Learning-based Robust Motion Planning with Guaranteed Stability: A Contraction Theory Approach

Hiroyasu Tsukamoto and Soon-Jo Chung

Abstract—This paper presents Learning-based Autonomous Guidance with RObustness and Stability guarantees (LAG-ROS), which provides machine learning-based nonlinear motion planners with formal robustness and stability guarantees, by designing a differential Lyapunov function using contraction theory. LAG-ROS utilizes a neural network to model a robust tracking controller independently of a target trajectory, for which we show that the Euclidean distance between the target and controlled trajectories is exponentially bounded linearly in the learning error, even under the existence of bounded external disturbances. We also present a convex optimization approach that minimizes the steady-state bound of the tracking error to construct the robust control law for neural network training. In numerical simulations, it is demonstrated that the proposed method indeed possesses superior properties of robustness and nonlinear stability resulting from contraction theory, whilst retaining the computational efficiency of existing learning-based motion planners.

Index Terms—Machine Learning for Robot Control, Robust/Adaptive Control, and Optimization & Optimal Control.

I. INTRODUCTION

In the near future of robotic exploration, teams of robots are expected to perform complex decision-making tasks autonomously in extreme environments, where their motions are typically governed by nonlinear dynamics with external disturbances. For such operations to be successful, they need to compute optimal motion plans online while robustly guaranteeing convergence to the target trajectory, both with their limited onboard computational resources. Thus, this work aims to propose a learning-based robust motion planning and control algorithm that meets these challenging requirements.

Related Work: Learning-based control designs have been an emerging area of research since the rise of neural networks and reinforcement learning [1], [2]. Model-free approaches learn optimal policies using data obtained in real-world environments, making them robust but not suitable for situations where sampling large training datasets is difficult. Also, proving the robustness and stability properties of such data-driven systems is challenging in general, although some techniques

do exist [3], [4]. In contrast, model-based methods allow sampling as much data as we want to design the policies by, e.g., imitation learning [5], [6], reinforcement learning [7], [8], or both [9]–[11]. However, the learned controller could yield cascading errors in the real-world environment if its nominal model poorly represents the true underlying dynamics [12].

Control theoretical approaches to circumvent such difficulties include robust tube-based motion planning [13]–[26] equipped with a tracking control law for robustness and stability certificates. Among these are contraction theory-based robust control [13]–[21], which tracks a target trajectory computed externally by existing motion planners, and thereby robustly keeps the system trajectories in a control invariant tube that satisfies given state constraints. Although these provable guarantees are promising, they still assume that the target trajectory can be computed online solving some planning problems, unlike the aforementioned learning frameworks.

Contributions: In this paper, we present Learning-based Autonomous Guidance with RObustness and Stability guarantees (LAG-ROS) as a novel way to bridge the gap between the learning-based and robust tube-based motion planners. In particular, while LAG-ROS requires one neural network evaluation to get its control input as in the learning schemes [5]–[11], its contraction theory-based architecture still allows obtaining formal robustness and stability guarantees as in [13]–[21]. This framework depicted in Fig. 1 is summarized as follows.

1) *Robustness and Stability Guarantees:* The theoretical foundation of LAG-ROS rests on contraction theory, which utilizes a contraction metric to characterize a necessary and sufficient condition of exponential incremental stability of nonlinear system trajectories [27]. The central result of this paper is that, if there exists a control law which renders a nonlinear system contracting, or equivalently, the closed-loop system has a contraction metric, then LAG-ROS trained to model the controller ensures the Euclidean distance between the target and controlled trajectories to be bounded exponentially with time, linearly in the learning error and size of perturbation. This property helps quantify how small the learning error should be in practice, giving some guidance in choosing design parameters of neural net training. We further show that such a contracting control law and its corresponding contraction metric can be designed explicitly via convex optimization, using the method of CV-STEM [15]–[17] to minimize a steady-state upper bound of the LAG-ROS tracking error.

2) *State Constraint Satisfaction:* We further exploit the computed bound on the tracking error in generating expert

Manuscript received: February, 24, 2021; Revised: May, 14, 2021; Accepted: June, 1, 2021.

This paper was recommended for publication by Editor Dana Kulic upon evaluation of the Associate Editor and Reviewers' comments.

The authors are with the Graduate Aerospace Laboratories, California Institute of Technology, Pasadena, CA, {htsukamoto, sjchung}@caltech.edu (Code: <https://github.com/astrohiro/lagros>).

Digital Object Identifier (DOI): see top of this page.

Fig. 1. Illustration of LAG-ROS: See Theorems 1 – 4 and Algorithm 1 for detailed explanation on each block of the LAG-ROS design. Note that LAG-ROS requires only one neural net evaluation to get u_L , and Theorem 1 provides its robustness and stability guarantees.

TABLE I
COMPARISON OF THE PROPOSED METHOD WITH THE LEARNING-BASED AND ROBUST TUBE-BASED MOTION PLANNERS.

Motion planning scheme	Policy to be learned	State tracking error $\ x - x_d\ $	Computational load
(a) Learning-based motion planner [5]–[11]	$(x, o_\ell, t) \mapsto u_d$	Increases exponentially (Lemma 1)	One neural net evaluation
(b) Robust tube-based motion planner [13]–[26]	$(x, x_d, u_d, t) \mapsto u^*$	Exponentially bounded (Theorem 2)	Computation required to get (x_d, u_d)
(c) Proposed method (LAG-ROS)	$(x, o_\ell, t) \mapsto u^*$	Exponentially bounded (Theorem 1)	One neural net evaluation

demonstrations for training, so that the learned policy will not violate given state constraints even with the learning error and external disturbances. In this phase, LAG-ROS learns the contracting control law independently of a target trajectory, making it implementable without solving any motion planning problems online unlike [13]–[26]. The performance of LAG-ROS is evaluated in cart-pole balancing [28] and nonlinear motion planning of multiple robotic agents [29] in a cluttered environment, demonstrating that LAG-ROS indeed satisfies the formal exponential bound as in [13]–[21] with its computational load as low as that of existing learning-based motion planners [5]–[11]. In particular, LAG-ROS requires less than 0.1s for computation in all of these tasks and achieves higher control performances and task success rates (see Sec. V-A3), when compared with the existing motion planners in Table I which outlines the differences of these schemes from our proposed method.

Notation: For $x \in \mathbb{R}^n$ and $A \in \mathbb{R}^{n \times m}$, we let $\|x\|$, δx , and $\|A\|$, denote the Euclidean norm, infinitesimal variation of x , and induced 2-norm, respectively. For a square matrix A , we use the notation $A \succ 0$, $A \succeq 0$, $A \prec 0$, and $A \preceq 0$ for the positive definite, positive semi-definite, negative definite, negative semi-definite matrices, respectively, and $\text{sym}(A) = (A + A^\top)/2$. Also, $I \in \mathbb{R}^{n \times n}$ denotes the identity matrix.

