaluated with these configurations. In Fig. 2 we show the normalized tracking error $x_e(t) = \|x(t) - x^*(t)\|_2 / \|x(0) - x^*(0)\|_2$ for each model and each available method. In Tab. 1, we evaluate the tracking controllers using two quality metrics as discussed in Sec. 3.2. (1) Total tracking error: Given the normalized tracking error curve $x_e(t)$ for $t \in [0, T]$, the total tracking error is just the area under this curve (AUC). (2) Convergence rate: Given a set of tracking error curves, first we do the following for the error curve $x_e(t)$ that possesses the highest overshoot: We search for the convergence rate $\lambda > 0$ and the overshoot $C \geq 1$ such that $x_e(t) \leq Ce^{-\lambda t}$ for all $t \in [0, T]$ and the AUC of $Ce^{-\lambda t}$ is minimized. Then, C is fixed, and the convergence rate λ is calculated for each curve. After obtaining the quality scores for all the error curves, we report the $(1 - \alpha)$ -th quantile with $\alpha = 0.05$ as in Sec. 3.2.

Some observations are in order. (1) On PVTOL and the quadrotor models, both SOS and our method successfully found tracking controllers such that the tracking error decreases rapidly. However, the SOS-based method entails the calculation of geodesics on the fly for control synthesis, while our method learns a neural controller offline, which makes a difference in running time. Also, due to the aforementioned limitations on control matrix sparsity, the SOS-based method cannot be applied to the neural lander and Segway. (2) The performance of the RL method varies with tasks. On PVTOL and the neural lander models, RL achieved comparable results. However, since the state space of the quadrotor model is larger, RL failed to find a contracting tracking controller within a reasonable time. As for the Segway model, although we have tried our best to tune the hyperparameters, RL failed to find a reasonable controller. (3) Running time of MPC is not practical for online control tasks. Also, the MPC library used in experiments failed to handle the Segway model due to numerical issues. (4) The results of our methods with and without penalty term \mathcal{L}_{w2} on Segway verify the sufficiency of this cost term when the sparsity assumption for B is not satisfied.

5 Conclusion

In this paper, we showed that a certified tracking controller with guaranteed convergence and tracking error bounds can indeed be learned by learning a control contraction metric simultaneously as guidance. Results indicated that such a learned tracking controller can better track the reference or nominal trajectory, and the execution time of such a controller is much shorter than online control methods such as MPC. In future works, we plan to extend our learning framework to systems with unknown dynamics and control non-affine systems.

Acknowledgments

The authors acknowledge support from the DARPA Assured Autonomy under contract FA8750-19-C-0089, U.S. Army Research Laboratory Cooperative Research Agreement W911NF-17-2-0196, U.S. National Science Foundation(NSF) grants #1740079 and #1750009. The views, opinions and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

References

- [1] S. M. LaValle. *Planning algorithms*. Cambridge university press, 2006.
- [2] L. E. Kavraki, M. N. Kolountzakis, and J.-C. Latombe. Analysis of probabilistic roadmaps for path planning. *IEEE Transactions on Robotics and Automation*, 14(1):166–171, 1998.
- [3] S. L. Herbert, M. Chen, S. Han, S. Bansal, J. F. Fisac, and C. J. Tomlin. FaSTrack: A modular framework for fast and guaranteed safe motion planning. In *56th Annual Conference on Decision and Control (CDC)*. IEEE, 2017.
- [4] S. Vaskov, S. Kousik, H. Larson, F. Bu, J. Ward, S. Worrall, M. Johnson-Roberson, and R. Vasudevan. Towards provably not-at-fault control of autonomous robots in arbitrary dynamic environments. *arXiv preprint arXiv:1902.02851*, 2019.
- [5] A. Majumdar and R. Tedrake. Funnel libraries for real-time robust feedback motion planning. *The International Journal of Robotics Research*, 36(8):947–982, 2017.
- [6] S. Jha, S. Raj, S. K. Jha, and N. Shankar. Duality-based nested controller synthesis from STL specifications for stochastic linear systems. In *International Conference on Formal Modeling and Analysis of Timed Systems*, pages 235–251. Springer, 2018.
- [7] B. Amos, I. Jimenez, J. Sacks, B. Boots, and J. Z. Kolter. Differentiable mpc for end-to-end planning and control. In *Advances in Neural Information Processing Systems*, 2018.
- [8] S. Singh, S. M. Richards, V. Sindhvani, J.-J. E. Slotine, and M. Pavone. Learning stabilizable nonlinear dynamics with contraction-based regularization. *arXiv:1907.13122*, 2019.
- [9] Y.-C. Chang, N. Roohi, and S. Gao. Neural lyapunov control. In *Advances in Neural Information Processing Systems*, pages 3245–3254, 2019.
- [10] H. Ravanbakhsh and S. Sankaranarayanan. Robust controller synthesis of switched systems using counterexample guided framework. In *2016 international conference on embedded software (EMSOFT)*, pages 1–10. IEEE, 2016.
- [11] A. D. Ames, J. W. Grizzle, and P. Tabuada. Control barrier function based quadratic programs with application to adaptive cruise control. In *53rd IEEE Conference on Decision and Control*, pages 6271–6278. IEEE, 2014.
- [12] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada. Control barrier functions: Theory and applications. In *2019 18th European Control Conference (ECC)*, pages 3420–3431. IEEE, 2019.
- [13] A. Taylor, A. Singletary, Y. Yue, and A. Ames. Learning for safety-critical control with control barrier functions. *arXiv preprint arXiv:1912.10099*, 2019.
- [14] R. Tedrake. LQR-Trees: Feedback motion planning on sparse randomized trees. 2009.
- [15] S. Bansal, M. Chen, J. F. Fisac, and C. J. Tomlin. Safe sequential path planning of multi-vehicle systems under presence of disturbances and imperfect information. In *American Control Conference*, 2017.
- [16] J. Kapinski, J. V. Deshmukh, S. Sankaranarayanan, and N. Arechiga. Simulation-guided lyapunov analysis for hybrid dynamical systems. In *Proceedings of the 17th international conference on Hybrid systems: computation and control*, pages 133–142, 2014.
- [17] S. M. Khansari-Zadeh and A. Billard. Learning control lyapunov function to ensure stability of dynamical system-based robot reaching motions. *Robotics and Autonomous Systems*, 2014.
- [18] S. Jha and P. Lincoln. Data efficient learning of robust control policies. In *2018 56th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 856–861. IEEE, 2018.
- [19] S. Jha, A. Tiwari, S. A. Seshia, T. Sahai, and N. Shankar. Telex: Passive stl learning using only positive examples. In *International Conference on Runtime Verification*, pages 208–224. Springer, 2017.

