

Planning with Learned Dynamics: Probabilistic Guarantees on Safety and Reachability via Lipschitz Constants

Craig Knuth^{*1}, Glen Chou^{*1}, Necmiye Ozay¹ and Dmitry Berenson¹

Abstract—We present a method for feedback motion planning of systems with unknown dynamics which provides probabilistic guarantees on safety, reachability, and goal stability. To find a domain in which a learned control-affine approximation of the true dynamics can be trusted, we estimate the Lipschitz constant of the difference between the true and learned dynamics, and ensure the estimate is valid with a given probability. Provided the system has at least as many controls as states, we also derive existence conditions for a one-step feedback law which can keep the real system within a small bound of a nominal trajectory planned with the learned dynamics. Our method imposes the feedback law existence as a constraint in a sampling-based planner, which returns a feedback policy around a nominal plan ensuring that, if the Lipschitz constant estimate is valid, the true system is safe during plan execution, reaches the goal, and is ultimately invariant in a small set about the goal. We demonstrate our approach by planning using learned models of a 6D quadrotor and a 7DOF Kuka arm. We show that a baseline which plans using the same learned dynamics *without* considering the error bound or the existence of the feedback law can fail to stabilize around the plan and become unsafe.

I. INTRODUCTION

Planning and control with guarantees on safety and reachability for systems with unknown dynamics has long been sought-after in the robotics and control community. Model-based optimal control can achieve this if the dynamics are precisely modeled, but modeling assumptions inevitably break down when applied to real physical systems due to unmodeled effects from friction, slip, flexing, etc. To account for this gap, data-driven machine learning methods and robust control seek to sidestep the need to precisely model the dynamics *a priori*. While robust control can provide strong guarantees when the unmodeled component of the dynamics is small and satisfies strong structural assumptions [1], such methods requires an accurate prior which may not be readily available. In contrast, machine learning methods are flexible but often lack formal guarantees, precluding their use in safety-critical applications. For instance, small perturbations from training data cause drastically poor and costly predictions in stock prices and power consumption [2]. Since even small perturbations from the training distribution can yield untrustworthy results, applying AI systems to predict dynamics can lead to unsafe, unpredictable behavior.

To address this gap, we propose a method for planning with learned dynamics which yields probabilistic guarantees on safety, reachability, and goal invariance in execution on the true system. Our core insight is that we can determine where a learned model can be trusted for planning using the Lipschitz constant of the error (the difference between the true and learned dynamics), which also informs how well the training data covers the task-relevant domain. Under the assumption of deterministic true dynamics, we can plan trajectories in this trusted domain with strong safety guarantees for an important class of learned dynamical systems.

Specifically, with a Lipschitz constant, we can bound the difference in dynamics between a novel point (that our model was not trained on) and a training point. Since the bound grows with the distance to training points, we can naturally define a domain where the model can be trusted as the set of points within a certain distance to training points. Conversely, to obtain a small bound over a desired domain, it is necessary to have good training data coverage in the task-relevant domain. At a high level, to obtain a small bound on the error in a domain, we want to have good coverage over the domain and regularity of the learned model via the Lipschitz constant of the error.

Our safety and reachability guarantees ultimately rely on an overestimate of the smallest Lipschitz constant. To find an estimate that exceeds the smallest Lipschitz constant with a given probability ρ , we use a statistical approach based on Extreme Value Theory [3] and validate its result with a Kolmogorov-Smirnov goodness-of-fit test [4]. If the test validates our estimate, we can choose a confidence interval [4] with an upper bound that overestimates the true Lipschitz constant with probability ρ . Our method requires the estimation of three Lipschitz constants, translating to system safety and reachability guarantees which hold with a probability of at least ρ^3 . This guarantee is fairly strong as it holds for all time, unlike many methods offering probabilistic guarantees on a per trajectory or episode basis [5] [6] [7].

If the learned dynamics have at least as many controls as states and are control-affine (note we do not assume the true dynamics are also control-affine), then we also determine conditions for the existence of a feedback controller that tightly tracks the planned trajectory in execution under the true dynamics. The tight tracking error bound yields favorable properties for our planner and controller: if we have a valid Lipschitz constant estimate for a sufficiently-accurate learned model, 1) we guarantee safety if no obstacle is within the tracking error of the trajectory, 2) we guarantee we can reach the goal within a small tolerance, and 3) if we

^{*}C. Knuth and G. Chou contributed equally to this work.

¹All authors affiliated with the University of Michigan, Ann Arbor, MI 48109, US, {cknuth, gchou, necmiye, dmitryb}@umich.edu
This work is funded by NSF grants IIS-1750489 and ECCS-1553873, ONR grants N00014-21-1-2118 and N00014-18-1-2501, and a National Defense Science and Engineering Graduate (NDSEG) fellowship.

can assert a feedback law that keeps the system at the goal exists, then the closed-loop system is guaranteed to remain in a small region around the goal. In this paper, we assume the learned dynamics are control-affine, deterministic, and have at least as many controls as states (such as a robotic arm under velocity control), the true dynamics are deterministic, and that independent samples of the true dynamics can be taken in the domain of interest. Our contributions are:

- 1) A method to bound error between two general dynamics functions in a domain by using a Lipschitz constant
- 2) A condition for uncertain control-affine systems that guarantees the existence of a feasible feedback law
- 3) A planner that probabilistically guarantees safety and closed-loop stability-like properties about the goal for learned dynamics with as many controls as states
- 4) Evaluation on a 7DOF Kuka arm and a 6D quadrotor

II. RELATED WORK

Prior work has used data coverage and Lipschitz constant regularization to ensure properties of a learned function. [8] shows that a linearization of a nonlinear state estimator generalizes with bounded error by estimating the maximum slope (related to the Lipschitz constant) of the error. [9] learns a control barrier function (CBF) from data, ensuring its validity through Lipschitz constant regularity and by checking the CBF conditions at a finite set of points. In contrast, we apply and extend these ideas to plan with learned dynamics, using the Lipschitz constant of the error dynamics to provide safety, reachability, and stability-like guarantees.

More broadly, our work is related to methods for planning and control of unknown dynamics with performance guarantees. A traditional approach is robust control, which assumes a good prior on the true dynamics and that unmodeled components are tightly bounded in some set [1]. Robust control has been applied in model predictive control [10] and Hamilton-Jacobi (HJ) reachability analysis [11]. While these methods have strong guarantees, the assumption that the unmodeled dynamics can be tightly bounded requires a accurate prior whereas our method does not, by actively keeping both planned and executed trajectories in a domain where the model can be trusted without *a priori* knowledge.