II. LEARNING-BASED ROBUST MOTION PLANNING WITH GUARANTEED STABILITY (LAG-ROS)

In this paper, we consider the following nonlinear systems with a controller $u \in \mathbb{R}^m$ (f & B are known but d is unknown):

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

$$\dot{x}_d(o_g, t) = f(x_d(o_g, t), t) + B(x_d(o_g, t), t)u_d(x_d(o_g, t), o_g, t) \quad (2)$$

where $t \in \mathbb{R}_{\geq 0}$, $f: \mathbb{R}^n \times \mathbb{R}_{\geq 0} \mapsto \mathbb{R}^n$, $B: \mathbb{R}^n \times \mathbb{R}_{\geq 0} \mapsto \mathbb{R}^{n \times m}$, $x: \mathbb{R}_{\geq 0} \mapsto \mathbb{R}^n$ is the state trajectory of the true system (1) perturbed by the bounded disturbance $d: \mathbb{R}^n \times \mathbb{R}_{\geq 0} \mapsto \mathbb{R}^n$ s.t.

$\sup_{x,t} \|d(x,t)\| = \bar{d}$, $o_g \in \mathbb{R}^g$ is a vector containing global environment information such as initial and terminal states, states of obstacles and other agents, etc., and $x_d: \mathbb{R}^g \times \mathbb{R}_{\geq 0} \mapsto \mathbb{R}^n$ and $u_d: \mathbb{R}^n \times \mathbb{R}^g \times \mathbb{R}_{\geq 0} \mapsto \mathbb{R}^m$ are the target trajectories given by existing motion planning algorithms, e.g., [13]–[26].

A. Problem Formulation of LAG-ROS

We seek to find u that is computable with one neural network evaluation and guarantees exponential boundedness of $\|x - x_d\|$ in (1) and (2), robustly against the learning error and external disturbances. The objective is thus not to develop new learning-based planners that compute (x_d, u_d) , but to augment them with formal robustness and stability guarantees. To this end, let us review the following existing planning techniques:

- (a) Learning-based motion planner [5] or [6]–[11]: $(x, o_\ell(x, o_g), t) \mapsto u_d(x, o_g, t)$, modeled by a neural network, where $o_\ell: \mathbb{R}^n \times \mathbb{R}^g \mapsto \mathbb{R}^\ell$ with $\ell \leq g$ is local environment information extracted from $o_g \in \mathbb{R}^g$ [5].
- (b) Robust tube-based motion planner [13] or [14]–[26]: $(x, x_d, u_d, t) \mapsto u^*(x, x_d, u_d, t)$, where u^* is a contraction theory-based tracking controller, e.g., in Theorem 2.

The robust tube-based motion planner (b) ensures that the perturbed trajectories x of (1) stay in an exponentially bounded error tube around the target trajectory x_d of (2) [13]–[21] (see Theorem 2). However, it requires the online computation of (x_d, u_d) as an input to their control policy, which is not realistic for systems with limited computational resources.

The learning-based motion planner (a) circumvents this issue by modeling the target policy $(x, o_\ell, t) \mapsto u_d$ by a neural network. In essence, our approach, to be proposed in Theorem 1, is for providing (a) with the contraction theory-based stability guarantees (b). We remark that (a) can only assure the tracking error $\|x - x_d\|$ to be bounded by a function which exponentially increases with time, as to be shown in Lemma 1 for comparison with LAG-ROS of Theorem 1.

B. Stability Guarantees of LAG-ROS

The approach of LAG-ROS bridges the gap between (a) and (b) by ensuring that the distance between the target and controlled trajectories to be exponentially bounded.

(c) Proposed approach (LAG-ROS, see Fig. 1):

$(x, o_\ell(x, o_g), t) \mapsto u^*(x, x_d(o_g, t), u_d(x_d(o_g, t), o_g, t), t)$ with o_ℓ of (a), modeled by a neural network u_L of Theorem 1, where u^* of (b) is viewed as a function of (x, o_ℓ, t) .

Theorem 1: Suppose that (1) is controlled to track (2) by the LAG-ROS $u_L = u_L(x, o_\ell(x, o_g), t)$, learned to satisfy

$$\|u_L - u^*(x, x_d(o_g, t), u_d(x_d(o_g, t), o_g, t), t)\| \leq \varepsilon_\ell, \quad \forall x, o_g, t \quad (3)$$

where $\varepsilon_\ell \in [0, \infty)$ is the learning error, u^* is the robust control policy of (b) (to be designed in Theorem 2), and (x_d, u_d) is given by the robust motion planner (b) (to be discussed in Theorem 3). Now consider the following virtual system of $y(\mu, t)$ parameterized by $\mu \in [0, 1]$, which has $y(\mu = 1, t) = x$ of (1) and $y(\mu = 0, t) = x_d$ of (2) as its particular solutions:

$$\dot{y} = \zeta(y, x, x_d, u_d, t) + d_y(\mu, x, x_d, u_d, o_g, t) \quad (4)$$

where ζ is parameterized by y to verify $\zeta|_{y=x} = f(x, t) + B(x, t)u^*(x, x_d, u_d, t)$ and $\zeta|_{y=x_d} = f(x_d, t) + B(x_d, t)u_d$ (see, e.g., (10)), $d_y = \mu(B(x, t)(u_L(x, o_\ell(x, o_g), t) - u^*(x, x_d, u_d, t)) + d(x, t))$, and x, x_d, u_d , and d are as defined in (1) and (2). Note that $y(\mu = 1, t) = x$ and $y(\mu = 0, t) = x_d$ are indeed particular solutions of (4). If $\exists \bar{b} \in [0, \infty)$ s.t. $\|B(x, t)\| \leq \bar{b}$, $\forall x, t$, and if u^* satisfies the following contraction conditions [30] with respect to y , for a contraction metric $M(y, x, x_d, u_d, t) = \Theta^\top \Theta \succ 0$ and $\alpha, \underline{\omega}, \bar{\omega} \in (0, \infty)$:

$$\dot{M} + 2 \text{sym} \left(M \frac{\partial \zeta}{\partial y} \right) \preceq -2\alpha M, \quad \forall y, x, x_d, u_d, t \quad (5)$$

$$\bar{\omega}^{-1} I \preceq M \preceq \underline{\omega}^{-1} I, \quad \forall y, x, x_d, u_d, t \quad (6)$$

then we have the following for $e(t) = x(t) - x_d(o_g, t)$, $\forall o_g$:

$$\|e(t)\| \leq \mathcal{R}(0) \sqrt{\bar{\omega}} e^{-\alpha t} + \frac{\bar{b}\varepsilon_\ell + \bar{d}}{\alpha} \sqrt{\frac{\bar{\omega}}{\underline{\omega}}} (1 - e^{-\alpha t}) = r_\ell(t) \quad (7)$$

where $\mathcal{R}(t) = \int_{x_d}^x \|\Theta \delta y\|$ for $M = \Theta^\top \Theta$, $x_d = x_d(o_g, t)$, and $u_d = u_d(x_d(o_g, t), o_g, t)$.