- [20] S. Jha, V. Raman, D. Sadigh, and S. A. Seshia. Safe autonomy under perception uncertainty using chance-constrained temporal logic. *Journal of Automated Reasoning*, 60(1):43–62, 2018.
- [21] A. J. Taylor, V. D. Dorobantu, H. M. Le, Y. Yue, and A. D. Ames. Episodic learning with control lyapunov functions for uncertain robotic systems. *arXiv:1903.01577*, 2019.
- [22] A. Robey, H. Hu, L. Lindemann, H. Zhang, D. V. Dimarogonas, S. Tu, and N. Matni. Learning control barrier functions from expert demonstrations. *arXiv preprint arXiv:2004.03315*, 2020.
- [23] J. Choi, F. Castañeda, C. J. Tomlin, and K. Sreenath. Reinforcement learning for safety-critical control under model uncertainty, using control lyapunov functions and control barrier functions. *arXiv preprint arXiv:2004.07584*, 2020.
- [24] A. S. Polydoros and L. Nalpantidis. Survey of model-based reinforcement learning: Applications on robotics. *Journal of Intelligent & Robotic Systems*, 86(2):153–173, 2017.
- [25] M. Ohnishi, L. Wang, G. Notomista, and M. Egerstedt. Barrier-certified adaptive reinforcement learning with applications to brushbot navigation. *IEEE Transactions on robotics*, 2019.
- [26] M. Deisenroth and C. E. Rasmussen. Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pages 465–472, 2011.
- [27] K. Chua, R. Calandra, R. McAllister, and S. Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems*, pages 4754–4765, 2018.
- [28] I. R. Manchester and J.-J. E. Slotine. Control contraction metrics: Convex and intrinsic criteria for nonlinear feedback design. *IEEE Transactions on Automatic Control*, 2017.
- [29] I. R. Manchester, J. Z. Tang, and J.-J. E. Slotine. Unifying classical and optimization-based methods for robot tracking control with control contraction metrics. In *International Symposium on Robotics Research (ISRR)*, pages 1–16, 2015.
- [30] S. Singh, B. Landry, A. Majumdar, J.-J. Slotine, and M. Pavone. Robust feedback motion planning via contraction theory. *The International Journal of Robotics Research*, 2019.
- [31] H. Tsukamoto and S.-J. Chung. Neural contraction metrics for robust estimation and control: A convex optimization approach. *arXiv preprint arXiv:2006.04361*, 2020.
- [32] S. M. LaValle. Rapidly-exploring random trees: A new tool for path planning. 1998.
- [33] C. Fan, K. Miller, and S. Mitra. Fast and guaranteed safe controller synthesis for nonlinear vehicle models. In *International Conference on Computer Aided Verification*, 2020.
- [34] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [35] E. F. Camacho and C. B. Alba. *Model predictive control*. Springer, 2013.
- [36] S. J. Anderson, S. C. Peters, T. E. Pilutti, and K. Iagnemma. An optimal-control-based framework for trajectory planning, threat assessment, and semi-autonomous control of passenger vehicles in hazard avoidance scenarios. *International Journal of Vehicle Autonomous Systems*, 8(2-4):190–216, 2010.
- [37] W. Lohmiller and J.-J. E. Slotine. On contraction analysis for non-linear systems. *Automatica*, 34(6):683–696, 1998.
- [38] M. Fazlyab, A. Robey, H. Hassani, M. Morari, and G. J. Pappas. Efficient and accurate estimation of lipschitz constants for deep neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [39] V. Vovk, A. Gammerman, and G. Shafer. *Algorithmic learning in a random world*. Springer Science & Business Media, 2005.

- [40] Y. Tassa, N. Mansard, and E. Todorov. Control-limited differential dynamic programming. In *IEEE International Conference on Robotics and Automation*, pages 1168–1175. IEEE, 2014.
- [41] A. Liu, G. Shi, S.-J. Chung, A. Anandkumar, and Y. Yue. Robust regression for safe exploration in control. *arXiv preprint arXiv:1906.05819*, 2019.
- [42] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [43] C. Fan and S. Mitra. Bounded verification with on-the-fly discrepancy computation. *arXiv preprint arXiv:1502.01801*, 2015.
- [44] J. Lei, M. G’Sell, A. Rinaldo, R. J. Tibshirani, and L. Wasserman. Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 2018.
- [45] H. K. Khalil and J. W. Grizzle. *Nonlinear systems*, volume 3. Prentice hall Upper Saddle River, NJ, 2002.
- [46] Q. Yang and S. Jagannathan. Reinforcement learning controller design for affine nonlinear discrete-time systems using online approximators. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 42(2):377–390, 2011.
- [47] L. Janson, E. Schmerling, A. Clark, and M. Pavone. Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions. *The International journal of robotics research*, 34(7):883–921, 2015.
- [48] C. Richter, A. Bry, and N. Roy. Polynomial trajectory planning for aggressive quadrotor flight in dense indoor environments. In *Robotics Research*, pages 649–666. Springer, 2016.
- [49] S. Karaman, M. R. Walter, A. Perez, E. Frazzoli, and S. Teller. Anytime motion planning using the rrt. In *2011 IEEE International Conference on Robotics and Automation*, pages 1478–1483. IEEE, 2011.
- [50] M. Kobilarov. Cross-entropy motion planning. *The International Journal of Robotics Research*, 31(7):855–871, 2012.
- [51] C. J. Tomlin, I. Mitchell, A. M. Bayen, and M. Oishi. Computational techniques for the verification of hybrid systems. *Proceedings of the IEEE*, 91(7):986–1001, 2003.
- [52] J. H. Gillula, H. Huang, M. P. Vitus, and C. J. Tomlin. Design of guaranteed safe maneuvers using reachable sets: Autonomous quadrotor aerobatics in theory and practice. In *2010 IEEE International Conference on Robotics and Automation*, pages 1649–1654. IEEE, 2010.
- [53] H. Huang, J. Ding, W. Zhang, and C. J. Tomlin. A differential game approach to planning in adversarial scenarios: A case study on capture-the-flag. In *2011 IEEE International Conference on Robotics and Automation*, pages 1451–1456. IEEE, 2011.
- [54] D. Fridovich-Keil, S. L. Herbert, J. F. Fisac, S. Deglurkar, and C. J. Tomlin. Planning, fast and slow: A framework for adaptive real-time safe trajectory planning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 387–394. IEEE, 2018.
- [55] M. Chen, S. L. Herbert, M. S. Vashishtha, S. Bansal, and C. J. Tomlin. Decomposition of reachable sets and tubes for a class of nonlinear systems. *IEEE Transactions on Automatic Control*, 63(11):3675–3688, 2018.
- [56] S. Prajna, A. Jadbabaie, and G. J. Pappas. A framework for worst-case and stochastic safety verification using barrier certificates. *IEEE Transactions on Automatic Control*, 52(8):1415–1428, 2007.
- [57] A. J. Barry, A. Majumdar, and R. Tedrake. Safety verification of reactive controllers for uav flight in cluttered environments using barrier certificates. In *2012 IEEE International Conference on Robotics and Automation*, pages 484–490. IEEE, 2012.
- [58] A. Majumdar and R. Tedrake. Robust online motion planning with regions of finite time invariance. In *Algorithmic foundations of robotics X*, pages 543–558. Springer, 2013.

