

# Statistical Safety and Robustness Guarantees for Feedback Motion Planning of Unknown Underactuated Stochastic Systems

Craig Knuth<sup>1</sup>, Glen Chou<sup>2</sup>, Jamie Reese<sup>1</sup>, Joseph Moore<sup>1,3</sup>

**Abstract**—We present a method for providing statistical guarantees on runtime safety and goal reachability for integrated planning and control of a class of systems with unknown nonlinear stochastic underactuated dynamics. Specifically, given a dynamics dataset, our method jointly learns a mean dynamics model, a spatially-varying disturbance bound that captures the effect of noise and model mismatch, and a feedback controller based on contraction theory that stabilizes the learned dynamics. We propose a sampling-based planner that uses the mean dynamics model and simultaneously bounds the closed-loop tracking error via a learned disturbance bound. We employ techniques from Extreme Value Theory (EVT) to estimate, to a specified level of confidence, several constants which characterize the learned components and govern the size of the tracking error bound. This ensures plans are guaranteed to be safely tracked at runtime. We validate that our guarantees translate to empirical safety in simulation on a 10D quadrotor, and in the real world on a physical CrazyFlie quadrotor and Clearpath Jackal robot, whereas baselines that ignore the model error and stochasticity are unsafe.

## I. INTRODUCTION

To deploy robots in the real world, we need assurances on runtime safety and robustness. However, real environments and hardware present a myriad of challenges to attaining such guarantees, such as unknown dynamics, system stochasticity arising from non-repeatable environmental interactions, and underactuation. These challenges establish safe real-world planning and control as a difficult open problem. To handle unknown system dynamics, a model is often obtained via black-box machine learning, e.g., via neural networks (NNs). Such methods are flexible and can accurately model complex phenomena, like driving on rough terrain or moving through fluid media like water or air. However, model prediction error can vary across the relevant domain. This can be catastrophic for planning and control, which can exploit inaccuracies to produce unreasonable plans that cannot be safely tracked.

To ensure that plans computed using the learned model can be safely tracked on the true system, we must characterize the noise profile of the true system together with the model mismatch, and identify how these errors can impact the performance of downstream planning and control. For unknown stochastic systems, this ultimately requires a statistical argument, as only a finite set of sample rollouts can be obtained on the true system and statements on safety must be made using this data. Therefore, we apply




Fig. 1: Comparison of our noise bounding approach compared to [2], [3] for example system  $h(x) = \dot{x}$ . From data (green dots), our method estimates a bound (shaded blue) around a learned model (blue) in a domain  $D$  (between black lines). This bound is verified with the addition of a small buffer (shaded green). Note this bound is not valid outside of  $D$ . Pairwise slopes on the data are used to estimate the Lipschitz constant in [2], [3], but the ratio of change in  $\dot{x}$  to change in  $x$  is unbounded due to noise, e.g. the red star data points.

methods from Extreme Value Theory (EVT) [1], a suite of statistical methods for estimating the maxima and minima of unknown distributions from data, to bound the model error in a trusted domain  $D$  and propagate its effect through the planner and controller to make statistical guarantees on closed-loop behavior. Our method builds upon prior work that has applied EVT for safe planning with learned models of *deterministic* fully actuated [2] and underactuated [3] systems with *noiseless* training data. These methods use the *Lipschitz constant* (i.e., the maximum rate of change) of the model error to inform planning. However, these approaches are unusable for stochastic systems, where the Lipschitz constant becomes unbounded due to noise (see Fig. 1). This corresponds to infinite model error bounds which are useless for informing downstream planning. This limitation precludes [2], [3] from scaling to real-world robots where noise and stochasticity are inevitable.

To address this gap, we propose a novel method to bound the error in the learned model which can be applied on stochastic systems and *on real robots*. We bound how this model error can impact a downstream tracking controller based on contraction theory, which applies to a broad class of underactuated systems, providing statistically-guaranteed closed-loop tracking tubes for any plan in the trusted domain. Our planner is a modified RRT that considers the tracking tube in calculating collision checks, ensures feedback control can be faithfully executed, and ensures that plans stay within the trusted domain where the EVT-based analysis is valid. This ensures runtime safety and goal reachability up to the statistical confidence. This confidence does not degrade with time, as it is made with respect to the model and controller, rather than specific plans. We summarize our method in Fig. 2. Our contributions:

- A novel model error bound that can be applied on stochastic systems with noisy training data

<sup>1</sup>Johns Hopkins University Applied Physics Laboratory, Laurel, MD, USA. [Craig.Knuth@jhuapl.edu](mailto:Craig.Knuth@jhuapl.edu), [Jamie.Reese@jhuapl.edu](mailto:Jamie.Reese@jhuapl.edu), [Joseph.Moore@jhuapl.edu](mailto:Joseph.Moore@jhuapl.edu)

<sup>2</sup>University of Michigan, Ann Arbor, MI, USA. [gchou@umich.edu](mailto:gchou@umich.edu)

<sup>3</sup> Johns Hopkins University Department of Mechanical Engineering, Baltimore, MD, USA.

- Derivations for propagating this error bound through the planner and controller, and designing a planner that uses the resulting tracking bound to return safer plans
- Validation on a simulated 10D quadrotor, a real 6D CrazyFlie quadrotor, and a real 5D Clearpath Jackal, outperforming baselines in safety and robustness

## II. RELATED WORK

Our work is related to feedback motion planning of uncertain systems. One body of work [5], [6], [7], [8] ensures safe feedback motion planning under worst-case error via tools from reachability analysis and contraction theory, under known, constant model error bounds. In contrast, we learn a tighter spatially-varying error bound. Adaptive contraction-based controllers are another solution, but existing methods [9], [10] assume a known model uncertainty structure, i.e., that it lies in the span of known basis functions. We consider uncertainty arising from the model error and noise distribution, which have no known structure. Stochastic contraction [11], [12], i.e., contraction in expectation, has been studied, but does not easily provide high-probability tracking tubes. Learning-based method to contraction-based control are well-suited for controlling the deep dynamics models we consider; however, existing methods [13], [14] assume known dynamics and constant error bounds, or do not consider tracking error [15].