Other methods use Gaussian Processes (GPs) to estimate the mean and covariance of the dynamics, providing probabilistic bounds on safety and reachability. For example, [12] probabilistically bounds the reachable set of a fixed horizon trajectory. Similarly [5] explores the environment while ensuring (with some probability) safety via HJ reachability analysis. In many contexts, GPs are used to derive confidence bounds that can provide probabilistic safety guarantees [13] [6]. These methods can model dynamics with stochasticity, which our method cannot handle. However, the GP-based methods are incapable of long-horizon planning due to the unbounded growth of the covariance ellipse unless a known feedback controller exists. Our method does not require any prior controller, and we can plan trajectories of arbitrary length without unbounded growth of the reachable set.

Other work performs long-horizon planning with learned models without safety or reachability guarantees. [14] plans in learned latent spaces. [15] estimates the confidence that a controller can move between states to guide planning. [16] learns when a reduced-order model can be used. Unlike [14]-[16], our method provides safety and reachability guarantees.

Our safety guarantees rely on proving the existence of a stabilizing feedback controller in execution, like LQR-trees [17], funnel libraries [18], and LQG-MP [7]. Unlike these approaches, our method requires no *a priori* model and can prove a feedback law exists without structural assumptions on the true dynamics (e.g. that they are polynomial).

III. PRELIMINARIES

Let $f : \mathcal{X} \times \mathcal{U} \rightarrow \mathcal{X}$ be the true unknown discrete-time dynamics where \mathcal{X} is the state space and \mathcal{U} is the control space, which we assume are deterministic. We define $g : \mathcal{X} \times \mathcal{U} \rightarrow \mathcal{X}$ to be an approximation of the true dynamics that is control-affine and therefore can be written as follows

$$g(x, u) = g_0(x) + g_1(x)u. \quad (1)$$

In this paper, we represent the approximate dynamics with a neural network, though our method is agnostic to the structure of the model and how it is derived. Let $\mathcal{S} = \{(x_i, u_i, f(x_i, u_i))\}_{i=1}^N$ be the training data for g , and let $\mathcal{P} = \{(x_j, u_j, f(x_j, u_j))\}_{j=1}^M$ be another set of samples collected near \mathcal{S} that will be used to estimate the Lipschitz constant. We use $\bar{\cdot}$ to refer to data points from \mathcal{S} or \mathcal{P} . We place no assumption on how \mathcal{S} is obtained; any appropriate method (uniform sampling, perturbations from expert trajectories, etc.) may be employed, although we require independent and identically distributed (i.i.d.) samples for \mathcal{P} . A single state-control pair is written as (x, u) . With some abuse of notation, we write $(\bar{x}, \bar{u}) \in \mathcal{S}$ if $(\bar{x}, \bar{u}) = (x_i, u_i)$ for some $1 \leq i \leq N$ (similarly for \mathcal{P}).

A Lipschitz constant bounds how much outputs change with respect to a change in the inputs. For some function h , a Lipschitz constant over a domain \mathcal{Z} is any number L such that for all $z_1, z_2 \in \mathcal{Z}$

$$\|h(z_1) - h(z_2)\| \leq L\|z_1 - z_2\| \quad (2)$$

Norms $\|\cdot\|$ are always the 2-norm or induced 2-norm. We define L_{f-g} , L_{g_0} , and L_{g_1} as the smallest Lipschitz constants of the error $f - g$, g_0 , and g_1 . The input to $f - g$ is a state-control pair (x, u) and its output is a state. For g_0 , both the input and output are a state. For g_1 , its input is a state and its output is a $\dim(\mathcal{X}) \times \dim(\mathcal{U})$ matrix where $\dim(\cdot)$ is the dimension of the space. A ball $\mathcal{B}_r(x)$ of radius r about a point x is defined as the set $\{y \mid \|y - x\| < r\}$, also referred to as a r -ball about x . We suppose the state space \mathcal{X} is partitioned into safe $\mathcal{X}_{\text{safe}}$ and unsafe $\mathcal{X}_{\text{unsafe}}$ sets (e.g., the states in collision with an obstacle).

The method consists of two major components. First, we determine a trusted domain $D \subseteq \mathcal{X} \times \mathcal{U}$ and estimate the Lipschitz constants. Second, we use D to find a path to the goal satisfying our safety and reachability requirements.

Problem 1. Given a learned model g , unknown dynamics f , and datasets Ψ and \mathcal{S} , determine the trusted domain D where $\|f(x, u) - g(x, u)\| \leq \epsilon$, for some $\epsilon > 0$. Additionally determine the Lipschitz constants L_{f-g} , L_{g_0} , and L_{g_1} in D .

Problem 2. Given control-affine g , unknown f , start x_I , goal x_G , goal tolerance λ , D , L_{f-g} , L_{g_0} , L_{g_1} , and \mathcal{X}_{unsafe} , plan a trajectory (x_0, \dots, x_K) , (u_0, \dots, u_{K-1}) such that $x_0 = x_I$, $x_{k+1} = g(x_k, u_k)$, $K < \infty$, and $\|x_K - x_G\| \leq \lambda$. Additionally, under the true dynamics f , guarantee that closed loop execution does not enter \mathcal{X}_{unsafe} , converges to $\mathcal{B}_{\epsilon+\lambda}(x_G)$, and remains in $\mathcal{B}_{\epsilon+\lambda}(x_G)$ after reaching x_K .

IV. METHOD

Secs. IV-A - IV-B and IV-C - IV-D cover our approaches to Probs. 1 and 2, respectively. In Sec. IV-A, we show how L_{f-g} can establish a trusted domain and how L_{f-g} can be estimated in Sec. IV-B. In Sec. IV-C, we design a planner that ensures safety, that the system remains in the trusted domain, and that a feedback law maintaining minimal tracking error exists. We present the full algorithm in Sec. IV-D.

A. The trusted domain

For many systems, we are only interested in a task-relevant domain, and it is often impossible to collect data everywhere in state space, especially for high-dimensional systems. Hence, it is natural that our learned model is only accurate near training data. With a Lipschitz constant of the error, we can precisely define how accurate the learned dynamics are in a domain constructed from the training data. We note this derivation can also be done for systems without the control-affine assumption on the learned dynamics, and thus it can still be useful for determining where a broader class of learned models can be trusted. However, removing the control-affine structure makes controller synthesis much more difficult, and is the subject of future work.