- [59] W. Langson, I. Chrysoschoos, S. Raković, and D. Q. Mayne. Robust model predictive control using tubes. *Automatica*, 40(1):125–133, 2004.
- [60] D. Q. Mayne, M. M. Seron, and S. Raković. Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41(2):219–224, 2005.
- [61] M. Farina and R. Scattolini. Tube-based robust sampled-data mpc for linear continuous-time systems. *Automatica*, 48(7):1473–1476, 2012.
- [62] M. Kögel and R. Findeisen. Discrete-time robust model predictive control for continuous-time nonlinear systems. In *2015 American Control Conference (ACC)*, pages 924–930. IEEE, 2015.
- [63] S. Yu, C. Maier, H. Chen, and F. Allgöwer. Tube mpc scheme based on robust control invariant set with application to lipschitz nonlinear systems. *Systems & Control Letters*, 62(2):194–200, 2013.
- [64] D. Q. Mayne, E. C. Kerrigan, E. Van Wyk, and P. Falugi. Tube-based robust nonlinear model predictive control. *International Journal of Robust and Nonlinear Control*, 21(11):1341–1353, 2011.
- [65] E. M. Aylward, P. A. Parrilo, and J.-J. E. Slotine. Stability and robustness analysis of nonlinear systems via contraction metrics and sos programming. *Automatica*, 44(8):2163–2170, 2008.
- [66] K. Leung and I. R. Manchester. Nonlinear stabilization via control contraction metrics: A pseudospectral approach for computing geodesics. In *2017 American Control Conference (ACC)*, pages 1284–1289. IEEE, 2017.

Appendix

A Justification for the CCM regularization terms

In theory, minimizing \mathcal{L}_u is enough for learning a tracking controller, since inequality (4) implies the contraction of the closed-loop system. However, in the experiments, we found that minimizing \mathcal{L}_u solely is hard. The training loss stayed on a relatively high level, and the accuracy of inequality (4) was extremely low ($\sim 1\%$). Here, \mathcal{L}_{w1} and \mathcal{L}_{w2} are used as auxiliary loss terms to mitigate the difficulty of minimizing \mathcal{L}_u . After adding these terms, the accuracy approached $\sim 100\%$.

Furthermore, conditions (2) and (3) are sufficient but not necessary for the existence of contracting controller. In the derivation of conditions (2) and (3), the inequality (inequality (4) in the main paper) which indicates contracting is required to be hold for all unbounded u , which is not necessary in the sense that u is usually bounded in a control problem. Thus, even if conditions (2) and (3) do not exactly hold, a contracting tracking controller may still exist.

All the benchmark models used in experiments are under-actuated where the rank of the control matrix $r(B) < n$. In cases where $r(B) = n$ for all x and thus there does not exist a non-zero annihilator matrix B_\perp such that $B_\perp^\top B = 0$, one can easily verify that any uniformly positive definite metric function $M(x)$ can be a valid CCM. The control synthesis is also straightforward in this case. One can directly compute the instantaneous control u given the desired \dot{x} since $B(x)$ is invertible. Moreover, as shown in [28], existence of a CCM is guaranteed for feedback linearizable systems. In that case, a valid CCM can be constructed analytically from the variable and feedback transformation. We refer the interested readers to [28] for more details.

B Proof of Proposition 2 and CCM validation process example

B.1 Proof of Proposition 2

We denote the operator 2-norm for matrices by $\|\cdot\|_2$. The following lemmas are going to be used for the proof of Proposition 2.

Lemma 5. *For any two symmetric matrices $A, B \in \mathbb{R}^{n \times n}$, the difference of their eigenvalues satisfies:*

$$|\lambda_{\max}(A) - \lambda_{\max}(B)| \leq \|A - B\|_2.$$

Lemma 5 is a well-known result that uses the Courant-Fischer minimax theorem. The detailed proof can be found at [43].

Lemma 6. *For any two functions $A : \mathbb{R}^n \rightarrow \mathbb{R}^{k \times \ell}$ and $B : \mathbb{R}^n \rightarrow \mathbb{R}^{\ell \times s}$, if A and B both have Lipschitz constants L_A and L_B with respect to 2-norm respectively, moreover, if $\forall x, \|A(x)\|_2 \leq S_A$ and $\|B(x)\|_2 \leq S_B$, then for any $x, y \in \mathbb{R}^n$,*

$$\|A(x)B(x) - A(y)B(y)\|_2 \leq (S_A L_B + S_B L_A) \|x - y\|_2.$$

Proof.

$$\begin{aligned} \|A(x)B(x) - A(y)B(y)\|_2 &= \|A(x)B(x) - A(x)B(y) + A(x)B(y) - A(y)B(y)\|_2 \\ &\leq L_B \|A(x)\|_2 \|x - y\|_2 + L_A \|B(y)\|_2 \|x - y\|_2 \\ &\leq (S_A L_B + S_B L_A) \|x - y\|_2. \end{aligned} \tag{9}$$

□

The proof of Proposition 2 follows from the above lemmas.

Proof. For all X and $Y \in \mathcal{X} \times \mathcal{X} \times \mathcal{U}$, we have

$$\begin{aligned} & \left| \lambda_{\max} \left(\dot{M}(X) + (A(X) + B(X)K(X))^{\top} M(X) + M(X)(A(X) + B(X)K(X)) + 2\lambda M(X) \right) \right. \\ & \left. - \lambda_{\max} \left(\dot{M}(Y) + (A(Y) + B(Y)K(Y))^{\top} M(Y) + M(Y)(A(Y) + B(Y)K(Y)) + 2\lambda M(Y) \right) \right| \\ & \leq \|\dot{M}(X) - \dot{M}(Y)\|_2 + 2\|M(X)A(X) - M(Y)A(Y)\|_2 \\ & \quad + 2\|M(X)B(X)K(X) - M(Y)B(Y)K(Y)\|_2 + 2\lambda\|M(X) - M(Y)\|_2 \\ & \leq (L_{\dot{M}} + 2(S_M L_A + S_A L_M + S_M S_B L_K + S_B S_K L_M + S_M S_K L_B + \lambda L_M)) \|X - Y\|_2. \end{aligned}$$

□

B.2 CCM validation on Dubin’s vehicle

First, the Lipschitz constants and bounds on 2-norms are derived as follows.

L_A and S_A : By definition, $A = \frac{\partial f}{\partial x} + \sum \frac{\partial b_i}{\partial x} u^i = \begin{bmatrix} 0 & 0 & -v \cos(\theta) & \cos(\theta) \\ 0 & 0 & v \sin(\theta) & \sin(\theta) \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$. Considering a

vector-valued function $\tilde{A} = [-v \cos(\theta), v \sin(\theta), \cos(\theta), \sin(\theta)]$ and its Lipschitz constant $L_{\tilde{A}}$ w.r.t.

the Frobenius norm, since the Jacobian of \tilde{A} is given by $\frac{\partial \tilde{A}}{\partial x} = \begin{bmatrix} -\cos(\theta) & v \sin(\theta) \\ \sin(\theta) & v \cos(\theta) \\ 0 & -\sin(\theta) \\ 0 & \cos(\theta) \end{bmatrix}$, we have

$L_{\tilde{A}} = \sup \|\frac{\partial \tilde{A}}{\partial x}\|_F = \sqrt{2 + v_h^2}$, where v_h is the upper bound of the velocity. Obviously, under the Frobenius norm, the Lipschitz constants of A and \tilde{A} coincide. Moreover, for any matrix A , we have $\|A\|_2 \leq \|A\|_F$, and thus $L_A \leq L_{\tilde{A}} = \sqrt{2 + v_h^2}$. As for the norm bound, $S_A = \sup \|A\|_2 \leq \sup \|A\|_F = \sqrt{1 + v_h^2}$.