Our work is also related to safe planning for stochastic systems via tools like chance constraints [16], [17], [18], [19], [20], [21], [22] and occupation measures [23], [24]. These methods reason about an (approximate) distribution over states, but in doing so, require accurate knowledge of the noise distribution, and have difficulties scaling (in state dimension and plan length [16], [17]). To sidestep these issues, we do not seek out detailed distributional information, instead solving a simpler robust planning problem that uses a high-probability estimate on the *support* of the disturbance.

Finally, our work is related to safe learning-based control. One class of methods learns stability certificates: [25],[26] learn Lyapunov functions from data; however, it is difficult to apply a single Lyapunov function for general point-to-point motion planning. In [27], the method derives a model-free stability certificate, but does not consider control design. Other methods use Gaussian processes (GPs) for reachable tube estimation [28] and safe exploration [29], [30], but assume that a controller is given; we do not make these assumptions. [31] learns trajectory tracking tubes; however, plans must remain near full training *trajectories* to be accurate, which can heavily restrict planning; we only require plans to be near *state/controls* seen in training. Finally, most relevant are [2], [3], which plan safely with learned dynamics by deriving an invariant tube around a plan inside a “trusted domain” near the training data. A restrictive key assumption of [2], [3] is that the unknown system is *deterministic*, preventing these approaches from scaling to real-world systems with stochasticity and noise. In this paper, we make these methods applicable to a broad class of stochastic systems by making fundamental improvements

to [2], [3], i.e., in designing new a model error bound and deriving its impact on tracking error.

## III. PRELIMINARIES AND PROBLEM FORMULATION

Let  $\dot{x} \sim h(x, u)$  be the true continuous stochastic dynamics where  $h : \mathcal{X} \times \mathcal{U} \rightarrow \mathcal{X}$ ,  $x \in \mathcal{X}$  is the state,  $u \in \mathcal{U}$  is the control, and  $\dot{x}$  is the time derivative of  $x$ , i.e., for each  $x, u$ , there is a distribution of possible state derivatives. Let  $\dim(\mathcal{X}) = n_x$  and  $\dim(\mathcal{U}) = n_u$ . Furthermore, let  $\dot{x} = \hat{h}(x, u) = f(x) + g(x)u$  be a deterministic control-affine approximation of the dynamics that is paired with an estimate of the maximum disturbance at  $(x, u)$  denoted  $w(x, u) : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}_{n_x}^+$  where  $\mathbb{R}_{n_x}^+$  is the set of positive real vectors with dimension  $n_x$ . Let  $S = \{(x_i, u_i, \dot{x}_i)\}_{i=0}^N$  and  $V = \{(x_i, u_i, \dot{x}_i)\}_{i=0}^M$  be a training and validation dataset of  $N$  and  $M$  transitions collected on the true system respectively. Let  $D = D_x \times D_u \subset \mathcal{X} \times \mathcal{U}$  be the trusted domain partitioned into state and control subsets.

Let  $\bar{Q}$  be the symmetric part of  $Q$ , i.e.  $\bar{Q} = \frac{1}{2}(Q + Q^\top)$ . We denote  $\bar{\lambda}(Q)$  and  $\underline{\lambda}(Q)$  as the largest and smallest eigenvalues of  $Q$  respectively. Furthermore for matrix valued functions  $Q(x)$  we define  $\bar{\lambda}_D(Q(x)) = \max_{x \in D} \bar{\lambda}(Q(x))$ , similarly for  $\underline{\lambda}_D(Q(x))$ . Let  $|v|$  refer to the element-wise absolute value of a vector  $v$  and  $v_k$  denote the  $k$ th element of  $v$ . We define addition between a vector  $v$  and scalar  $a$  as  $a$  added to each element of  $v$ . Let  $\mathbf{0}_{n \times m}$  denote a zero matrix and  $\mathbf{I}_{n \times n}$  denote the identity matrix.

Let  $\mathcal{S}_{>0}$  be the set of positive definite matrices with dimension  $n_x \times n_x$ . For a smooth manifold  $\mathcal{X}$ , a Riemannian metric is defined as  $M(x) : \mathcal{X} \rightarrow \mathcal{S}_{>0}$ . The distance between two points  $x_1$  and  $x_2$ ,  $d(x_1, x_2)$ , is defined as the length of the geodesic between them (see [3]). We calculate geodesics using the psuedospectral method of [32]. In this paper  $d$  refers to Riemannian distance,  $\|\cdot\|$  refers to the Euclidean norm, and  $\|\cdot\|_{\bar{M}}$  refers to the metric defined under  $\bar{M} \in \mathcal{S}_{>0}$ . We also note that the Euclidean distance between two points  $x_1, x_2$  on a Riemannian manifold with metric  $M(x)$  can be bounded as follows:  $\|x_1 - x_2\| \leq d(x_1, x_2) / \underline{\lambda}_D(M(x))$  (see [3]). Let a ball be defined as  $\mathcal{B}_{\bar{M}}(\chi, r) \doteq \{\chi' \mid \|\chi' - \chi\|_{\bar{M}} \leq r\}$ . If  $\bar{M}$  is omitted, then the Euclidean norm is used.

Note the specific form of  $D$  is left unspecified. We only require the ability to check if a ball of some radius  $r_1$  fits in the domain, i.e.  $\mathcal{B}_{\bar{M}}((x, u), r_1) \subset D$  (see [2]). In practice, we define  $D$  either as a union of  $\ell_2$ -balls with radius  $r_2 \geq r_1$  about a subset  $\bar{S}$  of our dataset  $S$ , i.e.  $\cup_{(x,u,\dot{x}) \in \bar{S}} \mathcal{B}_{\bar{M}}((x, u), r_2)$  or a hyper-rectangle in the state-control space, depending on the data collection procedure.