Consider a single training point (\bar{x}, \bar{u}) and a novel point (x, u) . We derive a bound on the error between the true and estimated dynamics at (x, u) using the triangle inequality and Lipschitz constant of the error:

$$\begin{aligned} & \|f(x, u) - g(x, u)\| \\ &= \|f(x, u) - g(x, u) - f(\bar{x}, \bar{u}) + g(\bar{x}, \bar{u}) \\ &\quad + f(\bar{x}, \bar{u}) - g(\bar{x}, \bar{u})\| \\ &\leq L_{f-g}\|(x, u) - (\bar{x}, \bar{u})\| + \|f(\bar{x}, \bar{u}) - g(\bar{x}, \bar{u})\|. \end{aligned} \quad (3)$$

The above relation describes the error at a novel point, but we can also generalize to any domain D . Define b_T to be the dispersion [19] of $\mathcal{S} \cap D$ in D and define e_T to be the maximum training error of the learned model. Explicitly,

$$b_T \doteq \sup_{(x, u) \in D} \min_{(\bar{x}, \bar{u}) \in \mathcal{S} \cap D} \|(x, u) - (\bar{x}, \bar{u})\| \quad (4)$$

$$e_T \doteq \max_{(\bar{x}, \bar{u}) \in \mathcal{S} \cap D} \|f(\bar{x}, \bar{u}) - g(\bar{x}, \bar{u})\| \quad (5)$$

Then, we can uniformly bound the error across the entire set D to yield a simple and exact relation between f and g .

$$\epsilon \doteq L_{f-g}b_T + e_T \quad (6)$$




Fig. 1. An example with $f(x) = x'$, $\dim(\mathcal{X}) = 1$, and $\dim(\mathcal{U}) = 0$. True dynamics: yellow; learned linear dynamics: orange; \mathcal{S} : green crosses; Ψ : light blue crosses; domain D : interval $[-1, 2]$, bordered in black. Here, $b_T = 0.3633$ (purple) and $e_T = 0.1161$ (blue). The Lipschitz constant of the error is $L_{f-g} = 0.1919$, yielding $\epsilon = 0.1859$. We can use this bound to ensure the difference between the learned and true dynamics is no more than ϵ in D (shaded orange area). Note L_{f-g} can be larger outside of D .

$$\forall (x, u) \in D \quad f(x, u) = g(x, u) + \delta, \quad \|\delta\| \leq \epsilon \quad (7)$$

See Fig. 1 for an example of these quantities. For the remainder of the method, we select D to be the union of r -balls about a subset of the training data $\mathcal{S}_D \subset \mathcal{S}$:

$$D = \bigcup_{(\bar{x}, \bar{u}) \in \mathcal{S}_D} \mathcal{B}_r(\bar{x}, \bar{u}) \quad (8)$$

In the next section we discuss selection of \mathcal{S}_D , its role in estimating L_{f-g} , and how r is selected.

B. Estimating the Lipschitz constant

For (7) to hold over all of D , we require that L_{f-g} is a Lipschitz constant for the error. We use results from Extreme Value Theory to obtain an estimate \hat{L}_{f-g} that overestimates L_{f-g} , i.e. $\hat{L}_{f-g} \geq L_{f-g}$, with a user-defined probability ρ .

We build on [20]-[21], which find an estimate \hat{L}_h of the Lipschitz constant L_h for a function $h(z)$ over a domain \mathcal{Z} by estimating the location parameter γ of a three-parameter reverse Weibull distribution, which for a random variable W has the cumulative distribution function (CDF)

$$F_W(w) = \begin{cases} \exp(-(\frac{\gamma-w}{\alpha})^\beta), & \text{if } w < \gamma \\ 1, & \text{if } w \geq \gamma. \end{cases}$$

Here, the location parameter γ is the upper limit on the support of the distribution, and α and β are the scale and shape parameters, respectively. Consider the random variable described by the maximum slope taken over N_L pairs of i.i.d. samples $\{(z_1^i, z_2^i)\}_{i=1}^{N_L}$ from \mathcal{Z} , i.e. $s = \max_i \frac{\|h(z_1^i) - h(z_2^i)\|}{\|z_1^i - z_2^i\|}$. From the Fisher-Tippett-Gnedenko Theorem [3], s follows one of the Frechet, reverse Weibull, or Gumbel distributions in the limit as N_L approaches infinity. If s follows the reverse Weibull distribution, which we validate in our results using the Kolmogorov-Smirnov (KS) goodness-of-fit test [4] with a significance value of 0.05 (the same threshold used in [21]), then L_h is finite and equals γ . We estimate L_h using the location parameter $\hat{\gamma}$ of a reverse Weibull distribution fit via maximum likelihood to N_S samples of s . Finally, we compute a confidence interval $c = \Phi^{-1}(\rho)\xi$ on $\hat{\gamma}$. Here ξ is the standard error of the fit $\hat{\gamma}$, which correlates with

the quality of the fit, and $\Phi(\cdot)$ is the standard normal CDF [4]. We select the upper end of the confidence interval as our estimate $\hat{L}_h = \hat{\gamma} + c$, which overestimates L_h with probability ρ . Note that increasing ρ increases c , improving the safety probability at the cost of loosening \hat{L}_h , which can make planning more conservative. We also note that this probability is valid in the limit as N_L approaches infinity, due to the Fisher-Tippett-Gnedenko theorem making claims only on the asymptotic distribution. We summarize the estimation method in Alg. 1.