L_B and S_B : By definition, $B = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}$, which is a constant. Thus, $L_B = 0$, and $S_B =$

$\sigma_{max}(B) = 1$, where σ_{max} is the largest singular value.

L_M and S_M : By definition, we have $W(x) = C(x)^{\top} C(x) + \underline{w} \mathbf{I}$ and $M(x) = W(x)^{-1}$, and thus $S_M = \sup \|M(x)\|_2 \leq \frac{1}{\underline{w}}$. Moreover, $\|M(x) - M(y)\|_2 = \|M(x)(W(x) - W(y))M(y)\|_2 \leq \frac{1}{\underline{w}^2} \|W(x) - W(y)\|_2$. Thus, we have $L_M \leq \frac{1}{\underline{w}^2} L_W$. As for L_W , by Lemma 6, $L_W \leq 2S_C L_C$, where L_C can be evaluated using tools for estimating the Lipschitz constants of NNs. In our experiments, we use the official implementation³ of [38]. After obtaining L_C , S_C can be upper bounded using samples and the Lipschitz constant L_C .

$L_{\dot{M}}$: By definition, $\dot{M} = \sum_{i=1}^n \frac{\partial M}{\partial x_i} \dot{x}_i$. We randomly sampled 10K points and evaluate the gradient of \dot{M} , then the maximum of the norm of the gradient is returned as a Lipschitz constant.

L_K and S_K : For simplicity of the analysis, we use a simpler neural controller $u = k(x, x^*)x_e + u^*$, where k is a two-layer neural network and $x_e := (x - x^*)$. Then, $K = \frac{\partial u}{\partial x} = k(x, x^*) + \sum_{i=1}^n x_e^i \frac{\partial k_i}{\partial x}$, where k_i is the i -th column of k and x_e^i is the i -th element of x_e . Thus, $S_K \leq S_k + \sum_{i=1}^n \sup |x_e^i| L_{k_i|x}$, where $L_{k_i|x}$ is the Lipschitz constant of k_i w.r.t. x . Also, $L_K \leq L_k + \sum_{i=1}^n \sup |x_e^i| L_{\frac{\partial k_i}{\partial x}} + \sqrt{2} \sum_{i=1}^n L_{k_i|x}$, where we use the facts that $x_e^i = x_i - x_i^* \implies L_{x_e^i} = \sqrt{2}$, and $S_{\frac{\partial k_i}{\partial x}} \leq L_{k_i|x}$. As for $L_{\frac{\partial k_i}{\partial x}}$, we have $\frac{\partial k_i}{\partial x} = A_2^i \cdot \text{diag}(\text{dtanh}(A_1[x, x^*])) \cdot A_{11}$, where A_1 is the weights in the first layer, A_{11} is the weights associated to x in the first layer, A_2^i is the weights

³<https://github.com/arobey1/LipSDP>

associated to the i -th column of k in the second layer, and $d\tanh$ is the derivative of the \tanh function. Thus, $L_{\frac{\partial k_i}{\partial x}} \leq \|A_2^i\| \cdot \|A_{11}\| \cdot \|A_1\| \cdot L_{d\tanh}$.

Validation process: We successfully verified the trained model for a state space where $[-5, -5, -\pi, 1] \leq x$, $x^* \leq [5, 5, \pi, 2]$, $[-0.1, -0.1, -0.1, -0.1] \leq x - x^* \leq [0.1, 0.1, 0.1, 0.1]$, and $[-1, 0] \leq u^* \leq [1, 0]$ with $\lambda = 0.1$. The derived Lipschitz constant using Proposition 2 for the trained model was 616.8. After that, we discretized the state space and evaluated the maximum eigenvalues of the LHS of inequality (4) at the grid points, and the results gave a theoretical guarantee on the satisfaction of inequality (4) for all x , x^* , and u^* .

C Proof of Proposition 3

The following lemma from [44] is going to be used.

Lemma 7. Consider a set of i.i.d. random variables $\{U_1, \dots, U_n\}$, then the following inequality holds for a new i.i.d. sample U_{n+1} ,

$$\Pr(U_{n+1} \geq q_{1-\alpha}) \leq 1 - \alpha,$$

where $q_{1-\alpha}$ is $(1 - \alpha)$ -th quantile of $\{U_i\}_{i=1}^n$, i.e.

$$q_{1-\alpha} = \begin{cases} U_{(\lceil (1-\alpha)(n+1) \rceil)} & \text{if } \lceil (1-\alpha)(n+1) \rceil \leq n, \\ \infty & \text{otherwise} \end{cases},$$

where $U_{(1)} \leq \dots \leq U_{(n)}$ is the order statistics of $\{U_i\}_{i=1}^n$.

The above lemma follows from an obvious fact that the rank of U_{n+1} in $\{U_i\}_{i=1}^{n+1}$ is uniformly distributed in $\{1, \dots, n+1\}$.

Since the evaluation configurations are all i.i.d., the corresponding quality scores $\{m_i\}_{i=1}^{n+1}$ are also i.i.d. The proposition immediately follows from the above lemma.

D Proof of Theorem 4

Proof. Since $\forall x, M(x) \in S_n^{>0}$, there exists $\Theta(x)$ such that $M(x) = \Theta(x)^\top \Theta(x)$ for all x . Now, let us consider the smallest path integral between $x_1(t)$ and $x_2(t)$ w.r.t. metric $M(x)$ and denote it by $R(t) = \int_{x_1(t)}^{x_2(t)} \sqrt{\delta_x^\top M(x) \delta_x} = \int_{x_1(t)}^{x_2(t)} \|\Theta(x) \delta_x\|_2$. Then differentiating $R(t)$ w.r.t. t yields $\dot{R} \leq -\lambda R + \|\Theta d\|_2 \leq -\lambda R + \epsilon \sqrt{m}$. (see [37], pp. 11 (vii)) By comparison lemma (see [45], pp. 102), $R(t)$ is bounded as $R(t) \leq R(0)e^{-\lambda t} + \frac{\epsilon}{\lambda} \sqrt{m}(1 - e^{-\lambda t})$. Since Θ is uniformly bounded as $\Theta(x) \succeq \sqrt{m}\mathbf{I}$, we have $R(t) \geq \sqrt{m}\|x_1(t) - x_2(t)\|_2$, which follows from the fact that the smallest path integral w.r.t. a constant metric is the path integral along the straight line. Combining these two inequalities yields the theorem. \square

E More details on the experimental results

E.1 Hyper-parameters

For all the models, we set $\underline{w} = 0.1$ and $\bar{w} = 10$. For PVTOL, Quadrotor, Neural lander, Segway, Quadrotor2, and Cart-pole, we set $\lambda = 0.5$. For Dubin's car, we set $\lambda = 1$. For TLPRA, we set $\lambda = 2$. For Pendulum, we set $\lambda = 3$.

For the evaluation of RL, we trained the model for 1K epochs (1K steps per epoch, and thus 1M steps in total). The best model was used for evaluation.