We make the following assumptions. One, the true noise distribution is bounded. Two, in order to construct a valid control contraction metric (CCM, [4]), we assume that  $\mathcal{X}$  is a smooth manifold and the true dynamics are locally incrementally exponentially stabilizable, i.e. there exists a  $\beta, \lambda$ , and a feedback controller such that  $\|x^*(t) - x(t)\| \leq \beta e^{-\lambda t} \|x^*(0) - x(0)\|$  in the trusted domain  $D$ . We do not assume the true system is control affine. Finally, we present the problem we wish to solve:

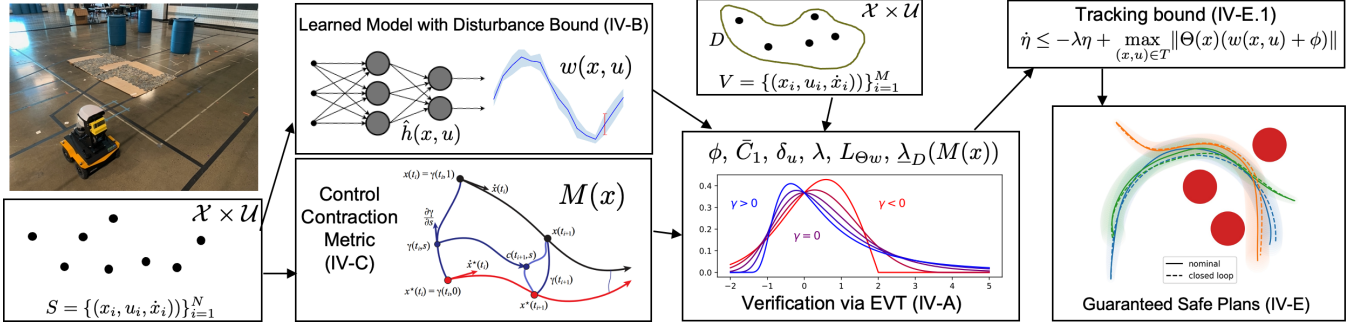


Fig. 2: Flow chart of method. CCM figure sourced from [4]. Sections given where applicable.

**Problem:** Given datasets  $S$  and  $V$ , learn a dynamics model, design a tracking controller, and construct a trusted domain  $D$  for planning. At planning time, given a start  $x_I$ , goal  $x_G$ , goal tolerance  $\mu$ , and safe set  $\mathcal{X}_{\text{safe}}$ , plan a nominal trajectory  $x^* : [0, T] \rightarrow \mathcal{X}$ ,  $u^* : [0, T] \rightarrow \mathcal{U}$  under the learned dynamics  $\hat{h}$  such that  $x^*(0) = x_I$ ,  $\dot{x}^* = \hat{h}(x^*, u^*)$ ,  $\|x^*(T) - x_G\| \leq \mu$ , and  $x^*(t)$ ,  $u^*(t)$  remains in  $D \cap \mathcal{X}_{\text{safe}}$  for all  $t \in [0, T]$ . Also, guarantee that in tracking  $(x^*(t), u^*(t))$  under the true dynamics  $f$ , the system remains in  $D \cap \mathcal{X}_{\text{safe}}$  and reaches  $\mathcal{B}(x_G, \hat{\mu} + \mu)$  where  $\hat{\mu}$  is the tracking error bound.

#### IV. METHOD

Our approach rests on learning three models: the approximate control affine dynamics  $\hat{h}(x, u)$ , the estimated disturbance bound  $w(x, u)$ , and the CCM  $M(x)$ . These models are used in a sampling based planner that computes tracking tubes around nominal plans. Certain properties of these models are verified using EVT which guarantees the tracking tube size and consequently safety in execution.

##### A. Use of Extreme Value Theory

The safety guarantees of this work rests on Extreme Value Theory, specifically the Fisher-Tippet-Gnedenko Theorem (FTG) [1] which guarantees that sample batch maximums of a random variable  $Z$  converge to the generalized extreme value distribution (GEV). We collect  $N_s$  batches of size  $N_b$  of i.i.d. samples,  $z_i^j \in Z$ , and fit the sample batch maximums,  $\{\max_{1 \leq i \leq N_b} z_i^j\}_{j=1}^{N_s}$ , to GEV with shape ( $\xi$ ), location ( $\gamma$ ), and scale parameters. If  $\xi < 0$  then the maximum of  $Z$  is finite and can be overestimated with a given confidence  $\psi$  using  $\gamma$ . Knowing that samples follow the GEV distribution, we perform a Kolmogorov-Smirnov goodness-of-fit test [33] to assure fit quality with only finite samples (in our results we use  $p = 0.05$ ). See [2] for more details.

##### B. Robust Model Characterization

We train  $\hat{h}$  using mean squared error loss and  $w$  with a hinge loss to encourage the following relationship.

$$|h(x, u) - \hat{h}(x, u)| \leq w(x, u) \quad (1)$$

Above,  $\leq$  is taken to mean element-wise. Merely training does not guarantee the above relationship is satisfied, so we employ EVT to estimate the maximum of  $Z$  over  $D$ .

$$Z \sim \max_k (|h(x, u) - \hat{h}(x, u)| - w(x, u))_k \quad (2)$$

Note that if  $Z \leq 0$  with probability 1, then  $w$  is never an underestimate of the maximum deviation from the learned model for any  $(x, u)$  pair. By fitting sample batch maximums from our domain  $D$  to GEV, we estimate the maximum of  $Z$  up to the confidence  $\psi$  which we call  $\phi$ . Then the following relationship is statistically guaranteed with confidence  $\psi$ :

$$|h(x, u) - \hat{h}(x, u)| \leq w(x, u) + \phi \quad (3)$$

This statement is much stronger than (1) in the sense that it is verified for all  $(x, u)$  in the domain  $D$ . We craft a robust planning and control system for the verified dynamics:

$$\hat{h}(x, u) = f(x) + g(x)u + \delta, \quad |\delta| \leq w(x, u) + \phi \quad (4)$$

##### C. Controller Formulation

Our controller formulation is based on Control Contraction Metrics [34]. The core approach requires a CCM  $M(x)$ , which is a Riemannian metric that satisfies some additional conditions (to be described shortly), and a corresponding tracking feedback controller  $k(x, x^*, u^*) : \mathcal{X} \times \mathcal{X} \times \mathcal{U} \rightarrow \mathcal{U}$ . We define the executed state trajectory as  $x(t) : [0, T] \rightarrow \mathcal{X}$ . The feedback controller  $k(x, x^*, u^*)$  stabilizes the current state  $x(t)$  to the nominal state  $x^*(t)$ . The existence of such a CCM and a controller ensures that the distance between the nominal and executed trajectory decays exponentially, at a contraction rate  $\lambda$ , i.e.,  $d(x^*(t), x(t)) \leq C e^{-\lambda(t-t_0)} d(x^*(t_0), x(t_0))$ ,  $t \geq t_0$ , for some constant  $C > 0$ , when model error and noise are excluded. There are several conditions that ensure that  $M(x)$  is a CCM (see [3], [7]); here we use equations 27 and 28 of [7]:

$$g_{\perp}(x)^{\top} \left( -\partial_f W(x) + 2 \frac{\partial f(x)}{\partial x} W(x) + 2\lambda W(x) \right) g_{\perp}(x) \doteq C_1(x) \leq 0 \quad (5)$$

$$g_{\perp}(x)^{\top} \left( \partial_{g^j} W(x) + 2 \frac{\partial g^j(x)}{\partial x} W(x) \right) g_{\perp}(x) = 0, \quad j = 1 \dots n_u$$

where  $W(x) = M^{-1}(x)$ ,  $g_{\perp}(x)$  is a matrix such that  $g_{\perp}^{\top}(x)g(x) = 0$  for all  $x$ ,  $g^j(x)$  is the  $j$ th column of  $g(x)$ ,  $\partial_f W(x)$  is the Lie derivative of  $W$  with respect to  $f$ , and similarly for  $\partial_{g^j} W(x)$ . Intuitively, the first constraint is a contraction condition that is simplified by the second ‘‘orthogonality’’ condition; together, they imply the system can be made to contract at rate  $\lambda$  along arbitrary dynamically-feasible trajectories. In our results, the second condition is satisfied by thoughtful construction of  $M(x)$  and  $g(x)$ : either  $M(x)$  and  $g(x)$  are constant, or we require  $M(x)$  to be a

function of only the first  $n_x - n_u$  states and select  $g(x)$  to be of the form  $[\mathbf{0}_{n_u \times n_x - n_u} \mathbf{I}_{n_u \times n_u}]^\top$  (see [3] for details). This construction is always possible by considering an augmented version of the original system where  $u$  is treated as part of the state and  $\dot{u}$  is treated as a new virtual input.

It is known [4] that if  $M(x)$  satisfies (5) over  $D_x$ , then there always exists a controller  $k(x, x^*, u^*)$  which makes trajectories contract at rate  $\lambda$  everywhere in  $D_x$ . To implement  $k$ , we use an optimization-based controller from [7, Eqn. (40)], which returns a minimum-norm modification  $u_{\text{fb}}$  to the nominal control  $u^*$ , such that the executed control  $u(t) = u_{\text{fb}}(t) + u^*(t)$  contracts the system at rate  $\lambda$ . We can bound the magnitude of this controller based on the distance between an executed and nominal trajectory:  $\|u_{\text{fb}}(t)\| \leq \|x(t) - x^*(t)\| \delta_u$ ; see (10) of [3] for the exact definition of  $\delta_u$ . Since (5) needs to be true everywhere in  $D_x$ , but we can only evaluate (5) point-wise, we verify satisfaction with EVT. We estimate the maximum of  $\bar{\lambda}_D(C_1(x))$  over  $D_x$ , denoted  $\bar{C}_1$ , and verify that it is less than 0 up to the given confidence  $\psi$ . We also apply EVT to estimate  $\delta_u$  as in [3].

#### D. Robust Control under Disturbance

We find a valid CCM on the fixed learned dynamics by parameterizing  $W(x)$  and  $\lambda(x)$  as neural networks. In practice, learning  $\hat{h}$ ,  $w$ ,  $W$ , and  $\lambda$  concurrently yielded poor results. To satisfy (5) we only need the minimum value of  $\lambda(x)$  (which is verified with EVT) but we found in practice that parameterizing  $\lambda$  based on the state improved training performance. Our training loss is equal to  $L(B) = \sum_{i=1}^3 L_i(B)$  where  $B = \{x_1, \dots, x_{N_b}\}$  is a training batch and  $L_1$ ,  $L_2$ , and  $L_3$  are defined as

$$L_1(B) = -\alpha_1 \frac{1}{N_b} \sum_{i=1}^{N_b} \exp[\max(\bar{\lambda}(C_1(x_i)), \tau)], \quad (6)$$

$$L_2(B) = \alpha_2 \frac{1}{N_b} \sum_{i=1}^{N_b} \sqrt{\frac{\bar{\lambda}(M(x_i))}{\underline{\lambda}(M(x_i))}}, \quad (7)$$

$$L_3(B) = -\alpha_3 \min_{1 \leq i \leq N_b} \lambda(x_i), \quad (8)$$

where  $\max(\cdot, \cdot)$  returns the maximum of the two inputs. We select the threshold  $\tau$  to be  $-0.01$ . We hand-tune  $\alpha_1$ ,  $\alpha_2$ , and  $\alpha_3$  to improve performance. Furthermore,  $L_1$  promotes the satisfaction of (5) on the training data,  $L_2$  promotes the CCM to be well-conditioned, which decreases the tracking tube size, and  $L_3$  promotes large contraction rates.

While the obtained CCM provides a contracting controller for the *learned dynamics*, the closed-loop behavior on the *true system* will differ due to the model mismatch and stochastic disturbances. Fortunately, under bounded perturbations, the tracking error  $\eta(t) \doteq d(x(t), x^*(t))$  satisfies the following relation (see [8]). Here,  $\Theta(x)$  is the Cholesky decomposition of  $M(x)$ , i.e.,  $M(x) = \Theta(x)^\top \Theta(x)$ :

$$\dot{\eta}(t) \leq -\lambda \eta(t) + \|\Theta(x(t))(w(x(t), u(t)) + \phi)\|. \quad (9)$$

In Sec. IV-E, we will leverage this relation to compute disturbance-informed plans that can be safely tracked.

#### E. Planner Formulation

We now discuss our planner, which uses the learned dynamics  $\hat{h}$ , verified error bound  $w(x, u) + \phi$ , CCM  $M(x)$ , and minimum contraction rate  $\lambda$  to provide runtime assurances. We use a modified kinodynamic RRT (see Alg. 1) that maintains the accumulated tracking error along transitions. We name this algorithm Safety Analysis for Learned Stochastic Underactuated Dynamics (SAFED) RRT. In the algorithm block, `SampleState` samples a random state in  $D$ , `NearestNeighbor` returns the nearest neighbor in the tree according to the standard Euclidean norm, `SampleCandidateControl` samples a random control in  $\mathcal{U}$ , `IntegrateLearnedDyn` integrates the learned dynamics returning the state at each  $t \in [0, dt]$ .