Proof: Let $V = \int_{x_d}^x \delta y^\top M \delta y = \int_{x_d}^x \|\Theta \delta y\|^2$ for (x_d, u_d) in (7). Since $\|\partial d_y / \partial \mu\| \leq \bar{b}\varepsilon_\ell + \bar{d} = \bar{d}_{\varepsilon_\ell}$, $\forall \mu, x, o_g, t$, for d_y in (4) with such (x_d, u_d) , the contraction condition (5) gives

$$\begin{aligned} \dot{V} &\leq \int_{x_d}^x \delta y^\top \left(\dot{M} + 2 \text{sym} \left(M \frac{\partial \zeta}{\partial y} \right) \right) \delta y + 2\bar{d}_{\varepsilon_\ell} \int_{x_d}^x \|M \delta y\| \\ &\leq -2\alpha V + (2\bar{d}_{\varepsilon_\ell} / \sqrt{\bar{\omega}}) \mathcal{R}(t). \end{aligned}$$

Since $d(\|\Theta \delta y\|^2)/dt = 2\|\Theta \delta y\|(d\|\Theta \delta y\|/dt)$, this implies that $\mathcal{R} \leq -\alpha \mathcal{R} + \bar{d}_{\varepsilon_\ell} / \sqrt{\bar{\omega}}$. Therefore, applying the comparison lemma [31, pp.102-103, pp.350-353] (i.e., if $\dot{v}_1 \leq h(v_1, t)$ for $v_1(0) \leq v_2(0)$ and $\dot{v}_2 = h(v_2, t)$, then $v_1(t) \leq v_2(t)$), along with the relation $\mathcal{R}(t) \geq \|e(t)\| / \sqrt{\bar{\omega}}$, results in (7) ■

Theorem 1 implies that the bound (7) decreases linearly in the learning error ε_ℓ and disturbance $\bar{d}$, and (1) controlled by LAG-ROS is exponentially stable when $\varepsilon_\ell = 0$ and $\bar{d} = 0$, showing a great improvement over (a) which only gives an exponentially diverging bound as to be derived in

Lemma 1 [5]–[11]. This property permits quantifying how small $\varepsilon_\ell = 0$ should be to meet the required guidance and control performance, giving some intuition on the neural network architecture (see Sec. V). Also, since we model u^* by $u_L(x, o_\ell, t)$ independently of x_d , it is indeed implementable without solving any motion planning problems online unlike the robust motion planners (b) [13]–[26], as outlined in Table I. If we can sample training data of u^* explicitly considering the bound (7), the LAG-ROS control enables guaranteeing given state constraints even with the learning error ε_ℓ and external disturbance $d(x, t)$, as will be seen in Sec. III and IV.

Remark 1: ε_ℓ of (3) can only be found empirically in practice, and thus we propose one way to generate training data with (7) and use the test error of u_L as ε_ℓ (see Sec. IV-B and Sec. V). Note that models $f(x, t)$ learned by system identification for a more accurate description of the nominal dynamics, e.g., [32], is still utilizable in (4) as long as the learned system is contracting [3], and the modeling error is bounded. Other types of perturbations, such as stochastic or parametric uncertainty, could be handled using [17], [33].

To appreciate the importance of the guarantees in Theorem 1, let us additionally show that (a), which models $(x, o_\ell, t) \mapsto u_d$, only leads to an exponentially diverging bound.

Lemma 1: Suppose that u of (1) is learned to satisfy

$$\|u(x, o_\ell(x, o_g), t) - u_d(x, o_g, t)\| \leq \varepsilon_\ell, \quad \forall x, o_g, t \quad (8)$$

for u_d of (2) with the learning error $\varepsilon_\ell \in [0, \infty)$, and that $\exists \bar{b}$ s.t. $\|B(x, t)\| \leq \bar{b}$, $\forall x, t$. If $f_{cl} = f(x, t) + B(x, t)u_d(x, o_g, t)$ is Lipschitz, i.e., $\exists L_f \in [0, \infty)$ s.t. $\|f_{cl}(x_1, o_g, t) - f_{cl}(x_2, o_g, t)\| \leq L_f \|x_1 - x_2\|$, $\forall x_1, x_2 \in \mathbb{R}^n$, then we have the following bound:

$$\|e(t)\| \leq \|e(0)\| e^{L_f t} + L_f^{-1} (\bar{b}\varepsilon_\ell + \bar{d}) (e^{L_f t} - 1). \quad (9)$$

where $e = x - x_d$, and x, x_d , and $\bar{d}$ are given in (1) and (2).

Proof: Integrating (1) and (2) for u in (8) yields $\|e(t)\| \leq \|e(0)\| + L_f \int_0^t \|e(\tau)\| d\tau + (\bar{b}\varepsilon_\ell + \bar{d})t$. Applying the Gronwall-Bellman inequality [31, pp. 651] gives

$$\|e(t)\| \leq \|e(0)\| + \bar{d}_{\varepsilon_\ell} t + L_f \int_0^t (\|e(0)\| + \bar{d}_{\varepsilon_\ell} \tau) e^{L_f(t-\tau)} d\tau$$

where $\bar{d}_{\varepsilon_\ell} = \bar{b}\varepsilon_\ell + \bar{d}$. Thus, integration by parts results in the desired relation (9). ■

Lemma 1 indicates that if there exists either a learning error ε_ℓ or external disturbance d , the tracking error bound grows exponentially with time, and thus (9) becomes no longer useful for large t . Section V demonstrates how the computed bounds of (7) ($\lim_{t \rightarrow \infty} e^{-\alpha t} = 0$) and (9) ($\lim_{t \rightarrow \infty} e^{L_f t} = \infty$) affect the control performance in practice.

III. CONTRACTION THEORY-BASED ROBUST AND OPTIMAL TRACKING CONTROL

Theorem 1 is subject to the assumption that we have a contraction theory-based robust tracking control law u^* , which satisfies (5) and (6) for a given (x_d, u_d) . This section thus delineates one way to extend the method called ConVex Optimization-based Steady-state Tracking Error Minimization (CV-STEM) [15]–[17] to find a contraction metric M of Theorem 1, which minimizes an upper bound of the steady-state error of (7) via convex optimization. Minimizing the

bound renders the tube-based planning of Theorem 3 to be given in Sec. IV for sampling training data less conservative, resulting in a better optimal solution for (x_d, u_d) .

In addition, we modify the CV-STEM in [15] to derive a robust control input u^* which also greedily minimizes the deviation of u^* from the target u_d , using the computed contraction metric M to construct a differential Lyapunov function $V = \delta y^\top M \delta y$. Note that u^* is to be modeled by a neural network which maps (x, o_ℓ, t) to u^* implicitly accounting for (x_d, u_d) as described in Theorem 1, although u^* takes (x, x_d, u_d, t) as its inputs (see Sec. IV-B).

A. Problem Formulation of CV-STEM Tracking Control

For given (x_d, u_d) , we assume that u^* of Theorem 1 can be decomposed as $u^*(x, x_d, u_d, t) = u_d + K(x, x_d, u_d, t)(x - x_d)$, the generality of which is guaranteed by the following lemma.

Lemma 2: Consider a general tracking controller u defined as $u = k(x, x_d, u_d, t)$ with $k(x_d, x_d, u_d, t) = u_d$, where $k: \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}_{\geq 0} \mapsto \mathbb{R}^m$. If k is piecewise continuously differentiable, then $\exists K: \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}_{\geq 0} \mapsto \mathbb{R}^{m \times n}$ s.t. $u = k(x, x_d, u_d, t) = u_d + K(x, x_d, u_d, t)(x - x_d)$.