E.2 Dynamics of the 4 benchmarks discussed in Section 4

As mentioned in the main paper, for each benchmark model, we have to define the state space \mathcal{X} and the reference control space \mathcal{U} . Also, for evaluation purpose, the initial state of the reference trajectory is sampled from a set \mathcal{X}_0 , and the initial error is sampled from a set \mathcal{X}_{e0} . All these sets are defined as hyper-rectangles. For two vectors $x, y \in \mathbb{R}^n$, $x \leq y$ or $x \geq y$ means element-wise inequality.

PVTOL models a planar vertical-takeoff-vertical-landing (PVTOL) system for drones and is adopted from [30, 8]. The system includes 6 variables $x := [p_x, p_z, \phi, v_x, v_z, \dot{\phi}]$, where p_x and p_z are the positions, v_x and v_z are the corresponding velocities, ϕ is the roll, and $\dot{\phi}$ is the angular rate. The 2-dimensional control input u corresponds to the thrusts of the two rotors. The dynamics of the system are as follows:

$$\dot{x} = \begin{bmatrix} v_x \cos(\phi) - v_z \sin(\phi) \\ v_x \sin(\phi) + v_z \cos(\phi) \\ \dot{\phi} \\ v_z \dot{\phi} - g \sin(\phi) \\ -v_x \dot{\phi} - g \cos(\phi) \\ 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 1/m & 1/m \\ l/J & -l/J \end{bmatrix} u,$$

where $m = 0.486$, $J = 0.00383$, $g = 9.81$, and $l = 0.25$. For the sets, we set $\mathcal{X} = \{[-35, -2, -\frac{\pi}{3}, -2, -1, -\frac{\pi}{3}] \leq x \leq [0, 2, \frac{\pi}{3}, 2, 1, \frac{\pi}{3}]\}$, $\mathcal{U} = \{[mg/2 - 1, mg/2 - 1] \leq u \leq [mg/2 + 1, mg/2 + 1]\}$, $\mathcal{X}_0 = \{[0, 0, -0.1, 0.5, 0, 0] \leq x_0 \leq [0, 0, 0.1, 1, 0, 0]\}$, and $\mathcal{X}_{e0} = \{[-0.5, -0.5, -0.5, -0.5, -0.5, -0.5] \leq x_{e0} \leq [0.5, 0.5, 0.5, 0.5, 0.5, 0.5]\}$.

Quadrotor is adopted from [30]. The state variables are given by $x := [p_x, p_y, p_z, v_x, v_y, v_z, f, \phi, \theta, \psi]$, where p_k and v_k for $k \in \{x, y, z\}$ are positions and velocities, and ϕ , θ and ψ are roll, pitch, and yaw respectively, and $f > 0$ is the net (normalized by mass) thrust generated by the four rotors. The control input is considered as $u := [f, \phi, \theta, \psi]$. The dynamics of the system are as follows:

$$\dot{x} = \begin{bmatrix} v_x \\ v_y \\ v_z \\ -f \sin(\theta) \\ f \cos(\theta) \sin(\phi) \\ g - f \cos(\theta) \cos(\phi) \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} u,$$

where $g = 9.81$. For the sets, we set $\mathcal{X} = \{[-30, -30, -30, -1.5, -1.5, -1.5, 0.5g, -\frac{\pi}{3}, -\frac{\pi}{3}, -\frac{\pi}{3}] \leq x \leq [30, 30, 30, 1.5, 1.5, 1.5, 2g, \frac{\pi}{3}, \frac{\pi}{3}, \frac{\pi}{3}]\}$, $\mathcal{U} = \{[-1, -1, -1, -1] \leq u \leq [1, 1, 1, 1]\}$, $\mathcal{X}_0 = \{[-5, -5, -5, -1, -1, -1, g, 0, 0, 0] \leq x_0 \leq [5, 5, 5, 1, 1, 1, g, 0, 0, 0]\}$, and $\mathcal{X}_{e0} = \{[-0.5, -0.5, -0.5, -0.5, -0.5, -0.5, -0.5, -0.5, -0.5, -0.5] \leq x_{e0} \leq [0.5, 0.5, 0.5, 0.5, 0.5, 0.5, 0.5, 0.5, 0.5, 0.5]\}$.

Neural lander models a drone flying close to the ground so that ground effect is prominent and therefore could not be ignored [41]. The state is give by $x := [p_x, p_y, p_z, v_x, v_y, v_z]$, where p_k and v_k for $k \in \{x, y, z\}$ are the positions and velocities. The 3-dimensional control input is given by $u := [a_x, a_y, a_z]$ which are the acceleration in three directions. Note that in this model the ground effect is learned from empirical data of a physical drone using a 4-layer neural network [41]. Due to the neural network function in the dynamics, the SoS-based method as in [30] cannot handle such a model since neural networks are hard to be approximated by polynomials. However, C3M can handle this model since it doesn't impose hard constraints and only requires the ability to evaluate the dynamics and its gradient. The dynamics of the system are as follows:

$$\dot{x} = \begin{bmatrix} v_x \\ v_y \\ v_z \\ Fa_1/m \\ Fa_2/m \\ Fa_3/m - g \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} u,$$

where $m = 1.47$, $g = 9.81$, and $Fa_i = Fa_i(z, v_x, v_y, v_z)$ for $i = 1, 2, 3$, are neural networks. For the sets, we set $\mathcal{X} = \{[-5, -5, 0, -1, -1, -1] \leq x \leq [5, 5, 2, 1, 1, 1]\}$, $\mathcal{U} = \{[-1, -1, -3] \leq u \leq [1, 1, 9]\}$, $\mathcal{X}_0 = \{[-3, -3, 0.5, 1, 0, 0] \leq x_0 \leq [3, 3, 1, 1, 0, 0]\}$, and $\mathcal{X}_{e0} = \{[-1, -1, -0.4, -1, -1, 0] \leq x_{e0} \leq [1, 1, 1, 1, 1, 0]\}$.

SEGWAY models a real-world Segway robot and is adopted from [13]. The state variables are given by $x := [p_x, \theta, v_x, \omega]$ where p_x and v_x are the position and velocity, and θ and ω are the angle and angular velocity. The control input is the torque of the motor. The dynamics of the system are as follows:

$$\dot{x} = \begin{bmatrix} v_x \\ \omega \\ \frac{\cos(\theta)(9.8 \sin(\theta) - 1.8u + 11.5v) - 10.9u + 68.4v - 1.2\omega^2 \sin(\theta)}{\cos(\theta) - 24.7} \\ \frac{(9.3u - 58.8v) \cos(\theta) + 38.6u - 243.5v - \sin(\theta)(208.3 + \omega^2 \cos(\theta))}{\cos^2(\theta) - 24.7} \end{bmatrix}.$$

For the sets, we set $\mathcal{X} = \{[-5, -\frac{\pi}{3}, -1, -\pi] \leq x \leq [5, \frac{\pi}{3}, 1, \pi]\}$, $\mathcal{U} = \{0 \leq u \leq 0\}$, $\mathcal{X}_0 = \{[0, 0, 0, 0] \leq x_0 \leq [0, 0, 0, 0]\}$, and $\mathcal{X}_{e0} = \{[-1, -\frac{\pi}{3}, -0.5, -\pi] \leq x_{e0} \leq [1, \frac{\pi}{3}, 0.5, \pi]\}$.