1) *Propagating Tracking Error*: For the following section, we drop dependence on  $t$  for brevity. To obtain tubes, we cannot directly use (9), as it requires knowledge of the executed trajectory, which is unknown at planning time. We instead seek a tracking error bound that is valid for all possible executed trajectories in the tracking tube  $T \doteq \{(x, u) \mid d(x, x^*) \leq \eta \text{ and } \|u - u^*\| \leq \|x - x^*\| \delta_u\}$ :

$$\dot{\eta} \leq -\lambda \eta + \max_{(x, u) \in T} \|\Theta(x)(w(x, u) + \phi)\| \doteq H \quad (10)$$

Computing this maximum is difficult as  $\Theta$  and  $w$  are complex NNs. For smaller NNs, we can upper bound the maximum via [35], [36], but we opt for a derivation based on the triangle inequality with the nominal point on the path.

$$\begin{aligned} & \max_{(x, u) \in T} \|\Theta(x)(w(x, u) + \phi)\| \\ & \leq \|\Theta(x^*)(w(x^*, u^*) + \phi)\| + \|\Theta(x)w(x, u) - \Theta(x^*)w(x^*, u^*)\| \\ & \leq \|\Theta(x^*)(w(x^*, u^*) + \phi)\| + L_{\Theta w}(d(x, x^*) + \|u - u^*\|) \\ & \leq \|\Theta(x^*)(w(x^*, u^*) + \phi)\| + L_{\Theta w}(1 + \frac{\delta_u}{\underline{\lambda}_D(M(x))})d(x, x^*) \end{aligned} \quad (11)$$

Here,  $L_{\Theta w}$  is the mixed-norm Lipschitz constant of  $\|\Theta(x)w(x, u)\|$  as follows:

$$L_{\Theta w} \geq \sup_{(x, u) \neq (x', u') \in D} \frac{\|\Theta(x)w(x, u) - \Theta(x')w(x', u')\|}{d(x, x') + \|u - u'\|}. \quad (12)$$

For systems with swift changes between small and large noise bounds,  $L_{\Theta w}$  may be large. However, in our results, this approach provided tight tracking bounds.  $L_{\Theta w}$  and  $\underline{\lambda}_D(M(x))$  are also estimated via EVT, similar to  $\phi$ ,  $\bar{C}_1$ ,  $\delta_u$ , and  $\lambda$ . Plugging (11) into (10) and integrating over time yields an upper bound on the tracking error  $\bar{\eta}(t)$  for any plan  $x^* : [0, T] \rightarrow \mathcal{X}$  at any time  $t \in [0, T]$ , i.e.,

$$d(x(t), x^*(t)) \leq \bar{\eta}(t) \doteq \bar{\eta}(0) + \int_0^t H(\tau) d\tau. \quad (13)$$

Evaluating (13) requires integrating a scalar ODE, and can be done efficiently in planning (see line 7 of Alg. 1).

2) *Collision Checking*: We upper bound the nonconstant Riemannian metric  $M(x)$  with a constant matrix  $\bar{M} \in \mathcal{S}_{>0}$  as in [7]. The collision check for a point  $x^*(t)$  on the nominal plan reduces to checking  $\mathcal{B}_{\bar{M}}(x^*(t), \eta(t)/\underline{\lambda}_D(M(x))) \in \mathcal{X}_{\text{safe}}$ . In theory, this needs to be checked for all  $t$ , but in practice we check at discrete samples. See line 8 of Alg. 1.

3) *Staying in the Trusted Domain:* We ensure that the nominal plan remains in the trusted domain where the estimates of  $\phi$ ,  $\bar{C}_1$ ,  $\delta_u$ ,  $\lambda$ ,  $L_{\Theta w}$ , and  $\underline{\lambda}_D(M(x))$  are statistically guaranteed. This is equivalent to the ball in domain check  $\mathcal{B}_{\bar{M}}(x^*(t), \eta(t)/\underline{\lambda}_D(M(x))) \in D_x$ . See line 10 of Alg. 1.

4) *Respecting Controller Bounds:* To guarantee that the nominal plan can be safely tracked, we must ensure that the feedback control can be faithfully executed, i.e.  $u(t) \in D_u$ . Since we have an upper bound on the feedback term, this corresponds to the ball in domain check  $\mathcal{B}(u^*(t), \delta_u \eta(t)/\underline{\lambda}_D(M(x))) \in D_u$ . This also ensures our statistical guarantees hold for all feedback control values. See line 9 of Alg. 1.

*Theorem 1:* Given  $\phi$ ,  $\bar{C}_1$ ,  $\delta_u$ ,  $\lambda$ ,  $L_{\Theta w}$ , and  $\underline{\lambda}_D(M(x))$  have been appropriately overestimated (or underestimated for  $\underline{\lambda}_D(M(x))$  and  $\lambda$ ) and a plan  $x^*(t), u^*(t)$  found by Alg. 1, the following holds for the executed trajectory:  $x(t) \in \mathcal{X}_{\text{safe}}$ ,  $x(t) \in D_x$ , and  $u(t) \in D_u \forall t \in [0, T]$ .

*Proof:* Since the tracking error is appropriately bounded following the argument in Section IV-E.1, and the planner verifies that the tracking tube stays with  $\mathcal{X}_{\text{safe}}$  and  $D$ , the resulting trajectory is guaranteed to be safe and feasible. ■

As the independently estimated constants are each valid with confidence  $\psi$ , the plan is guaranteed safe and feasible with confidence  $\psi^6$  (similarly to [2]).

---

#### Algorithm 1: SAFED RRT

---

**Input:**  $x_I, x_G, \eta_I, \lambda, dt, \mu$

- 1  $\mathcal{T} \leftarrow \{(x_I, \eta_I)\}, \mathcal{P} \leftarrow \{\emptyset\}$
- 2 **while** True **do**
- 3      $x_d \leftarrow \text{SampleState}()$
- 4      $x_n, \bar{\eta}_n \leftarrow \text{NearestNeighbor}(\mathcal{T}, x_d)$
- 5      $u_c \leftarrow \text{SampleCandidateControl}()$
- 6      $x_c^*(t) \leftarrow \text{IntegrateLearnedDyn}(x_n, u_c, dt)$
- 7      $\bar{\eta}(t) \leftarrow \text{TrkErrBnd}(\bar{\eta}_n, x_c^*(t), u_c)$
- 8     **if**  $\neg \text{InCollision}(x_c^*(t), u_c, \bar{\eta}(t)) \forall t \in [0, dt] \wedge$
- 9          $\text{InDU}(x_c^*(t), u_c, \bar{\eta}(t)) \forall t \in [0, dt] \wedge$
- 10          $\text{InDX}(x_c^*(t), u_c, \bar{\eta}(t)) \forall t \in [0, dt] \wedge$
- 11          $\|x_c^*(dt) - x_d\| < \|x_n - x_d\|$  **then**
- 12              $\mathcal{T} \leftarrow \mathcal{T} \cup (x_c^*(dt), \bar{\eta}(dt)), \mathcal{P} \leftarrow \mathcal{P} \cup \{u_c\}$  ;
- 12     **if**  $\exists t, x_c^*(t) \in \mathcal{B}(x_G, \mu)$  **then** return plan;