Proof: We have $u = u_d + (k(x, x_d, u_d, t) - k(x_d, x_d, u_d, t))$ for $x_d = (x_d, u_d)$ due to $k(x_d, x_d, u_d, t) = u_d$. Since $k(x, x_d, u_d, t) - k(x_d, x_d, u_d, t) = \int_0^1 (dk(cx + (1-c)x_d, x_d, t)/dc)dc$, choosing K as $\int_0^1 (\partial k(cx + (1-c)x_d, x_d, t)/\partial x)dc$ gives the desired relation. ■

Lemma 2 implies that designing optimal k of $u^* = k(x, x_d, u_d, t)$ reduces to designing optimal $K(x, x_d, u_d, t)$ of $u^* = u_d + K(x, x_d, u_d, t)(x - x_d)$. When (1) is controlled by the LAG-ROS u_L of Theorem 1 with such u^* , the virtual system of (1) and (2) which has $y = x$ and $y = x_d$ as its particular solutions can be given by (4), where ζ is defined as follows:

$$\zeta = \dot{x}_d + (A(x, x_d, u_d, t) + B(x, t)K(x, x_d, u_d, t))(y - x_d) \quad (10)$$

where A is the State-Dependent Coefficient (SDC) form of the dynamical system (1) given by Lemma 2 of [17], which verifies $A(x, x_d, u_d, t)(x - x_d) = f(x, t) + B(x, t)u_d - f(x_d, t) - B(x_d, t)u_d$. Note that ζ indeed satisfies $\zeta|_{y=x} = f(x, t) + B(x, t)u^*$ and $\zeta|_{y=x_d} = f(x_d, t) + B(x_d, t)u_d$ for such A , to have $y = x$ and $y = x_d$ as the particular solutions to (4).

B. CV-STEM Contraction Metrics as Lyapunov Functions

The remaining task is to construct M so that it satisfies (5) and (6). The CV-STEM approach suggests that we can find such M via convex optimization to minimize an upper bound of (7) as $t \rightarrow \infty$ when α of (5) is fixed. Theorem 2 proposes using the metric M designed by the CV-STEM for a Lyapunov function, thereby augmenting u^* with additional optimality to greedily minimize $\|u^* - u_d\|^2$ for u_d in (2).

Theorem 2: Suppose that f and B are piecewise continuously differentiable, and let $B = B(x, t)$ and $A = A(x, x_d, u_d, t)$ in (10) for notational simplicity. Consider a contraction metric $M(x, x_d, u_d, t) = W(x, x_d, u_d, t)^{-1} \succ 0$ given by the following convex optimization (CV-STEM) [15]–[17] to minimize an upper bound on the steady-state tracking error of (7):

$$J_{CV}^* = \min_{v > 0, \chi \in \mathbb{R}, \bar{W} \succ 0} \frac{\bar{b}\epsilon_\ell + \bar{d}}{\alpha} \chi \quad \text{s.t. (12) and (13)} \quad (11)$$

with the convex constraints (12) and (13) given as

$$-\dot{\bar{W}} + 2\text{sym}(A\bar{W}) - 2vBR^{-1}B^\top \preceq -2\alpha\bar{W}, \quad \forall x, x_d, u_d, t \quad (12)$$

$$I \preceq \bar{W}(x, x_d, u_d, t) \preceq \chi I, \quad \forall x, x_d, u_d, t \quad (13)$$

where $\alpha, \omega, \bar{\omega} \in (0, \infty)$, $v = 1/\omega$, $\chi = \bar{\omega}/\omega$, $\bar{W} = vW$, and $R = R(x, x_d, u_d, t) \succ 0$ is a weight matrix on the control input. Suppose also that u^* of Theorem 1 is designed as $u^* = u_d + K^*(x, x_d, u_d, t)e$, where $e = x - x_d$, and K^* is found by the following convex optimization for given (x, x_d, u_d, t) :

$$K^* = \arg \min_{K \in \mathbb{R}^{m \times n}} \|u - u_d\|^2 = \arg \min_{K \in \mathbb{R}^{m \times n}} \|K(x, x_d, u_d, t)e\|^2 \quad (14)$$

$$\text{s.t. } \dot{M} + 2\text{sym}(MA + MBK(x, x_d, u_d, t)) \preceq -2\alpha M. \quad (15)$$

Then M satisfies (5) and (6) for ζ defined in (10), and thus (7) holds, i.e., we have the exponential bound on the tracking error $\|x - x_d\|$ when the dynamics (1) is controlled by the LAG-ROS control input u_L of Theorem 1. Furthermore, the problem (14) is always feasible.

Proof: Since the differential dynamics of (4) with (10) is given as $\delta \dot{y} = (\partial \zeta / \partial y) \delta y = (A - BK) \delta y$, substituting this into (5) verifies that (5) and (15) are equivalent. For $\bar{K} = -R^{-1}B^\top M$, (15) can be rewritten as

$$v^{-1}M(-\dot{\bar{W}} + 2\text{sym}(A\bar{W}) - vBR^{-1}B^\top)M \preceq -2\alpha v^{-1}M\bar{W}M.$$

Since this is clearly feasible as long as M satisfies the condition (12), this implies that the problem (14) is always feasible. Also, multiplying (6) by W_s s.t. $W = W_s^2$ with $W_s \succ 0$ from both sides gives (13) [15]. These facts indicate that the conditions (5) and (6) are satisfied for M and u^* constructed by (11) and (14), respectively, and thus we have the exponential bound (7) as a result of Theorem 1. Furthermore, the problem (11) indeed minimizes an upper bound of (7) as $t \rightarrow \infty$ due to the relation $0 \leq \sqrt{\bar{\omega}/\omega} = \sqrt{\chi} \leq \chi$. We remark that (11) is convex as the objective is affine in χ and (12) and (13) are linear matrix inequalities in terms of v , χ , and $\bar{W}$. ■

Remark 2: (11) and (14) are convex in terms of their respective decision variables and thus can be solved computationally efficiently [34, pp. 561]. For systems with a known Lyapunov function (e.g. Lagrangian systems [35, pp. 392]), we could simply use it to get robust tracking control u^* in Theorem 2 without solving (14), although optimality may no longer be guaranteed in this case.

Remark 3: The contraction metric construction itself can be performed using a neural network [16]–[18], [33], leading to an analogous incremental stability and robustness results to those of Theorem 1 [33], [36].

IV. GENERATING EXPERT DEMONSTRATIONS

This section presents how to sample training data, i.e., the target trajectory (x_d, u_d) of (2) and corresponding perturbed state and CV-STEM robust control (x, u^*) , so that x of (1) controlled by the LAG-ROS control u_L will stay in given admissible state space as long as we have $\|u_L - u^*\| \leq \epsilon_\ell$ as in (3) of Theorem 1. To be specific, since u^* of Theorem 2 solves (11) to obtain an optimal error tube around x_d , i.e., (7) of Theorem 1, we exploit it in the tube-based motion planning [13] for generating training data, which satisfies

Fig. 2. Illustration of state sampling in robust bounded error tube.

given state constraints even under the existence of the learning error ε_ℓ and disturbance d in (7) and (11). Note that the learning error ε_ℓ and disturbance upper bound $\bar{d}$ of Theorem 1 are assumed to be selected *a priori* [32].