Algorithm 1: Lipschitz estimation for $h(z)$ over \mathcal{Z}

Input: N_S, N_L, ρ
1 for $j = 1, \dots, N_S$ **do**
2 | sample $\{(z_1^{i,j}, z_2^{i,j})\}_{i=1}^{N_L}$ uniformly in \mathcal{Z}
3 | compute $s_j = \max_i \|h(z_1^{i,j}) - h(z_2^{i,j})\| / \|z_1^{i,j} - z_2^{i,j}\|$
4 fit reverse Weibull to $\{s_j\}$ to obtain $\hat{\gamma}$ and standard error ξ
5 validate fit using KS test with significance level 0.05
6 if validated **return** $\hat{L}_h = \hat{\gamma} + \Phi^{-1}(\rho)\xi$ **else return** failure

We wish to choose D to be large enough for planning while also keeping L_{f-g} small. To achieve this, we use a filtering procedure to reduce the impact of outliers in \mathcal{S} . Let μ and σ be the mean and standard deviation of the error over \mathcal{S} . Then, let $\mathcal{S}_D = \{(\bar{x}, \bar{u}) \in \mathcal{S} \mid \|f(\bar{x}, \bar{u}) - g(\bar{x}, \bar{u})\| \leq \mu + a\sigma\}$ where a is a user-defined parameter. Then, we run Alg. 2 in order to grow D . This method works by proposing values of r , estimating L_{f-g} , and increasing r until $r > \epsilon$ or $L_{f-g} \geq 1$. Finding D with $r > \epsilon$ and $L_{f-g} < 1$ is useful for planning (described further in Sec. IV-C, see (10)). Note that, in Euclidean spaces, $r \geq b_T$. If no filtering is done, $r = b_T$, since no point in D is further than a distance r from \mathcal{S}_D and the furthest any point in D can lie from a point in \mathcal{S}_D is r ; however, filtering shrinks D and thus decreases the dispersion, making it possible that $r \geq b_T$. The parameter a should be chosen to balance the size of D against the magnitude of L_{f-g} , which we tune heuristically.

This filtering lets us exclude regions where our learned model is less accurate, yielding smaller e_T . Note that filtering does not affect the i.i.d. property of the samples needed for Alg. 1; it only applies a mask to the domain. We also note that Alg. 2 returns a minimum value for r , but a larger r can be chosen as long as L_{f-g} is estimated with Alg. 1. A larger r makes planning easier by expanding the trusted domain.

Algorithm 2: Selecting r and D

Input: $\mu, \sigma, a, \mathcal{S}_D, \Psi, \alpha > 0$
1 $r \leftarrow \mu + a\sigma$
2 while True **do**
3 | construct D using equation (8)
4 | estimate L_{f-g} using Alg. 1 and Ψ
5 | calculate ϵ using equation (6)
6 | **if** $L_{f-g} \geq 1$ **then return** failure
7 | **if** $r > \epsilon$ **then return** r and D
8 | **else** $r \leftarrow r + \alpha$ // α is a small constant

While we never explicitly address the assumption that




Fig. 2. Visualizing D (boundary in black), D_ϵ (yellow), and D^c (complement of D). Each point in D_ϵ is at least ϵ distance away from D^c . If the system is controlled to a point in D_ϵ from anywhere in D under the learned dynamics, then it remains in D under the true dynamics.

the true dynamics are deterministic, the estimated Lipschitz constant may be unbounded in the stochastic case, such as when two samples have the same inputs but different outputs due to noise, causing a division by 0 in line 3 of Alg. 1.

L_{g_0} and L_{g_1} may also be estimated with Alg. 1, which we employ in the results. Alternatively, [22] can give tight upper bounds on the Lipschitz constant of neural networks, though it could not scale to the networks used in our results. Other approaches [23] improve scalability at the cost of looser Lipschitz upper bounds, and will be examined in the future.

C. Planning

We want to plan a trajectory from start x_I to goal x_G using the learned dynamics while remaining in $\mathcal{X}_{\text{safe}}$ in execution. We constrain the system to stay inside D , as model accuracy may degrade outside of the trusted domain. We develop a planner similar to a kinodynamic RRT [24], growing a search tree \mathcal{T} by sampling controls that steer towards novel states until we reach the goal. If a path is found we can ensure the goal is reachable with safety guarantees.

1) *Staying inside D :* To remain inside the set D , we introduce another set $D_\epsilon := D \ominus \mathcal{B}_\epsilon(0)$, which is the Minkowski difference between D and a ball of radius ϵ . Every point in D_ϵ is at least a distance of ϵ from any point in the complement of D . Since the learned dynamics differs from the true dynamics by at most ϵ in D , controlling to a point in D_ϵ under the learned dynamics ensures the system remains within D under the true dynamics (see Fig. 2).

How do we determine if a query point (x, u) is inside of D_ϵ ? Since we define D to be a union of balls (8), it would suffice to find a subset of training points $\mathcal{W} \subset \mathcal{S}_D$ such that the union of r -balls about the training points completely covers an ϵ -ball about (x, u) . Explicitly,

$$\bigcup_{(\bar{x}, \bar{u}) \in \mathcal{W}} \mathcal{B}_r(\bar{x}, \bar{u}) \supset \mathcal{B}_\epsilon(x, u) \quad (9)$$

In general, checking (9) is difficult, but if $L_{f-g} < 1$ and

$$r > \frac{e_T}{1 - L_{f-g}}, \quad (10)$$

then only one training point within a distance $r - \epsilon$ is needed to ensure a query point is in D_ϵ (see Fig. 3). Note by Alg.




Fig. 3. Illustrating the advantage of $L_{f-g} < 1$ and r selected according to (10). An ϵ -ball about a query point is shown in black; r -balls about training data are shown in blue. **Left:** $L_{f-g} > 1$, therefore requiring many training points to cover an ϵ -ball about the query point. **Right:** $L_{f-g} < 1$ and r is selected according to (10). Under these conditions, only one training point within a $r - \epsilon$ distance ensures an ϵ -ball about the (x, u) is entirely in D , ensuring that the query point is in D_ϵ .

2 lines 6-7, either (10) is guaranteed or $L_{f-g} \geq 1$, in which we return failure.

Lemma 1. *If $L_{f-g} < 1$ and r is selected according to equation (10), then a point (x, u) is in D_ϵ if there exists $(\bar{x}, \bar{u}) \in S$ such that $\|(x, u) - (\bar{x}, \bar{u})\| \leq r - \epsilon$.*

Proof. To prove, note we can rearrange terms in equation (10) to get $r > L_{f-g}r + e_T \geq \epsilon$. If there exists $(\bar{x}, \bar{u}) \in S_D$ such that $\|(x, u) - (\bar{x}, \bar{u})\| \leq r - \epsilon$, then $\mathcal{B}_\epsilon(x, u) \subset \mathcal{B}_r(\bar{x}, \bar{u}) \subset D$ since no point in $\mathcal{B}_\epsilon(x, u)$ is further than r distance from (\bar{x}, \bar{u}) . Since $\mathcal{B}_\epsilon(x, u) \subset D$, (x, u) is at least ϵ distance from any point in D^c and therefore $(x, u) \in D_\epsilon$. \square

In order to ensure $L_{f-g} < 1$, since it is derived from the training data and learned model, we must train a learned model that is sufficiently accurate (i.e. low error on $\mathcal{S} \cup \Psi$). In our experiments, it was enough to minimize mean squared error over the training set to learn models with this property.