---

## V. RESULTS

We present results on three systems: 1) a 10D simulated quadrotor, 2) a physical CrazyFlie quadrotor, and 3) a physical Clearpath Jackal. Our full method satisfies tracking bounds in all runs. We present the results of our method and comparisons to a naïve planner that does not consider the trusted domain nor tracking tubes (i.e. it excludes lines 7, 9, and 10 and uses  $\bar{\eta}(t) \doteq 0$  in line 8 of Alg. 1). We evaluate the baseline [3] by estimating a Lipschitz constant of the model error. In each case, the shape parameter  $\xi$  of GEV is positive, corresponding to an infinite Lipschitz constant, and thus, infinite tracking tubes and infeasible plans.

### A. 10D simulated quadrotor

We select the underactuated 10D quadrotor model from [7] to evaluate our approach in moderately-high dimensions. Here, we assume knowledge of the nominal

	SAFED	Naïve planner
Avg. trk. err.	0.021 $\pm$ 0.009 (0.052)	5.47e3 $\pm$ 2.88e4 (2.48e5)
Goal error	0.038 $\pm$ 0.025 (0.117)	7.24e3 $\pm$ 3.52e4 (3.95e5)

TABLE I: 10D quadrotor tracking errors  $\|x^*(t) - x(t)\|_2$ . Mean  $\pm$  standard deviation (worst case).




Fig. 3: Five plans computed between random start/goals using SAFED; associated tracking tubes are color-coded to the plan. Five executions for each plan are overlaid in black (with similar results); obstacles are in red; zoomed-in view on right.

dynamics and fit a CCM using it (here,  $\lambda = 1.29$ ); the true dynamics are perturbed with additive noise, sampled uniformly within state-dependent bounds:  $\pm[0.01, 0.01, 0.01, |0.05s_1|, |0.05c_2|, |0.05s_3|, |0.07s_7|, |0.07c_8|, |0.07s_9|]^T$  (where  $s_i, c_i$  are  $\sin(x_i), \cos(x_i)$ ), if within  $[-5, 5]$ .  $w$  is represented as a neural network with 3 hidden layers, each with 256 neurons. Using EVT on a validation dataset of size 100000, we establish a constant bound of  $\phi = 0.0075$ , with confidence  $\psi = 0.99$ . We select  $D_x = \bigcup_{i=1}^N B(x_i, r_x)$ , where  $r_x = 0.3$ , and  $D_u = \mathcal{U}$  from [7].

Using our method, we compute 20 plans in an obstacle-filled environment (cf. Fig. 3) and execute each five times. See Table II for numerical results and Figs. 3 for visualizations. The naïve planner leads to substantially larger tracking errors due to the larger model error. Comparing against [3], the fitted GEV distribution has shape  $\xi = 0.6154 > 0$  which corresponds to an infinite Lipschitz constant.

### B. Physical CrazyFlie quadrotor

We consider feedback motion planning for a learned model of a physical CrazyFlie quadrotor, moving in 3D. We select the state as the linear positions and velocities  $x = [p_x, p_y, p_z, \dot{p}_x, \dot{p}_y, \dot{p}_z]$ , and represent the learned dynamics with a linear model,  $\dot{x} = Ax + Bu$ , where  $A$  and  $B$  were fit to 5700 datapoints using linear least squares. Here,  $w$  is represented as a neural network with five hidden layers, each with 1024 neurons. Using EVT on a validation dataset of size 10800, we establish a constant bound of  $\phi = 0.0178$ , with confidence  $\rho = 0.99$ ; to obtain approximately i.i.d. samples in  $D$  required by EVT, we subsample the trajectory rollouts (justified by [37]). We select  $D_x = \bigcup_{i=1}^N B(x_i, r_x)$  and  $D_u = \bigcup_{i=1}^N B(u_i, r_u)$ , and  $r_x = 1.5$  and  $r_u = 0.1$ ; these values are selected as the 90th percentile of the minimum distance of each validation datapoint to the training data. Given the learned linear dynamics, the CCM conditions simplify to a single semidefinite constraint (i.e., the condition (5) holds uniformly for all  $x$ ); hence, the CCM is valid everywhere for the learned dynamics; we set  $\lambda = 2.0$ .

Using our method, we compute 20 plans in an obstacle-filled environment (cf. Fig. 4a) and execute each three times. See Table II for numerical results and Figs. 4a and 4b for

	SAFED	Naïve planner
Avg. trk. err.	0.066 ± 0.012 (0.121)	0.198 ± 0.044 (0.294)
Goal error	0.107 ± 0.042 (0.203)	—

TABLE II: Crazyflie tracking errors  $\|x^*(t) - x(t)\|_2$ . Mean ± standard deviation (worst case).




Fig. 4: (a) Three plans computed using SAFED; the corresponding tracking tubes are color-coded to the plan. The three executions for each plan are overlaid in black; the obstacles are in red. (b) One execution of the blue plan.

visualizations. In executing these plans, we terminate automatically if the drone deviates from the planned trajectory by more than 0.2 in any position dimension for platform safety. The naïve plans lead to premature terminations in each run (cf. Table II). Comparing against [3], the resulting fitted GEV distribution has shape  $\xi = 0.3272 > 0$  which corresponds to an infinite Lipschitz constant.