A. Tube-based State Constraint Satisfaction in LAG-ROS

Given the learning error ε_ℓ and disturbance upper bound $\bar{d}$ of Theorem 1, we sample (x_d, u_d) of (2) by solving the following tube-based global motion planning problem, for each element in a training set containing randomized o_g of (1):

$$\min_{\substack{\bar{x}=\bar{x}(o_g, t) \\ \bar{u}=\bar{u}(\bar{x}, o_g, t)}} \int_0^T c_1 \|\bar{u}\|^2 + c_2 P(\bar{x}, \bar{u}, t) dt \quad (16)$$

s.t. $\dot{\bar{x}} = f(\bar{x}, t) + B(\bar{x}, t)\bar{u}$, $\bar{x} \in \bar{\mathcal{X}}(o_g, t)$, and $\bar{u} \in \bar{\mathcal{U}}(o_g, t)$

where $c_1 \geq 0$, $c_2 \geq 0$, $P(\bar{x}, \bar{u}, t)$ is some performance-based cost function, $T > 0$ is a given time horizon, $\bar{\mathcal{X}}$ is robust admissible state space defined as $\bar{\mathcal{X}}(o_g, t) = \{v(t) \in \mathbb{R}^n \mid \forall \xi(t) \in \{\xi(t) \mid \|v(t) - \xi(t)\| \leq r_\ell(t)\}, \xi(t) \in \mathcal{X}(o_g, t)\}$, $\mathcal{X}(o_g, t)$ is given admissible state space, $r_\ell(t)$ is the right-hand side of (7) in Theorem 1, and $\bar{u} \in \bar{\mathcal{U}}(o_g, t)$ is an input constraint. The following theorem shows that the LAG-ROS control ensures the perturbed state x of (1) to satisfy state constraints $x \in \mathcal{X}$, due to the contracting property of u^* in Theorem 2.

Theorem 3: If the solution of (16) yields (x_d, u_d) of (2), the LAG-ROS control u_L of Theorem 1, where M and u^* are designed by Theorem 2, ensures the perturbed solution $x(t)$ of (1) to stay in the admissible state space $\mathcal{X}(o_g, t)$, i.e., $x(t) \in \mathcal{X}(o_g, t), \forall t$, even with the learning error ε_ℓ of u_L in (3).

Proof: Since $x_d \in \bar{\mathcal{X}}$, we have $\xi \in \mathcal{X}$ for all $\xi \in S(x_d) = \{\xi \mid \|x_d - \xi\| \leq r_\ell\}$ by the definition of $\bar{\mathcal{X}}$. Also, if we control (1) by u_L , we have $\|x - x_d\| \leq r_\ell$, or equivalently, $x \in S(x_d)$. These two statements indicate that $x \in \mathcal{X}$. ■

Remark 4: Theorem 3 implies that if x_d is sampled by (16), the perturbed trajectory (1) controlled by LAG-ROS u_L of Theorem 1 will not violate the given state constraints as long as LAG-ROS u_L of Theorem 1 satisfies (3). This helps greatly reduce the need for safety control schemes such as [37].

B. Learning Contraction Theory-based Robust Control

To meet the learning error requirement (3) for the sake of robustness and state constraint satisfaction proposed in Theorems 1 and 3, we should also sample x to get training data for the CV-STEM robust control inputs u^* of Theorem 2.

Theorem 4: If (3) of Theorem 1 is satisfied $\forall x(t) \in S(x_d) = \{\xi \in \mathbb{R}^n \mid \|x_d(o_g, t) - \xi\| \leq r_\ell(t)\}$ and $\forall o_g, t$ instead of $\forall x, o_g, t$, then u_L with such a constraint still guarantees (7). Also, if we always choose $x(0) = x_d(o_g, 0)$ for the perturbed state x in (1), the radius of the tube $S(x_d)$ is bounded $\forall o_g, t$.

Proof: Following the proof of Theorem 1, we still get (3) for all $x(t) \in S(x_d(o_g, t))$, and (3) implies that the perturbed state x of (1) indeed lies in $S(x_d(o_g, t))$. When $x(0) = x_d(o_g, 0)$, i.e., $\mathcal{R}(0) = 0$ in (7), $r_\ell(t)$ is independent of the state x and thus the radius of the tube $S(x_d)$ is bounded. ■

Using Theorem 4, we can sample x in the bounded error tube $S(x_d(o_g, t))$ with $x(0) = x_d(o_g, 0)$ to obtain training data for the CV-STEM robust control u^* by (14) of Theorem 2. These samples are to be used for training the LAG-ROS neural network, so we can satisfy the condition (3) of Theorem 1 for the pre-defined learning error ε_ℓ .

Remark 5: The localization method in [5] allows extracting o_ℓ of (a) by o_g of (16), to render LAG-ROS applicable to different environments in a distributed way (see Sec. V).

The pseudocode for the offline construction of the LAG-ROS control of Theorem 1 is presented in Algorithm 1 (see Fig. 1 for its visual description). Once we get $u_L(x, o_\ell, t)$ by Algorithm 1, $u = u_L$ in (1) can be easily computed by only evaluating the neural net u_L for a given (x, o_ℓ, t) observed at (x, t) , whilst ensuring robustness, stability, and state constraint satisfaction due to Theorems 1, 2, and 3.

Algorithm 1: LAG-ROS Algorithm

Inputs : Random environment information $\{(o_g)_i\}_{i=1}^N$
 Contraction metric M of Theorem 2
 Motion planner $\mathcal{P}$
 Learning error ε_ℓ and disturbance bound $\bar{d}$

Outputs: LAG-ROS control u_L of Theorem 1

for $i \leftarrow 1$ **to** N **do**

Solve (16) of Theorem 3 for $(o_g)_i$ using $\mathcal{P}$ and obtain a target trajectory and environment observation history $(x_d, u_d, o_g, o_\ell, t)_i$
 Sample D robust CV-STEM control $\{(x, u^*)_{ij}\}_{j=1}^D$ using Theorems 2 and 4 (see Fig. 2)

Model $(x, o_\ell, t)_{ij} \mapsto u_{ij}^*$ by a neural net to satisfy $\|u_L - u^*\| \leq \varepsilon_\ell$ as in (3) of Theorem 1 (see Fig. 1)

V. SIMULATION

Our proposed LAG-ROS framework is demonstrated using multiple motion planning problems under external disturbances (<https://github.com/astrohiro/lagros>). CVXPY [38] with the MOSEK solver [39] is used to solve optimization problems in Theorems 2 and 3 for sampling training data.

A. Simulation Setup

The maximum admissible control time interval is selected to be $\Delta t_{\max} = 0.1$ (s). The computational time of each framework is measured for the Macbook Pro laptop (2.2 GHz Intel Core i7, 16 GB 1600 MHz DDR3 RAM), and each simulation result is the average of 50 simulations for each random environment and disturbance realization.

1) *Neural Network Training:* We use a neural network u_L with 3 layers and 100 neurons. The network is trained using stochastic gradient descent with training data sampled by Theorems 2–4, and the loss function is defined as $\|u_L - u^*\|$

to satisfy the learning error bound ε_ℓ of Theorem 1. We use 50000 training samples and select $\varepsilon_\ell = 0.01$ for all the tasks, but these numbers can be modified accordingly depending on situations, considering the required performance for the error bound (3) in Theorem 1. Note that we can guarantee (3) only empirically, using the test error computed with a given test set as in standard neural network training.

2) *Environment Information*: o_g of (2) is selected as initial states, target terminal states, and states of obstacles and other agents, if any. In Sec. V-C, the deep set framework [40] is used to extract local information o_ℓ of (a) from o_g as in [5].