To ensure that the resulting trajectory remains in D_ϵ , we ensure that corresponding pairs of state and control lie in D_ϵ at each step. In growing the search tree \mathcal{T} , we break down this requirement into two separate checks, the first of which optimistically adds states to the search tree and the second that requires pairs of states and controls to lie in D_ϵ . To illustrate, suppose we sample a new configuration x_{new} and grow the tree from some x to x_{new} . At this point, when sampling a control u to steer from x to x_{new} we enforce that $(x, u) \in D_\epsilon$ (see line 11 in Alg. 3). However, how do we know the resulting state, $x' = g(x, u)$, will lie in D_ϵ ? Since x' is a state and not a state-control pair, the above question is not well defined. Instead, we perform an optimistic check in adding x' which requires that there exists some \hat{u} such that $(x', \hat{u}) \in D_\epsilon$ (see line 14 in Alg. 3). In turn, when growing the search tree from x' to some other sampled point x'_{new} we ensure that the pair of state and newly sampled control u' lies in D_ϵ , i.e. $(x', u') \in D_\epsilon$.

2) *One step feedback law:* To prevent drift in execution, we also seek to ensure the trajectory planned with RRT can be tracked with minimal error. One key requirement to guarantee a feedback law exists is that the system is sufficiently actuated under the learned dynamics. This requires




Fig. 4. The one-step feedback law: plan with the learned dynamics (dashed black); rollout with the true dynamics (blue); prediction with the learned dynamics using the feedback law (red). At each point, we use (12) to find a feedback control \tilde{u}_k so $x_{k+1} = g(x_k, \tilde{u}_k)$. We arrive within ϵ of the next state under the true dynamics. This repeats until we reach the goal.

that $\dim(\mathcal{U}) \geq \dim(\mathcal{X})$. The check for sufficient actuation is done on a per state basis and can be done as we grow \mathcal{T} . This feedback law ensures that, under the learned dynamics, we can return to a planned trajectory in exactly one step.

Suppose we are executing a trajectory (x_0, \dots, x_K) with corresponding control (u_0, \dots, u_{K-1}) planned with the learned dynamics, and the system is currently at x_{k-1} . Under the learned dynamics, the plan is to move to $x_k = g(x_{k-1}, u_{k-1})$, but, under the true dynamics, the system will end up at some $\tilde{x}_k = f(x_{k-1}, u_{k-1})$ which is no more than an ϵ distance from x_k . Our goal is to find an input \tilde{u}_k such that $x_{k+1} = g(\tilde{x}_k, \tilde{u}_k)$. If this one-step feedback law exists for all $1 \leq k \leq K - 1$, it ensures the executed trajectory stays within ϵ distance of the planned trajectory (see Fig. 4).

With Lipschitz constants L_{g_0} and L_{g_1} , we can bound how much the learned dynamics varies in the ϵ -ball about x_k .

$$\forall \tilde{x}_k \in \mathcal{B}_\epsilon(x_k) \quad g(\tilde{x}_k, u) = g_0(x_k) + \Delta_0 + (g_1(x_k) + \Delta_1)u \quad (11)$$

where $\|\Delta_0\| \leq L_{g_0}\epsilon$ and $\|\Delta_1\| \leq L_{g_1}\epsilon$. With (11), the existence of \tilde{u}_k is informed by a perturbed linear equation:

$$\begin{aligned} x_{k+1} &= g(\tilde{x}_k, \tilde{u}_k), \quad \tilde{x}_k \in \mathcal{B}_\epsilon(x_k) \\ \Rightarrow x_{k+1} &= g_0(x_k) + \Delta_0 + (g_1(x_k) + \Delta_1)\tilde{u}_k \\ \Rightarrow A\tilde{u}_k &= b \end{aligned} \quad (12)$$

with $A = g_1(x_k) + \Delta_1$ and $b = x_{k+1} - g_0(x_k) - \Delta_0$

Prior to execution, we seek to answer two questions: when does \tilde{u}_k exist and does \tilde{u}_k lie in the control space \mathcal{U} (for instance in the presence of box constraints)? Results from the literature [25] give a bound on the difference between the nominal solution u_k and perturbed solution \tilde{u}_k ,

$$\|u_k - \tilde{u}_k\| \leq \frac{\|g_1(x_k)^+\| (\|\Delta_1\| \|u_k\| + \|\Delta_0\|)}{1 - \|g_1(x_k)^+\| \|\Delta_1\|} \doteq u_{\text{pert}}, \quad (13)$$

where $g_1(x_k)^+$ is the pseudo-inverse of $g_1(x_k)$ (in general $g_1(x_k)$ is not square). We can use this bound to ensure that \tilde{u}_k is guaranteed to lie in \mathcal{U} by enforcing that $u_k + u_{\text{pert}} \mathbf{1}_\infty \subseteq \mathcal{U}$, where $\mathbf{1}_\infty$ is the unit infinity-norm ball. Furthermore, A may become singular if $1 - \|g_1(x_k)^+\| \|\Delta_1\| \leq 0$. In this case, \tilde{u}_k is not guaranteed to exist.

If \tilde{u}_k exists and satisfies the control constraints for all $1 \leq k \leq K-1$, then we ensure that the system will track the path up to an ϵ error under the one-step feedback law. In planning, we add the existence of a valid one step feedback law as a check when growing the search tree. Formally:

Theorem 1. *For trajectory (x_0, \dots, x_K) and (u_0, \dots, u_{K-1}) , if the solution to the perturbed linear equation (12), \tilde{u}_k , exists for all $k \in \{1, \dots, K-1\}$, then under the true dynamics $\|\tilde{x}_k - x_k\| \leq \epsilon$ for all k , given L_{f-g} , L_{g_0} , and L_{g_1} are each an overestimate of the true Lipschitz constant of $f-g$, g_0 , and g_1 , respectively.*

Proof. Proof by induction. For the induction step, assume $\|\tilde{x}_k - x_k\| \leq \epsilon$ for some k . Since $\tilde{x}_k \in \mathcal{B}_\epsilon(x_k)$, the perturbed linear equation (12) is valid. If a solution exists, then $x_{k+1} = g(\tilde{x}_k, \tilde{u}_k)$ and $\|f(\tilde{x}_k, \tilde{u}_k) - x_{k+1}\| \leq \epsilon$. This satisfies the induction step. For the base case, we have $g(x_0, u_0) = x_1$ and $\|f(x_0, u_0) - x_1\| \leq \epsilon$. Thus, for all k , $\|\tilde{x}_k - x_k\| \leq \epsilon$. \square

3) *Ensuring safety and invariance about the goal:* Since it is guaranteed by Thm. 1 that $\|\tilde{x} - x_k\| \leq \epsilon$, we check that $\mathcal{B}_\epsilon(x_k) \subset \mathcal{X}_{\text{safe}}$ for each x_k on the path to ensure safety.