E.3 Additional benchmarks

Cart-pole is a well-known model for evaluating control algorithms. The state variables are given by $x := [p_x, \theta, v_x, \omega]$. The dynamics are given by

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

where $m_c = 1$, $m_p = 1$, $g = 9.8$, and $l = 1$. For the sets, we set $\mathcal{X} = \{[-5, -\frac{\pi}{3}, -1, -1] \leq x \leq [5, \frac{\pi}{3}, 1, 1]\}$, $\mathcal{U} = \{0 \leq u \leq 0\}$, $\mathcal{X}_0 = \{[0, 0, 0, 0] \leq x_0 \leq [0, 0, 0, 0]\}$, and $\mathcal{X}_{e0} = \{[-0.3, -0.3, -0.3, -0.3] \leq x_{e0} \leq [0.3, 0.3, 0.3, 0.3]\}$.

Quadrotor2 is adopted from [3]. The state variables are given by $x := [p_x, p_y, p_z, v_x, v_y, \theta_x, \theta_y, \omega_x, \omega_y, v_z]$. The control input is given by $u := [a_x, a_y, a_z]$. The dynamics are given by

$$\dot{x} = \begin{bmatrix} v_x \\ v_y \\ v_z \\ g \tan(\theta_x) \\ g \tan(\theta_y) \\ -d_1 \theta_x + \omega_x \\ -d_1 \theta_y + \omega_y \\ -d_0 \theta_x + n_0 a_x \\ -d_0 \theta_y + n_0 a_y \\ k_T a_z - g \end{bmatrix},$$

where $d_0 = 10$, $d_1 = 8$, $n_0 = 10$, $k_T = 0.91$, and $g = 9.81$. For the sets, we set $\mathcal{X} = \{[-15, -15, -15, -2, -2, -\frac{\pi}{3}, -\frac{\pi}{3}, -\frac{\pi}{3}, -\frac{\pi}{3}, -2] \leq x \leq [15, 15, 15, 2, 2, \frac{\pi}{3}, \frac{\pi}{3}, \frac{\pi}{3}, \frac{\pi}{3}, 2]\}$, $\mathcal{U} = \{[-10, -10, 0] \leq u \leq [10, 10, 1.5g]\}$, $\mathcal{X}_0 = \{[-2, -2, -2, -1, -1, -0.5, -0.5, -0.5, -0.5, -1] \leq x_0 \leq [2, 2, 2, 1, 1, 0.5, 0.5, 0.5, 0.5, 1]\}$, and $\mathcal{X}_{e0} = \{[-0.5, -0.5, -0.5, -0.5, -0.5, -0.5, -0.5, -0.5, -0.5, -0.5] \leq x_{e0} \leq [0.5, 0.5, 0.5, 0.5, 0.5, 0.5, 0.5, 0.5, 0.5, 0.5]\}$.

Pendulum is a simple model for evaluating control methods. The state variables are given by $x := [\theta, \dot{\theta}]$, where θ is the angle, and $\dot{\theta}$ is the angular velocity. The one-dimensional control input corresponds to the torque of the motor. The dynamics are as follows:

$$\dot{x} = \begin{bmatrix} \dot{\theta} \\ \frac{mgl \sin(\theta) - 0.1\dot{\theta} + u}{ml^2} \end{bmatrix},$$

where $g = 9.81$, $m = 0.15$, and $l = 0.5$. As for the sets, we set $\mathcal{X} = \{[-\frac{\pi}{3}, -\frac{\pi}{3}] \leq x \leq [\frac{\pi}{3}, \frac{\pi}{3}]\}$, $\mathcal{U} = \{-1 \leq u \leq 1\}$, $\mathcal{X}_0 = \{[0, 0] \leq x_0 \leq [0, 0]\}$, and $\mathcal{X}_{e0} = \{[-\frac{\pi}{4}, -\frac{\pi}{4}] \leq x_{e0} \leq [\frac{\pi}{4}, \frac{\pi}{4}]\}$.

Dubin’s car. The state of the car is given by $x := [p_x, p_y, \theta, v]$, where p_x, p_y are the positions, θ is the heading direction, and v is the velocity. Its motion is controlled by two inputs: angular velocity ω and linear acceleration a . The car’s dynamics are given by:

$$\dot{x} = \begin{bmatrix} v \cos(\theta) \\ v \sin(\theta) \\ \omega \\ a \end{bmatrix}.$$

For the sets, we set $\mathcal{X} = \{[-5, -5, -\pi, 1] \leq x \leq [5, 5, \pi, 2]\}$, $\mathcal{U} = \{[-1, 0] \leq u \leq [1, 0]\}$, $\mathcal{X}_0 = \{[-2, -2, -1, 1.5] \leq x_0 \leq [2, 2, 1, 1.5]\}$, and $\mathcal{X}_{e0} = \{[-1, -1, -1, -1] \leq x_{e0} \leq [1, 1, 1, 1]\}$. There is a simpler Dubin’s car model with three state variables $x := [p_x, p_y, \theta]$ and a two-dimensional control input $u := [v, \omega]$. The dynamics are similar to the above ODEs without $\dot{v} = a$. However, we notice that in the simpler model $f(x) \equiv 0$ and $B(x)$ is not full-rank. Therefore, Eq. (2) in the paper is not feasible and the loss term \mathcal{L}_{w1} does not make sense. If we train the model only with the loss \mathcal{L}_u as in Eq. (5), the resulting controller is not performing well.

TLPRA models a Two-Link Planar Robot Arm system and is adopted from [46]. The state variables are given by $x := [q_1, q_2, \dot{q}_1, \dot{q}_2]$, where q and \dot{q} are the rotational angle and angular velocity of the joints respectively. The control input is given by $u := [\tau_1, \tau_2]$ which are the torques applied on the joints. The dynamics are given by

$$\begin{bmatrix} \ddot{q}_1 \\ \ddot{q}_2 \end{bmatrix} = \begin{bmatrix} \alpha + \beta + 2\eta \cos(q_2) & \beta + \eta \cos(q_2) \\ \beta + \eta \cos(q_2) & \beta \end{bmatrix}^{-1} \begin{bmatrix} \tau_1 + \eta(2\dot{q}_1\dot{q}_2 + \dot{q}_2^2) \sin(q_2) - \alpha e_1 \cos(q_1) - \eta e_1 \cos(q_1 + q_2) \\ \tau_2 - \eta\dot{q}_1^2 \sin(q_2) - \eta e_1 \cos(q_1 + q_2) \end{bmatrix},$$

where $\alpha = (m_1 + m_2)a_1^2$, $\beta = m_2a_2^2$, $\eta = m_2a_1a_2$, $e_1 = g/a_1$, $g = 9.8$, $m_1 = 0.8$, $m_2 = 2.3$, $a_1 = 1$, and $a_2 = 1$. For the sets, we set $\mathcal{X} = \{[-\frac{\pi}{2}, -\frac{\pi}{2}, -\frac{\pi}{3}, -\frac{\pi}{3}] \leq x \leq [\frac{\pi}{2}, \frac{\pi}{2}, \frac{\pi}{3}, \frac{\pi}{3}]\}$, $\mathcal{U} = \{[0, 0] \leq u \leq [0, 0]\}$, $\mathcal{X}_0 = \{[\frac{\pi}{2}, 0, 0, 0] \leq x_0 \leq [\frac{\pi}{2}, 0, 0, 0]\}$, and $\mathcal{X}_{e0} = \{[-0.3, -0.3, 0, 0] \leq x_{e0} \leq [0.3, 0.3, 0, 0]\}$.