3) *Performance Measure*: The objective function of Theorem 3 for sampling (x_d, u_d) of (2) is selected as $\int_0^T \|\ddot{u}\|^2 dt$. Since LAG-ROS is not for proposing a new trajectory optimization solver but for augmenting it with the formal guarantees of Theorem 1, one could use any other objective, e.g., information-based cost [41], as long as (x_d, u_d) is obtainable.

We define success of each task as the situation where the agent reaches, avoiding collisions, if any, a given target terminal state x_f within a given time horizon $\mathcal{T} \geq T$, i.e. $\exists t^* \in [0, \mathcal{T}]$ s.t. $\|x(t^*) - x_f\| \leq r_\ell(t^*)$ for r_ℓ in (7), where the value for r_ℓ is to be defined in the subsequent sections. The success rate is computed as the percentage of successful trials in the total 50 simulations. Also, we evaluate the performance of each planner, (a), (b), and (c) in Sec. II, by the objective function $\int_0^{t^*} \|\ddot{u}\|^2 dt$ if the task is successfully completed, and $\int_0^T \|\ddot{u}\|^2 dt$ otherwise, where T is the nominal time horizon.

4) *External Disturbances*: As shown in Lemma 1 and Theorem 1, the tracking error bound of learning-based motion planners (a) increases exponentially with time, whilst it decreases exponentially for the proposed approach (c). To make such a difference clear, we consider the situations where $d(x, t)$ of (1) is non-negligible, and thus (a) tends to fail the task of Sec V-A3 due to Lemma 1. This is not to argue that (a) [5]–[11] should be replaced by (c), but to imply that they can be improved further to have the guarantees of Theorem 1.

5) *Computational Complexity*: Since (a) and LAG-ROS (c) are implementable with one neural net evaluation at each t , their performance is also compared with (b) which requires solving motion planning problems to get its control input. Its time horizon is selected to make the trajectory optimization [42] solvable online considering the current computational power, for the sake of a fair comparison. We denote the computational time as Δt in this section, and it should be less than the maximum control time interval $\Delta t_{\max}$, i.e., $\Delta t \leq \Delta t_{\max} = 0.1(\text{s})$.

B. Cart-Pole Balancing

We first consider the cart-pole balancing task in Fig. 3 to demonstrate the differences of (a) – (c) summarized in Table I. Its dynamics is given in [28], [33], and we use $g = 9.8$, $m_c = 1.0$, $m = 0.1$, $\mu_c = 0.5$, $\mu_p = 0.002$, and $l = 0.5$.

1) *LAG-ROS Training*: T , $\mathcal{T}$, and x_f in Sec. V-A3 are selected as $T = \mathcal{T} = 9(\text{s})$ and $x_f = [p_f, 0, 0, 0]^\top$ for x in Fig. 3, where p_f is a random terminal position at each episode. We let $\mathcal{R}(0) = 0$ and $\bar{d}_\varepsilon = b\varepsilon_\ell + \bar{d} = 0.75$ in (7), and train the neural net of Sec. V-A1 to have $(\bar{d}_\varepsilon/\alpha)\sqrt{\chi} = 3.15$ with $\alpha = 0.60$ using Algorithm 1. The performance of LAG-ROS

Fig. 3. Cart-pole balancing task: $x = [p, \theta, \dot{p}, \dot{\theta}]^\top$, x and x_d are given in (1) and (2), (a) – (c) are given in Sec. II, and r_ℓ is given in (7). The shaded area denotes the standard deviation ($+\sigma$ and $-\sigma$).

TABLE II

CONTROL PERFORMANCES FOR CART-POLE BALANCING ($\bar{d}_\varepsilon = 0.75$). NOTE THAT THE NOTATIONS ARE AS GIVEN IN SEC. V-A3 AND V-A5.

	Success rate (%)	$\int \ u(t)\ ^2 dt$	Δt (s)
(a)	40.0	4.79×10^2	3.01×10^{-2}
(b)	100.0	5.19×10^2	1.00×10^{-1}
(c)	100.0	5.67×10^2	3.01×10^{-2}

(c) is compared with the learning-based planner (a) and robust tube-based planner (b).

2) *Simulation Results and Discussions*: Note that the robust tube-based motion planner (b) (i.e. u^* of Theorem 1) is computable with $\Delta t \leq \Delta t_{\max}$ (see Sec. V-A5) in this case, and thus we expect that the performance of LAG-ROS (c) should be worse than that of (b), as (c) is a neural net model that approximates (b). The right-hand side of Fig. 3 shows the tracking error $\|x - x_d\|$ of (1) and (2) averaged over 50 simulations at each time instant t . It is still interesting to see that LAG-ROS of (c) and u^* of (b) indeed satisfies the exponential bound (7) of Theorem 1 given as

$$r_\ell(t) = \mathcal{R}(0)\sqrt{\bar{\omega}} + (\bar{d}_\varepsilon/\alpha)\sqrt{\chi}(1 - e^{-\alpha t}) = 3.15(1 - e^{-0.60t})$$

for all t with a small standard deviation σ , unlike learning-based motion planner (a) with a diverging bound (9) and increasing deviation σ , as can be seen from Fig. 3. Contraction theory enables such quantitative analysis on robustness and stability of learning-based planners, which is one of the major advantages of our proposed technique.

Table II shows the control performance and computational cost of (a) – (c), which is a good summary of their differences and trade-offs aforementioned in Sec. II and in Table I:

- (a) approximates u_d of (1), and thus requires lower computational cost $\Delta t = 0.03(\text{s})$ with a smaller objective value $\int \|u\|^2 dt = 479$, but robustness is not guaranteed (Lemma 1) resulting in a 40% success rate.
- (b) computes u^* of Theorem 2, and thus possesses robustness as in Fig. 3 resulting in a 100% success rate, but requires larger $\Delta t = 0.1(\text{s})$ to compute x_d .
- (c) approximates u^* independently of x_d , and thus possesses robustness of Theorem 1 as in Fig. 3 resulting in a 100% success rate, even with $\Delta t = 0.03(\text{s})$ as small as that of (b). It yields 9.2% larger $\int \|u\|^2 dt$ than (b) as it models $(x, o_\ell, t) \mapsto u^*$, not $(x, o_\ell, t) \mapsto u_d$.

It is demonstrated that LAG-ROS indeed possesses the robustness and stability guarantees of u^* as in Theorem 1, unlike (a), with significantly lower computational cost than that of (b) as expected.

C. Multi-Agent Nonlinear Motion Planning

The advantages of (c) demonstrated in Sec. V-B are more appreciable in the problem settings where the robust tube-based motion planner (b) is no longer capable of computing a global solution with $\Delta t \leq \Delta t_{\max}$ (see Sec. V-A5). We thus consider motion planning and collision avoidance of multiple robotic simulators [29] in a cluttered environment with external disturbances, where the agents are supposed to perform tasks based only on local observations o_ℓ of (a). Note that its nonlinear equation of motion is given in [29] and all of its parameters are normalized to 1.