### C. Clearpath Jackal

We apply our method to a Clearpath Jackal. Our environment consisted primarily of cement but also included an approximately 1.25x2.5 meter area with artificial rough terrain, see Figure 5. The Jackal takes linear  $v$  and angular  $w$  velocities as input. We model the dynamics as:

$$\begin{bmatrix} x & y & \theta & v & w \end{bmatrix}^\top = \begin{bmatrix} f(x, u) \\ \mathbf{0}_{2 \times 1} \end{bmatrix} + \begin{bmatrix} \mathbf{0}_{3 \times 2} \\ \mathbf{I}_{2 \times 2} \end{bmatrix} \begin{bmatrix} a \\ \alpha \end{bmatrix} \quad (14)$$

where  $x$  and  $y$  correspond to the position and  $\theta$  the heading of the robot.  $a$  and  $\alpha$  are virtual acceleration inputs.  $f$  is parameterized as a NN with three hidden ReLU layers of sizes 128, 256, and 128. Our metric  $M(x)$  is parameterized as a NN, using only  $x$ ,  $y$ , and  $\theta$  as input in order to satisfy the orthogonality condition of (5). The network outputs the lower left Cholesky decomposition  $L$  of the dual metric  $W$ . The diagonal elements of  $L$  are computed via three hidden ReLU layers of sizes 256, 512, and 256 and a final SoftPlus output layer. Off-diagonal elements are computed with another network with three hidden ReLU layers of sizes 256, 512, and 256.  $w$  is composed of one hidden ReLU layer of size 128. We collected random rollouts on the data resulting in about 22000 training points and 2200 validation points over the hyperrectangle  $[-4.2, 3.2] \times [-3.1, 3.1] \times [-\pi, \pi] \times [0.3, 1.2] \times [-1.25, 1.25] \times [-0.6, 0.6] \times [-0.8, 0.8] = D \subset \mathcal{X} \times \mathcal{U}$ . Our




Fig. 5: Upper left shows jackal test environment showing two barrels, rough terrain simulant, and Clearpath Jackal. Remaining figures show an example rollout.

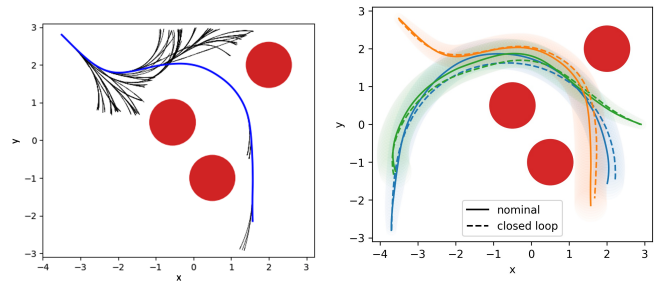


Fig. 6: Left: Example tree and found plan. Notice the tree stops expanding in the area with rough terrain. Right: Three example plans and corresponding rollouts using our full method. Tube used for collision checks is shown with shaded region. Obstacles are in red. Not shown is rough terrain region approximately spanning  $[-2.5, -1.25]$  in  $x$  and  $[-2.1, 0.3]$  in  $y$ .

	SAFED	Naïve planner
Avg. trk. err.	0.193 ± 0.0529 (0.256)	0.512 ± 0.450 (1.514)
Goal error	0.187 ± 0.0556 (0.262)	1.24 ± 1.06 (3.60)

TABLE III: Jackal tracking errors  $\|x^*(t) - x(t)\|_2$ . Mean ± standard deviation (worst case).

dynamics and CCM are trained with data without the rough terrain simulant. We collected an additional 11000 training points and 1100 validation points with the rough terrain simulant to train the bound  $w(x, u)$ . With confidence  $\psi = 0.99$ , we found  $\phi = 0.000204$ ,  $\lambda = 0.3277$ ,  $L_{\Theta w} = 0.046$ ,  $\delta_u = 0.106$ ,  $\lambda_D(M(x)) = 0.1508$ , and  $\bar{C}_1 = -0.0037$ .

Using our full method, we computed 8 plans in the environment with 3 additional barrel obstacles and executed each plan 3 times, see Figure 6. No plans were found over the rough terrain region, indicating that the model learned to avoid larger noise as it contributed to large and unsafe tracking tubes. We compare to the naïve planner and a second “unsafe” planner, computing 5 plans for each one and executing each plan 3 times. The unsafe planner only uses the nominal trajectory for collision checking rather than the full tracking tube. Tracking results are given in Table III for the full and naïve planners. The tracking results for the unsafe planner are similar to the full method, but 40% of rollouts collided with the barrels compared to no rollouts of the full method. Comparing against [3], the resulting fitted GEV distribution has shape  $\xi = 0.1582 > 0$  which corresponds to an infinite Lipschitz constant of the model error.

## VI. CONCLUSION

We present a method that is capable of planning safe, feasible paths for stochastic, underactuated, unknown dynamics. We verify the properties of a learned dynamics model, noise bound, and control contraction metric using Extreme Value Theory to bound tracking tubes around nominal plans. The planner ensures that the tubes do not collide, the tubes stay in the domain where our statistical guarantees hold, and feedback control does not exceed controller bounds. We experimentally verify on 10D simulated quadrotor, a CrazyFlie, and a Clearpath Jackal. While this method is capable of handling stochastic systems, it does so in a robust way by always considering worst case. Potentially the conservativeness can be reduced by learning the distribution of noise and utilizing probabilistic collision bounds in planning. This method could also be extended to a real-time receding horizon planner in order to plan in unknown environments.