The exact nature of this check depends on the system and definition of $\mathcal{X}_{\text{unsafe}}$. For example, in our experiments on quadrotor, the state includes the quadrotor’s position in \mathbb{R}^3 and $\mathcal{X}_{\text{unsafe}}$ is defined by unions of boxes in \mathbb{R}^3 . By defining a bounding sphere that completely contains the quadrotor, we can verify a path is safe via sphere-box intersection. With the Kuka arm, we randomly sample joint configurations in an ϵ -ball about states, transform the joint configurations via forward kinematics, and check collisions in workspace. While this method is not guaranteed to validate the entire ball around a state, in practice no collisions resulted from execution of plans. Another approach computes a free-space bubble [26] around a given state x and check if it contains $\mathcal{B}_\epsilon(x)$, however this is known to be conservative.

To stay near the goal after executing the trajectory, we use the same perturbed linear equation to ensure the existence of a one-step feedback law. Here, rather than checking the next state along the trajectory is reachable from the previous, we check that the final state is reachable from itself, i.e. x_K is reachable from x_K . Similar to the arguments above, we can repeatedly execute the feedback law to ensure the system remains in an $(\epsilon + \lambda)$ -ball about the goal. Formally, we have:

Theorem 2. *If the solution, denoted u_{st} , to the perturbed linear equation exists for $A = g_1(x_K) + \Delta_1$ and $b = x_K - g_0(x_K) - \Delta_0$ for all $x \in \mathcal{B}_\epsilon(x_K)$, then the closed loop system will remain in $\mathcal{B}_{\epsilon+\lambda}(x_G)$, given L_{f-g} , L_{g_0} , and L_{g_1} are each an overestimate of the true Lipschitz constant of $f-g$, g_0 , and g_1 , respectively.*

Proof. By Thm. 1, $\|\tilde{x}_K - x_K\| \leq \epsilon$. Thus, if the solution to the perturbed linear equation with $A = g_1(x_K) + \Delta_1$ and $b = x_K - g_0(x_K) - \Delta_0$ exists and is valid then $g(\tilde{x}_K, u_{\text{st}}) = x_K$ and $\|f(\tilde{x}_K, u_{\text{st}}) - x_K\| \leq \epsilon$. Since $\|x_K - x_G\| \leq \lambda$, the system remains in $\mathcal{B}_{\epsilon+\lambda}(x_K)$ by the triangle inequality. \square

To close, we note that the overall safety and invariance

probability of our method is ρ^3 , arising from our need to estimate three Lipschitz constants: L_{f-g} , L_{g_0} , and L_{g_1} . Given independent samples for overestimating each constant with probability ρ via Alg. 1, the overall correctness probability is the product of the correctness of each constant, i.e. ρ^3 .

D. Algorithm

We present our full method, **Learned Models in Trusted Domains (LMTD-RRT)**, in Alg. 3. In practice, we implemented `SampleState` and `SampleControl` in two different ways: uniform sampling and perturbations from training data. Sampling perturbations (up to a norm of $r - \epsilon$) does not exclude valid (x, u) pairs since all points in D_ϵ lie within $r - \epsilon$ from a training point, and, in cases where D_ϵ is a relatively small volume, can yield a faster search. However, it also biases samples near regions where training data is more dense. We define the set $\mathcal{S}_\mathcal{X} = \{\tilde{x} \mid \exists \tilde{u} \text{ s.t. } (\tilde{x}, \tilde{u}) \in \mathcal{S}_D\}$ to describe the optimistic check described in Sec. IV-C.1. NN finds the nearest neighbor and `OneStep` checks that a valid feedback exists as described in Sec. IV-C.2. `Model` evaluates the learned dynamics and `InCollision` checks if an ϵ -ball is in $\mathcal{X}_{\text{safe}}$ as described in Sec. IV-C.3.

Algorithm 3: LMTD-RRT

Input: $x_I, x_G, \mathcal{S}_\mathcal{X}, \mathcal{S}_D, r, \epsilon, \lambda, N_{\text{samples}}, \text{goal_bias}$

```

1  $\mathcal{T} \leftarrow \{x_I\}$ 
2 while True do
3   while  $\neg$ sampled do
4      $x_{\text{new}} \leftarrow \text{SampleState}(\text{goal\_bias})$ 
5     if  $\|x_{\text{new}} - \text{NN}(\mathcal{S}_\mathcal{X}, x_{\text{new}})\| \leq r - \epsilon$  then
6       sampled  $\leftarrow$  True
7    $x_{\text{near}} \leftarrow \text{NN}(\mathcal{T}, x_{\text{new}})$ 
8    $i \leftarrow 0, u_{\text{best}} \leftarrow \emptyset, x_{\text{best}} \leftarrow \emptyset, d \leftarrow \infty$ 
9   while  $i < N_{\text{samples}}$  do
10     $u \leftarrow \text{SampleControl}()$ 
11    if  $\|(x_{\text{near}}, u) - \text{NN}(\mathcal{S}_D, (x_{\text{near}}, u))\| \leq r - \epsilon$  then
12       $x_{\text{next}} \leftarrow \text{Model}(x_{\text{near}}, u)$ 
13      if OneStep( $x_{\text{near}}, u, x_{\text{next}}$ )  $\wedge$ 
14         $\|x_{\text{next}} - \text{NN}(\mathcal{S}_\mathcal{X}, x_{\text{next}})\| \leq r - \epsilon \wedge$ 
15         $\|x_{\text{next}} - x_G\| < d \wedge$ 
16         $\neg \text{InCollision}(x_{\text{next}}, \epsilon)$  then
17           $u_{\text{best}} \leftarrow u, x_{\text{best}} \leftarrow x_{\text{next}}$ 
18           $d \leftarrow \|x_{\text{next}} - x_G\|$ 
19     $i \leftarrow i + 1$ 
20   if  $u_{\text{best}}$  then
21      $\mathcal{T} \leftarrow \mathcal{T} \cup \{x_{\text{best}}\}$ 
22     if  $\|x_{\text{best}} - x_G\| \leq \lambda$  then
23       return ConstructPath( $\mathcal{T}, x_{\text{best}}$ )
```

Once a plan has been computed, it can be executed in closed-loop with Alg. 4. `ModelG0` and `ModelG1` evaluate g_0 and g_1 of the learned model. `SolveLE` solves the linear equation and `Dynamics` executes the true dynamics f .