1) *LAG-ROS Training*: We select T , $\mathcal{T}$, and x_f in Sec. V-A3 as $T = 30(\text{s})$, $\mathcal{T} = 45(\text{s})$, and $x_f = [p_{xf}, p_{yf}, 0, 0]^\top$, where (p_{xf}, p_{yf}) is a random position in $(0, 0) \leq (p_{xf}, p_{yf}) \leq (5, 5)$. We let $\mathcal{R}(0) = 0$, and train the neural network of Sec. V-A1 by Algorithm 1 to get $(\bar{d}_\varepsilon/\alpha)\sqrt{\chi} = 0.125$ with $\alpha = 0.30$ in (7). In particular, the following error tube (7):

$$r_\ell(t) = (\bar{d}_\varepsilon/\alpha)\sqrt{\chi}(1 - e^{-\alpha t}) = 0.125(1 - e^{-0.30t})$$

is used for sampling (x_d, u_d) by Theorem 3 with an input constraint $u_i \geq 0$, $\forall i$ to avoid collisions with a random number of multiple circular obstacles (0.5m in radius) and of other agents, even under the learning error and disturbances. When training LAG-ROS, the input constraint is satisfied by using a ReLU function for the network output, and the localization technique [5] is used to extract o_ℓ of (a) from o_g of (1) for its distributed implementation, with the communication radius 2.0m. Its performance is compared with (a), (b), and a centralized planner (d) which is not computable with $\Delta t \leq \Delta t_{\max}$, where Δt is given in Sec. V-A5:

(d) Centralized robust motion planner:

$$(x, x_d, u_d, t) \mapsto u^*, \text{ offline centralized solution of (b).}$$

2) *Remarks on Sub-optimal Trajectories*: Multi-agent problems have sub-optimal solutions with optimal values close to the global optimum [43], and thus (c) does not necessarily take the same x_d as that of the centralized planner (d) in the presence of disturbances, as depicted in Fig. 5. However, it implicitly guarantees tracking to its own (sub-)optimal x_d due to Theorem 1, which means we can still utilize their objective value $\int \|u\|^2 dt$ and success rate to evaluate their performance.

3) *Implication of Simulation Results*: Figure 4 shows one example of the trajectories of the motion planners (a) – (d), under external disturbances with $\bar{d} = \sup_{x,t} \|d(x,t)\| = 0.4$.

- For (a), the tracking error accumulates exponentially with time due to $\bar{d}$ (Lemma 1), which necessitates the use of safety control for avoiding collisions [5]. Such non-optimal control inputs also increase the error in (9) as (a) does not possess any robustness guarantees.
- For (b), robustness is guaranteed by Theorem 2 but can only obtain locally optimal (x_d, u_d) of (2), as its time horizon has to be small enough to make the problem solvable within $\Delta t \leq \Delta t_{\max} = 0.1(\text{s})$, and the agent only has access to local information. This renders some agents stuck in local minima as depicted in Fig. 4.
- LAG-ROS (c) tackles these two problems by providing formal robustness and stability guarantees of Theorems 1 – 4, whilst implicitly knowing the global solution (only

TABLE III
COMPUTATIONAL TIME OF EACH MOTION PLANNER FOR MULTI-AGENT NONLINEAR MOTION PLANNING.

	(a)	(b)	(c)	(d)
Δt (s)	4.63×10^{-2}	1.53×10^{-1}	4.66×10^{-2}	1.12×10^3

from the local information o_ℓ as in [5]) without computing it online. It satisfies the given state constraints due to Theorem 3 as can be seen from Fig. 4

4) *Simulation Results and Discussions*: Figure 5 and Table III summarize the simulation results which corroborate the arguments of Sec. V-C3 implied by Fig. 4:

- (a) satisfies the requirement on the computational cost since we have $\Delta t \leq \Delta t_{\max} = 0.1(\text{s})$ as shown in Table III, but its success rate remains the lowest for all $\bar{d} = \sup_{x,t} \|d(x,t)\|$ due to the cumulative error (9) of Lemma 1, as can be seen in Fig. 5 and Table III. Although its objective value is the smallest for $\bar{d} \leq 0.6$ since it models u_d , it gets larger for larger $\bar{d}$, due to the lack of robustness to keep x around x_d .
- (b) has a success rate higher than that of (a), but still lower than 50% as it can only compute sub-optimal x_d under the limited computational capacity. In fact, we have $\Delta t \geq \Delta t_{\max}$ in this case, which means it requires a slightly better onboard computer. Also, it uses excessive control effort larger than 10^3 due to such sub-optimality.
- (c) achieves more than 90% success rates for all $\bar{d}$, and its objective value remains only 1.64 times larger than that of the centralized planner (d), even without computing (x_d, u_d) of (2) online. Its computational cost is as low as that of (a), satisfying $\Delta t \leq \Delta t_{\max} = 0.1(\text{s})$ while retaining a standard deviation smaller than (a).

These results imply that LAG-ROS indeed enhances learning-based motion planners (a) with robustness and stability guarantees of contraction theory as in (d) (and (b)), thereby bridging the technical gap between them.

VI. CONCLUSION

In this work, we propose a new learning-based motion planning framework, called LAG-ROS, with the formal robustness and stability guarantees of Theorem 1. It extensively utilizes contraction theory to provide an explicit exponential bound on the distance between the target and controlled trajectories, even under the existence of the learning error and external disturbances. Simulation results demonstrate that it indeed satisfies the bound in practice, thereby yielding consistently high success rates and control performances in contrast to the existing motion planners (a) and (b). Note that other types of disturbances can be handled, using [17] for stochastic systems and [33] for parametric uncertain systems.

Acknowledgments: This work was in part funded by the Jet Propulsion Laboratory, California Institute of Technology, and benefited from discussions with J. Castillo-Rogez, M. D. Ingham, and J.-J. E. Slotine.

REFERENCES

- [1] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. MIT Press, 1998.

Fig. 4. Trajectories for the learning-based planner (a), robust tube-based planner (b), LAG-ROS (c), and offline centralized solution (d) (○: start, ●: goal).

Fig. 5. Control performances versus $\sup_{x,t} \|d(x,t)\|$ of (1) for multi-agent nonlinear motion planning. Note that (a) – (d) are as given in Fig. 4, and the performances are as defined in Sec. V-A3. The shaded area denotes the standard deviation ($\pm 10^{-1}\sigma$).