E.4 Experimental results on additional benchmarks

Results comparing the performance of C3M and other approaches on the additional 5 models are shown in Fig. 3 and Tab. 2. On some simple models including Pendulum, Dubin’s car, Cart-pole, and TLPRA, both RL and C3M achieve perfect results, and C3M demonstrates lower variance. On the higher-dimensional system, Quadrotor2, C3M outperforms RL with lower tracking error and lower variance. Among these models, the SoS-based method cannot handle Cart-pole since it does not satisfy the sparse assumption for the control matrix B and TLPRA since the dynamics is too complicated and hard to be approximated with low-degree polynomials. We did not evaluate the SoS-based method on other models since the case-by-case hyper-parameters tuning is too time consuming without domain expertise on the models. For Cart-pole, Pendulum, and TLPRA, the MPC-PyTorch library had numerical errors.

Model	Dim		Execution time per step (ms)			Tracking error (AUC)			Convergence rate (λ/C)		
	n	m	C3M	MPC	RL	C3M	MPC	RL	C3M	MPC	RL
CAR	4	2	0.46	2300	0.46	0.290	0.509	0.275	1.970/2.69	0.493/1.49	1.836/4.99
Quadrotor2	10	3	0.45	7460	0.45	1.239	1.673	2.081	0.400/3.52	0.451/3.98	0.237/3.78
Cart-pole	4	1	0.45	-	0.45	0.865	-	1.207	1.100/4.85	-	0.509/4.10
Pendulum	2	1	0.37	-	0.37	0.091	-	0.092	5.772/6.82	-	2.675/9.46
TLPRA	4	2	0.42	-	0.42	1.082	-	1.086	2.035/4.91	-	1.883/6.82

Table 2: Comparison results of C3M vs. other methods. Here, n and m are dimensions of the state and input space. The MPC library used in experiments failed to handle some models due to numerical issues (we got some **nan** errors).

E.5 Sampled trajectories under disturbance

We show some plots of actual trajectories of the controlled systems in Fig. 4. We simulate the controlled Dubin’s car and Quadrotor2 system using different control strategy and with different

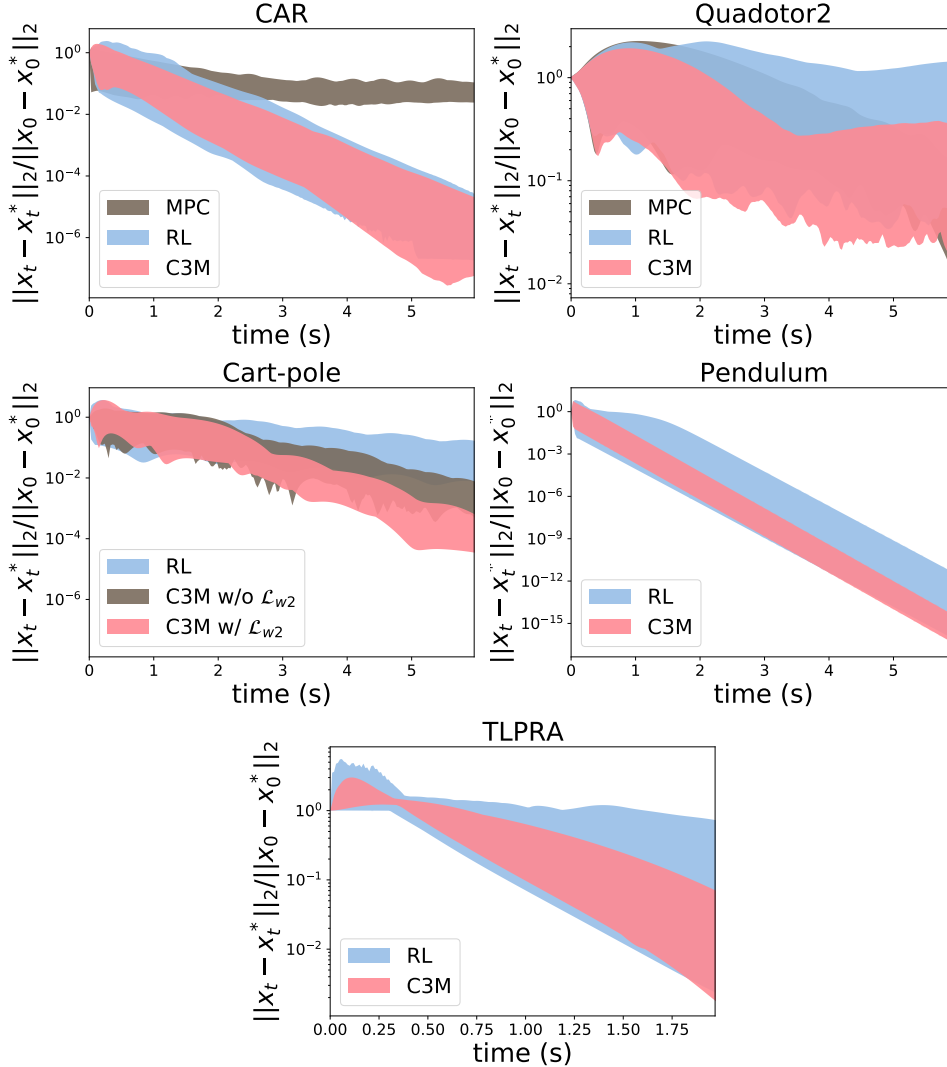


Figure 3: Normalized tracking error on the benchmarks using different approaches. The y axes are in log-scale. The tubes are tracking errors between mean plus and minus one standard deviation over 100 trajectories.

levels of disturbances. In order to model a more realistic disturbance (e.g. wind), we do not adopt white noise as disturbance. Instead, $d(t)$ is a piece-wise constant function, where the length of each constant interval and the magnitude on each interval are uniformly randomly sampled from $[0, 1]$ (seconds) and $[0, \sigma]$ respectively. Since the controlled system is proved to be contracting using C3M, the tracking error is still bounded even under additive random disturbances. We can also see from Fig. 4 that even with disturbance, the trajectories of the closed-loop system using C3M has less tracking error than the system with controllers using RL or MPC.

F More discussions on related work

In the broad area of safe motion planning, driving the robots to some waypoints or follow some nominal trajectory without hitting obstacles in dynamically changing environments is a very challenging task, especially when robotics systems are nonlinear and nonholonomic. In this section, we provide more details on some related techniques in this area.