V. RESULTS

We present results on 1) a 2D system to illustrate the need for remaining near the trusted domain, 2) a 6D quadrotor to show scaling to higher-dimensional systems, and 3) a 7DOF Kuka arm simulated in Mujoco [27] to show scaling to complex dynamics that are not available in closed form.

Algorithm 4: LMTD-Execute

Input: $\{x_k\}_{k=0}^K, \{u_k\}_{k=0}^{K-1}$

- 1 $\tilde{x}_0 \leftarrow x_0, k \leftarrow 0$
- 2 **for** $k = 1 \dots n - 1$ **do**
- 3 $b \leftarrow x_{k+1} - \text{ModelG0}(\tilde{x}_k), A \leftarrow \text{ModelG1}(\tilde{x}_k)$
- 4 $\tilde{u}_k \leftarrow \text{SolveLE}(A, b)$
- 5 $\tilde{x}_{k+1} \leftarrow \text{Dynamics}(\tilde{x}_k, \tilde{u}_k)$

Using $\rho = 0.975$, we plan with LMTD-RRT and rollout the plans in open-loop (no computation of \tilde{u}_k) and closed-loop (Alg. 4). We compare with a naïve kinodynamic RRT that skips the checks on lines 5, 11, 13-14 of Alg. 3 in both open and closed loop. See the video for experiment visualizations.

A. 2D Sinusoidal Model

To aid in visualization, we demonstrate LMTD-RRT on a 2D system with dynamics $f(x, u) = f_0(x) + f_1(x)u$:

$$f_0(x) = \begin{bmatrix} x \\ y \end{bmatrix} + \Delta T \begin{bmatrix} 3 \sin(0.3(x + 4.5)) \\ 3 \sin(0.3(y + 4.5)) \end{bmatrix} \begin{bmatrix} \sin(0.3(y + 4.5)) \\ \sin(0.3(x + 4.5)) \end{bmatrix}$$

$$f_1(x) = \Delta T \begin{bmatrix} 1 + 0.05 \cos(y) & 0 \\ 0 & 1 + 0.05 \sin(x) \end{bmatrix}$$

where $\Delta T = 0.2$. We are given 9000 training points $(x_i, u_i, f(x_i, u_i))$, where x_i is drawn uniformly from an ‘L’-shaped subset of \mathcal{X} (see Fig. 5) and u_i is drawn uniformly from $\mathcal{U} = [-1, 1]^2$. $g_0(x)$ and $g_1(x)$ are modeled with separate neural networks with one hidden layer of size 128 and 512, respectively. We select $a = 3$ in Alg. 2. 1000 more samples are used to estimate L_{f-g} via Alg. 1, which we validate with a KS test with a p value of 0.56, far above the 0.05 threshold significance value. We obtain $\hat{\gamma} = 0.117$ and $c = 6.85 \times 10^{-4}$, giving $\epsilon = 0.215$ over D .

See Fig. 5 for examples of the nominal, open-loop, and closed-loop trajectories planned with LMTD-RRT and a naïve kinodynamic RRT. The plan computed with LMTD-RRT remains in regions where we can trust the learned model (i.e. within D_ϵ) and the closed-loop execution of the trajectory converges to $\mathcal{B}_{\epsilon+\lambda}(x_G)$. In contrast, both the open-loop and closed-loop execution of the naïve RRT plan diverge. We provide statistics in Table I of maximum ℓ_2 tracking error $\max_{i \in \{1, \dots, T\}} \|\tilde{x}_i - x_i\|$ and final ℓ_2 distance to the goal $\|\tilde{x}_T - x_G\|_2$ for both the open loop (OL) and closed loop (CL) variants, averaged over 70 random start/goal states. To give the baseline an advantage, we fix the start/goal states and plan with naïve RRT using two different dynamics models: 1) the same learned dynamics model used in LMTD-RRT and 2) a learned dynamics model with the same hyperparameters trained on the full dataset (10^4 datapoints), and report the statistics on the minimum of the two errors. The worst case tracking error for the plan computed with LMTD-RRT was 0.199, which is within the guaranteed tracking error bound of $\epsilon = 0.215$, while despite the data advantage, plans computed with naïve RRT suffer from higher tracking error. Average planning times for LMTD-RRT and naïve RRT are 4.5 and 17 seconds, respectively. Overall, this suggests that planning




Fig. 5. 2D sinusoidal dynamics. The LMTD-RRT plan (magenta) stays in D and ensures a valid feedback law exists at each step. The plan can be tracked within ϵ under closed loop control (cyan). If feedback is not applied, the system drifts to the edge of the trusted domain, exits, and diverges (green). The naïve RRT plan (brown) does not consider D , and does not reach the goal under closed loop (grey) or open loop (red) control.

	LMTD-RRT	Naïve kino. RRT
Max. trck. err. (CL)	0.099 ± 0.036 (0.199)	8.746 ± 4.195 (15.21)
Goal error (CL)	0.039 ± 0.020 (0.113)	7.855 ± 3.851 (14.78)
Max. trck. err. (OL)	12.84 ± 4.444 (20.92)	10.39 ± 1.962 (15.31)
Goal error (OL)	12.40 ± 4.576 (20.92)	10.12 ± 1.762 (15.20)

TABLE I
SINUSOID ERRORS IN CLOSED LOOP (CL) AND OPEN LOOP (OL).
MEAN \pm STANDARD DEVIATION (WORST CASE).

with LMTD-RRT avoids regions where model error may lead to poor tracking, unlike planning with a naïve RRT.