- [2] D. P. Bertsekas and J. N. Tsitsiklis, *Neuro-Dynamic Programming*. Athena Scientific, 1996.
- [3] N. M. Boffi, S. Tu, N. Matni, J.-J. E. Slotine, and V. Sindhiani, "Learning stability certificates from data," in *CoRL*, Nov. 2020.
- [4] D. Ho and J. C. Doyle, "Robust model-free learning and control without prior knowledge," in *IEEE CDC*, 2019, pp. 4577–4582.
- [5] B. Rivière, W. Hönig, Y. Yue, and S. Chung, "GLAS: Global-to-local safe autonomy synthesis for multi-robot motion planning with end-to-end learning," *IEEE Robot. Automat. Lett.*, vol. 5, no. 3, pp. 4249–4256, 2020.
- [6] J. Carius, F. Farshidian, and M. Hutter, "Mpc-net: A first principles guided policy search," *IEEE Robot. Automat. Lett.*, vol. 5, no. 2, pp. 2897–2904, 2020.
- [7] M. Everett, Y. F. Chen, and J. P. How, "Motion planning among dynamic, decision-making agents with deep reinforcement learning," in *IEEE/RSJ IROS*, 2018, pp. 3052–3059.
- [8] F. Berkenkamp, M. Turchetta, A. Schoellig, and A. Krause, "Safe model-based reinforcement learning with stability guarantees," in *Adv. Neural Inf. Process. Syst.*, vol. 30. Curran Associates, Inc., 2017, pp. 908–918.
- [9] J. Ho and S. Ermon, "Generative Adversarial Imitation Learning," in *Adv. Neural Inf. Process. Syst.*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, Eds., vol. 29. Curran Associates, Inc., 2016, pp. 4565–4573.
- [10] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social GAN: Socially acceptable trajectories with generative adversarial networks," in *IEEE/CVF Conf. Comput. Vision Pattern Recognit.*, 2018, pp. 2255–2264.
- [11] A. Kuefler, J. Morton, T. Wheeler, and M. Kochenderfer, "Imitating driver behavior with generative adversarial networks," in *IEEE Intell. Vehicles Symp.*, 2017, pp. 204–211.
- [12] S. Ross and D. Bagnell, "Efficient reductions for imitation learning," in *13th Int. Conf. Artif. Intell. Statist.*, ser. Proc. Mach. Learn. Res., vol. 9, May 2010, pp. 661–668.
- [13] S. Singh, A. Majumdar, J.-J. E. Slotine, and M. Pavone, "Robust online motion planning via contraction theory and convex optimization," in *IEEE ICRA*, May 2017, pp. 5883–5890.
- [14] I. R. Manchester and J.-J. E. Slotine, "Control contraction metrics: Convex and intrinsic criteria for nonlinear feedback design," *IEEE Trans. Autom. Control*, vol. 62, no. 6, pp. 3046–3053, Jun. 2017.
- [15] H. Tsukamoto and S.-J. Chung, "Robust controller design for stochastic nonlinear systems via convex optimization," *IEEE Trans. Autom. Control*, to appear, Oct. 2021.
- [16] H. Tsukamoto and S.-J. Chung, "Neural contraction metrics for robust estimation and control: A convex optimization approach," *IEEE Control Syst. Lett.*, vol. 5, no. 1, pp. 211–216, 2021.
- [17] H. Tsukamoto, S.-J. Chung, and J.-J. E. Slotine, "Neural stochastic contraction metrics for learning-based control and estimation," *IEEE Control Syst. Lett.*, vol. 5, no. 5, pp. 1825–1830, 2021.
- [18] D. Sun, S. Jha, and C. Fan, "Learning certified control using contraction metric," in *CoRL*, Nov. 2020.
- [19] B. T. Lopez, J.-J. E. Slotine, and J. P. How, "Dynamic tube MPC for nonlinear systems," in *Proc. Amer. Control Conf.*, 2019, pp. 1655–1662.
- [20] A. Lakshmanan, A. Gahlawat, and N. Hovakimyan, "Safe feedback motion planning: A contraction theory and $\mathcal{L}_1$ -adaptive control based approach," in *IEEE CDC*, 2020, pp. 1578–1583.
- [21] D. Sun, M. J. Khojasteh, S. Shekhar, and C. Fan, "Uncertainty-aware safe exploratory planning using Gaussian process and neural control contraction metric," arXiv:2105.06567, 2021.
- [22] A. Bemporad and M. Morari, "Robust model predictive control: A survey," in *Robustness in identification and control*. London: Springer, 1999, pp. 207–226.
- [23] D. Mayne, M. Seron, and S. Raković, "Robust model predictive control of constrained linear systems with bounded disturbances," *Automatica*, vol. 41, no. 2, pp. 219 – 224, 2005.
- [24] D. Q. Mayne, E. C. Kerrigan, E. J. van Wyk, and P. Falugi, "Tube-based robust nonlinear model predictive control," *Int. J. Robust Nonlinear Control*, vol. 21, no. 11, pp. 1341–1353, 2011.
- [25] J. Schulman, Y. Duan, J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, "Motion planning with sequential convex optimization and convex collision checking," *Int. J. Robot. Res.*, vol. 33, no. 9, pp. 1251–1270, 2014.
- [26] Y. K. Nakka, A. Liu, G. Shi, A. Anandkumar, Y. Yue, and S. J. Chung, "Chance-constrained trajectory optimization for safe exploration and learning of nonlinear systems," *IEEE Robot. Automat. Lett.*, vol. 6, no. 2, pp. 389–396, 2021.
- [27] W. Lohmiller and J.-J. E. Slotine, "On contraction analysis for nonlinear systems," *Automatica*, no. 6, pp. 683 – 696, 1998.
- [28] A. G. Barto, R. S. Sutton, and C. W. Anderson, "Neuronlike adaptive elements that can solve difficult learning control problems," *IEEE Trans. Syst. Man Cybern.*, vol. SMC-13, no. 5, pp. 834–846, 1983.
- [29] Y. Nakka *et al.*, "A six degree-of-freedom spacecraft dynamics simulator for formation control research," in *Astrodyn. Special. Conf.*, 2018.
- [30] W. Wang and J.-J. E. Slotine, "On partial contraction analysis for coupled nonlinear oscillators," *Biol. Cybern.*, vol. 92, no. 1, pp. 38–53, Jan. 2005.
- [31] H. K. Khalil, *Nonlinear Systems*, 3rd ed. Prentice-Hall, 2002.
- [32] G. Shi *et al.*, "Neural lander: Stable drone landing control using learned dynamics," in *IEEE ICRA*, May 2019.
- [33] H. Tsukamoto, S.-J. Chung, and J.-J. E. Slotine, "Learning-based adaptive control via contraction theory," *IEEE Control Syst. Lett.*, under minor revision, Mar. 2021.
- [34] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, Mar. 2004.
- [35] J.-J. E. Slotine and W. Li, *Applied Nonlinear Control*. Upper Saddle River, NJ: Pearson, 1991.
- [36] H. Tsukamoto, S.-J. Chung, J.-J. E. Slotine, and C. Fan, "A tutorial on neural contraction metrics for learning-based control with guaranteed stability," in *IEEE CDC*, under review, 2021.
- [37] S. Kolathaya and A. D. Ames, "Input-to-state safety with control barrier functions," *IEEE Control Syst. Lett.*, vol. 3, no. 1, pp. 108–113, 2019.
- [38] S. Diamond and S. Boyd, "CVXPY: A Python-embedded modeling language for convex optimization," *J. Mach. Learn. Res.*, 2016.
- [39] MOSEK ApS, *MOSEK Optimizer API for Python 9.0.105*, 2020.
- [40] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Poczos, R. R. Salakhutdinov, and A. J. Smola, "Deep sets," in *Adv. Neural Inf. Process. Syst.*, vol. 30. Curran Associates, Inc., 2017, pp. 3391–3401.
- [41] Y. K. K. Nakka *et al.*, "Information-based guidance and control architecture for multi-spacecraft on-orbit inspection," in *AIAA Scitech*, 2021.
- [42] S. Singh, V. Sindhiani, J.-J. E. Slotine, and M. Pavone, "Learning stabilizable dynamical systems via control contraction metrics," in *Workshop Algorithmic Found. Robot.*, 2018.
- [43] W. Hönig, J. A. Preiss, T. K. S. Kumar, G. S. Sukhatme, and N. Ayanian, "Trajectory planning for quadrotor swarms," *IEEE Trans. Robot.*, vol. 34, no. 4, pp. 856–869, 2018.