Sample-based planning methods including rapidly-exploring random trees (RRT) [32], probabilistic road maps (PRM) [2], fast marching tree (FMT) [47], and many others [48, 49, 50], plan safe open-loop trajectories by exploring the environment using samples. However, it is difficult to use sample-

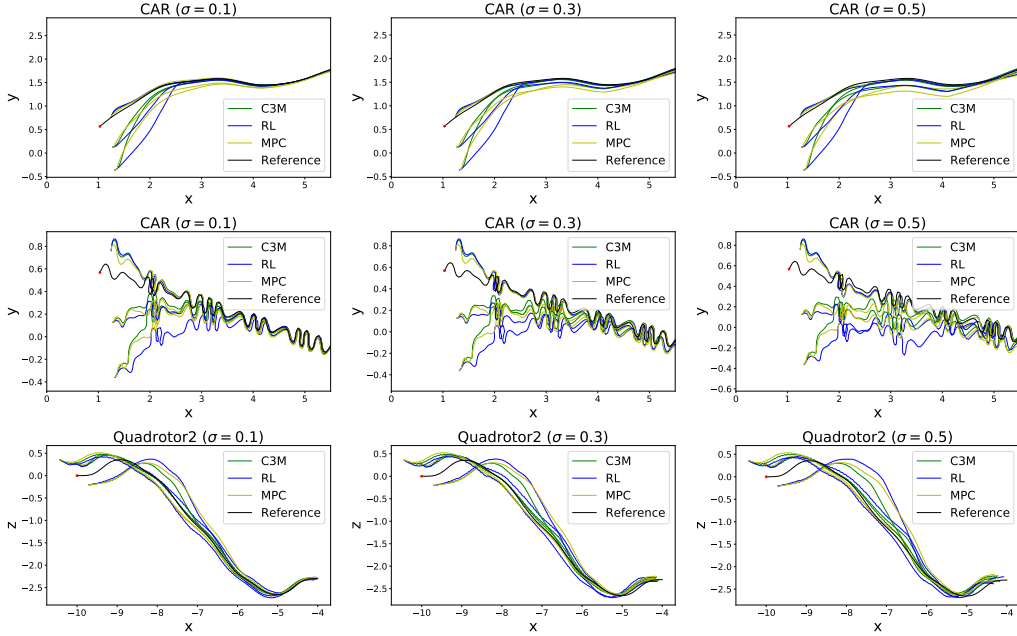


Figure 4: Actual trajectories of the controlled systems in the presence of disturbances. Red dots indicate the starting points of the reference trajectories. Moreover, in order to show that the proposed method can work for any reference trajectories, in the second row, we use a reference control u^* that consists of a very high-frequency component.

based planning to handle uncertainty and disturbances with safety guarantees. In order to plan safe trajectories in the presence of disturbances, the uncertainty should be handled appropriately. The disturbances are usually assumed to be bounded, and the safety can be guaranteed by dealing with the worst case safely.

One may model safe motion planning problems as differential games, where the disturbances are modeled as an adversarial agent. Backward reachable sets and the corresponding optimal controller can be computed by solving Hamilton-Jacobi PDE using level-set methods [51, 52, 53, 3, 54, 55]. However, this class of methods suffers from poor scalability. Solving Hamilton-Jacobi PDE for high-dimensional models is computationally expensive. Moreover, such high computational complexity limits their capability of solving online planning problems for dynamical environments. In comparison, our learned tracking controller can be executed in sub-millisecond level and therefore supports online planning. Moreover, because of the learned contraction metric, we can pre-compute the tracking error bound, which can be further used to understand how far the nominal trajectories should be from the obstacles.

Barrier functions are another widely-used class of certificates for safety. A barrier function is a function of state whose time-derivative on the zero-level set is uniformly negative, which makes the zero-level set an invariant. In [56, 57], barrier functions are computed for given controllers and any admissible disturbances and are used for verification of the closed-loop system. As an extension to controlled systems, control barrier functions [11, 12, 13] enable synthesis of safety-guaranteed controllers and thus have gained increasing popularity. In [11], the authors jointly leveraged control barrier functions (CBF) and control Lyapunov functions (CLF) as safety and objective certificates respectively for control synthesis in the context of adaptive cruise control and lane keeping. However, finding valid CBF for general models is hard. Recently, the advances of machine learning have been utilized to help search for CBF [13].

The idea of synthesizing a controller and bounding the tracking error meanwhile through pre-computation has gained increasing popularity. It enables safe motion planning for dynamically changing environments. Control Lyapunov functions can be used for this purpose [10]. Learning-based CLF has also been studied in some recent works [23, 17, 22]. However, in order to use CLF for synthesizing controller to track reference or nominal trajectories instead of driving the system to a fixed goal set, the original dynamics has to be converted into error dynamics depending on the

reference trajectory, which limits the diversity of available reference trajectories. Similarly, funnel library methods have been studied in [58, 5], where tracking controllers for a fixed set of reference trajectories and the corresponding tracking error are computed offline. In the online planning phase, the safety of the reference trajectories are examined using the funnels corresponding to them, and the safest one is chosen. Again, this class of methods lack diversity of available reference trajectories which is critical to safe motion planning in complicated environments. LQR-Trees [14] builds a tree of LQR controllers backward from the goal set such that the total contraction regions of the nodes cover the whole state space, however, it cannot handle scenarios where the environments are unknown until runtime. Tube Model Predictive Control (TMPC) is another class of related techniques, where a tracking controller is computed such that the actual trajectory remains in a tube centered at the planned MPC nominal trajectory in the presence of bounded disturbances. TMPC techniques have been studied for linear system in many works [59, 60, 61]. As for non-linear systems, linearization and Lipschitz-continuity-based reachability analysis are used [62, 63, 64]. However, in the latter case, the tube are often too conservative to be used in motion planning with limited freespace.

Contraction analysis [37] is another series of methods for analyzing the incremental stability of systems. Recently, it has been extended to controlled system in [28] and thus enabled tracking control synthesis for arbitrary reference trajectories with guaranteed bound on tracking error. The most challenging part in contraction analysis is the search for a valid contraction metric which entails solving Linear Matrix Inequalities (LMIs). In [65], the authors proposed to solve this feasibility problem with Sum-of-Squares (SoS) programming. In [30], the authors extended the SoS-based method to controlled systems, i.e. search for a *control* contraction metric (CCM) instead of an ordinary one, and proposed a more general method for synthesizing control given a valid CCM. However, in order to apply SoS-based methods, the dynamics of the system has to be represented by polynomials or can be approximated by polynomials. Furthermore, the method proposed in [30] relies on an assumption on the structure which encodes the controllability of the system. Different from ours, the method proposed in [30] searches for a metric using the structural assumption and synthesize the controller using geodesics, which entails heavy computation. Moreover, geodesics cannot be computed exactly, and thus the authors in [30] resorted to the optimization-based method proposed in [66] as an approximation of the geodesics.

Combining machine learning and contraction theory has also been studied in some recent work. In [8], the authors proposed a learning-based method for identifying an unknown system with an additional constraint that such learned dynamics admit a valid CCM. In [31], the authors proposed a control synthesis method for nonlinear systems where the dynamics is written as a convex combination of multiple state-dependent coefficients (i.e. $f(x, t)$ written as $A(x, t)x$) and made use of RNNs to model a time-varying metric. Again, different from our method, the controller is not jointly learned with the metric. In contrast, our approach can simultaneously synthesize the controller and the CCM certificate without any additional structural assumptions for the system.