## REFERENCES

- [1] L. De Haan and A. Ferreira, *Extreme value theory: an introduction*. Springer Science & Business Media, 2007.
- [2] C. Knuth, G. Chou, N. Ozay, and D. Berenson, "Planning with learned dynamics: Probabilistic guarantees on safety and reachability via lipschitz constants," *IEEE RA-L*, 2021. [Online]. Available: <https://arxiv.org/abs/2010.08993>
- [3] G. Chou, N. Ozay, and D. Berenson, "Model error propagation via learned contraction metrics for safe feedback motion planning of unknown systems," in *2021 60th IEEE Conference on Decision and Control (CDC), Austin, TX, USA, December 14-17, 2021*. IEEE, 2021, pp. 3576–3583.
- [4] I. R. Manchester and 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, 2017.
- [5] S. Bansal, M. Chen, S. L. Herbert, and C. J. Tomlin, "Hamilton-jacobi reachability: A brief overview and recent advances," in *56th IEEE Annual Conference on Decision and Control, CDC 2017, Melbourne, Australia, December 12-15, 2017*. IEEE, 2017, pp. 2242–2253.
- [6] A. Majumdar and R. Tedrake, "Funnel libraries for real-time robust feedback motion planning," *IJRR*, vol. 36, no. 8, pp. 947–982, 2017.
- [7] S. Singh, B. Landry, A. Majumdar, J. E. Slotine, and M. Pavone, "Robust feedback motion planning via contraction theory," 2019.
- [8] G. Chou, N. Ozay, and D. Berenson, "Safe output feedback motion planning from images via learned perception modules and contraction theory," in *Workshop on the Algorithmic Foundations of Robotics, WAFR*, ser. Springer Proceedings in Advanced Robotics, vol. 25. Springer, 2022, pp. 349–367.
- [9] A. Lakshmanan, A. Gahlawat, and N. Hovakimyan, "Safe feedback motion planning: A contraction theory and  $l_1$ -adaptive control based approach," *CDC*, 2020.
- [10] B. T. Lopez, J. E. Slotine, and J. P. How, "Robust adaptive control barrier functions: An adaptive & data-driven approach to safety," *IEEE Control. Syst. Lett.*, vol. 5, no. 3, pp. 1031–1036, 2021.
- [11] Q. Pham, N. Tabareau, and J. E. Slotine, "A contraction theory approach to stochastic incremental stability," *IEEE Trans. Autom. Control.*, vol. 54, no. 4, pp. 816–820, 2009.
- [12] H. Tsukamoto, S. Chung, and 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.
- [13] D. Sun, S. Jha, and C. Fan, "Learning certified control using contraction metric," *CoRL*, 2020.
- [14] H. Tsukamoto and S. 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.
- [15] S. Singh, S. M. Richards, V. Sindhvani, J. E. Slotine, and M. Pavone, "Learning stabilizable nonlinear dynamics with contraction-based regularization," *IJRR*, 2020.
- [16] W. Han, A. Jasour, and B. C. Williams, "Non-gaussian risk bounded trajectory optimization for stochastic nonlinear systems in uncertain environments," in *2022 International Conference on Robotics and Automation, ICRA 2022, Philadelphia, PA, USA, May 23-27, 2022*. IEEE, 2022, pp. 11 044–11 050.
- [17] G. Chou, H. Wang, and D. Berenson, "Gaussian process constraint learning for scalable chance-constrained motion planning from demonstrations," *IEEE Robotics Autom. Lett.*, vol. 7, no. 2, pp. 3827–3834, 2022.
- [18] M. Ono, L. Blackmore, and B. C. Williams, "Chance constrained finite horizon optimal control with nonconvex constraints," in *ACC*, 2010.
- [19] G. Chou, D. Berenson, and N. Ozay, "Uncertainty-aware constraint learning for adaptive safe motion planning from demonstrations," in *4th Conference on Robot Learning, CoRL 2020, 16-18 November 2020, Virtual Event / Cambridge, MA, USA*, ser. Proceedings of Machine Learning Research, vol. 155. PMLR, 2020, pp. 1612–1639.
- [20] M. P. Vitus, Z. Zhou, and C. J. Tomlin, "Stochastic control with uncertain parameters via chance constrained control," *IEEE Transactions on Automatic Control*, vol. 61, no. 10, pp. 2892–2905, 2016.
- [21] B. Luders, M. Kothari, and J. How, "Chance constrained rrt for probabilistic robustness to environmental uncertainty," in *GNC*, 2010.
- [22] S. Patil, J. van den Berg, and R. Alterovitz, "Estimating probability of collision for safe motion planning under gaussian motion and sensing uncertainty," in *ICRA*, 2012.
- [23] Y. Meng, D. Sun, Z. Qiu, M. T. B. Waez, and C. Fan, "Learning density distribution of reachable states for autonomous systems," in *Conference on Robot Learning, 8-11 November 2021, London, UK*, ser. Proceedings of Machine Learning Research, vol. 164. PMLR, 2021, pp. 124–136.
- [24] A. Majumdar, R. Vasudevan, M. M. Tobenkin, and R. Tedrake, "Convex optimization of nonlinear feedback controllers via occupation measures," *Int. J. Robotics Res.*, vol. 33, no. 9, pp. 1209–1230, 2014.
- [25] G. Manek and J. Z. Kolter, "Learning stable deep dynamics models," in *NeurIPS*, 2019, pp. 11 126–11 134.
- [26] N. M. Boffi, S. Tu, N. Matni, J. E. Slotine, and V. Sindhvani, "Learning stability certificates from data," *CoRL*, 2020.
- [27] Y. Chang, N. Roohi, and S. Gao, "Neural lyapunov control," *NeurIPS*, 2019.
- [28] T. Koller, F. Berkenkamp, M. Turchetta, and A. Krause, "Learning-based model predictive control for safe exploration," in *CDC*, 2018.
- [29] A. K. Akametalu, J. F. Fisac, J. H. Gillula, S. Kaynama, M. N. Zeilinger, and C. J. Tomlin, "Reachability-based safe learning with gaussian processes," in *CDC*, 2014, pp. 1424–1431.
- [30] F. Berkenkamp, R. Moriconi, A. P. Schoellig, and A. Krause, "Safe learning of regions of attraction for uncertain, nonlinear systems with gaussian processes," in *CDC*, 2016, pp. 4661–4666.
- [31] D. D. Fan, A. Agha-mohammadi, and E. A. Theodorou, "Deep learning tubes for tube MPC," *RSS*, 2020.
- [32] K. Leung and I. R. Manchester, "Nonlinear stabilization via control contraction metrics: A pseudospectral approach for computing geodesics," in *2017 American Control Conference (ACC)*. IEEE, 2017, pp. 1284–1289.
- [33] M. DeGroot and M. Schervish, *Probability & Statistics*. Pearson, 2013.
- [34] B. T. Lopez and J. E. Slotine, "Contraction metrics in adaptive nonlinear control," *CoRR*, vol. abs/1912.13138, 2019.
- [35] M. Jordan and A. G. Dimakis, "Exactly computing the local lipschitz constant of relu networks," in *NeurIPS*, 2020.
- [36] M. Fazlyab, A. Robey, H. Hassani, M. Morari, and G. Pappas, "Efficient and accurate estimation of lipschitz constants for deep neural networks," in *NeurIPS*. Curran Associates, Inc., 2019, pp. 11 427–11 438.
- [37] Y. Sattar and S. Oymak, "Non-asymptotic and accurate learning of nonlinear dynamical systems," 2020.