B. 6D Quadrotor Model

We evaluate our method on 6-dimensional fully-actuated quadrotor dynamics [28] with state $x = [\chi, y, z, \phi, \theta, \psi]^T$, where $f(x, u) = f_0(x) + f_1(x)u$, $f_0(x) = x$ and $f_1(x) =$

$$\Delta T \begin{bmatrix} c_\theta c_\psi & -c_\phi s_\psi + c_\psi s_\phi s_\theta & s_\psi s_\phi + c_\phi c_\psi s_\theta & 0 & 0 & 0 \\ c_\theta s_\psi & c_\phi c_\psi + s_\phi s_\psi s_\theta & -c_\psi s_\phi + c_\phi s_\psi s_\theta & 0 & 0 & 0 \\ -s_\theta & c_\theta s_\phi & c_\phi c_\theta & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & s_\phi t_\theta & c_\phi t_\theta \\ 0 & 0 & 0 & 0 & c_\phi & -s_\phi \\ 0 & 0 & 0 & 0 & s_\phi c_\theta & c_\phi / c_\theta \end{bmatrix},$$

where $\Delta T = 0.1$ and $s(\cdot)$, $c(\cdot)$, and $t(\cdot)$ are short for $\sin(\cdot)$, $\cos(\cdot)$, and $\tan(\cdot)$ respectively. We are given 9×10^6 training data tuples $(x_i, u_i, f(x_i, u_i))$, where x_i, u_i are generated with Halton sampling over $[-1, 1]^3 \times [-\frac{\pi}{20}, \frac{\pi}{20}]^3$ and $[-1, 1]^6$, respectively (data is collected near hover). As $f_0(x)$ is a simple integrator term, we assume it is known and we set $g_0(x) = x$, while $g_1(x)$ is learned with a neural network with one hidden layer of size 4000. We select $a = 6$ in Alg. 2. We use 10^6 more samples in Alg. 1 to estimate L_{f-g} , and conduct a KS test resulting in a p -value of 0.43 \gg 0.05. We obtain $\hat{\gamma} = 0.205$, $c = 0.011$, and $\epsilon = 0.134$.




Fig. 6. Quadrotor tracking example. The trajectory planned with LMTD-RRT (magenta) is tracked in closed loop (blue) and reaches the goal. The open loop (green) also converges near the goal, but not as close as the closed loop. The naïve RRT produces a plan (brown) that leaves the trusted domain. Thus, both the open (red) and closed (light blue) loop rapidly diverge.




Fig. 7. **Left:** Quadrotor obstacle (red) avoidance. Example plans (green, blue, black), tracking error bound ϵ overlaid (light blue). Closed-loop trajectories remain in the tubes, converging to the goal without colliding. **Right:** Naïve RRT plan (pink) fails to be tracked (cyan) and collides (red dots).

	LMTD-RRT	Naïve kino. RRT
Max. trck. err. (CL)	0.003 ± 0.001 (0.008)	10.59 ± 16.75 (153.26)
Goal error (CL)	0.001 ± 0.001 (0.004)	8.247 ± 8.434 (46.150)
Max. trck. err. (OL)	0.020 ± 0.007 (0.040)	4.289 ± 2.340 (12.986)
Goal error (OL)	0.019 ± 0.008 (0.040)	3.265 ± 2.028 (11.670)

TABLE II

QUADROTOR ERRORS (NO OBSTACLES) IN CLOSED LOOP (CL) AND OPEN LOOP (OL). MEAN \pm STANDARD DEVIATION (WORST CASE).

See Fig. 6 for examples of the planned, open-loop, and closed-loop trajectories planned with LMTD-RRT and a naïve RRT. The trajectory planned with LMTD-RRT remains close to the training data, and the closed-loop system tracks the planned path with ϵ -accuracy converging to $\mathcal{B}_{\epsilon+\lambda}(x_G)$. We note that using the feedback controller to track trajectories planned with naïve RRT tends to worsen the tracking error, implying our learned model is highly inaccurate outside of the domain. We provide statistics in Table II for maximum tracking error and distance to goal, averaged over 100 random start/goal states. The worst case closed-loop tracking error for trajectories planned with LMTD-RRT is 0.011, again much smaller than ϵ . As with the 2D example, we give the baseline an advantage in computing tracking error statistics by reporting the minimum of the two errors when planning with 1) the same model used in LMTD-RRT and 2) a model trained on the full dataset (10^7 points). Despite the data advantage, the plans computed using naïve RRT have much higher tracking error. Average planning times for our unoptimized code are 100 sec. for LMTD-RRT and 15 min. for naïve RRT, suggesting that sampling focused near the training data can improve planning efficiency.

We also evaluate LMTD-RRT on an obstacle avoidance problem (Fig. 7). We perform collision checking as described in Sec. IV-C.3. As the tracking error tubes (of radius $\epsilon = 0.134$) centered around the nominal trajectories never intersect with any obstacles, we can guarantee that the system never collides in execution. Empirically, in running Alg. 3 over 500 random seeds to obtain different nominal paths, the closed-loop trajectory never collides. In contrast, the naïve RRT plan fails to be tracked and collides (Fig. 7, right).

C. 7DOF Kuka Arm in Mujoco

We evaluate our method on a 7DOF Kuka iiwa arm simulated in Mujoco [27] using a Kuka model from [29]. We

	LMTD-RRT	Naïve kino. RRT
Max. trck. err. (CL)	0.010 ± 0.010 (0.038)	0.090 ± 0.250 (1.265)
Goal error (CL)	0.004 ± 0.004 (0.018)	0.082 ± 0.251 (1.265)
Max. trck. err. (OL)	0.036 ± 0.041 (0.142)	0.076 ± 0.084 (0.282)
Goal error (OL)	0.035 ± 0.040 (0.140)	0.071 ± 0.075 (0.225)

TABLE III

7DOF ARM ERRORS (NO OBSTACLES) IN CLOSED LOOP (CL) AND OPEN LOOP (OL). MEAN \pm STANDARD DEVIATION (WORST CASE).

train two models using different datasets, one for evaluating tracking error without the presence of obstacles (Table III), and the other for obstacle avoidance. For both models, $g_0(x)$ is again set to be x while $g_1(x)$ is learned with a neural network with one hidden layer of size 4000. For the results in Table III, we are provided 2475 training data tuples, which are collected by recording continuous state-control trajectories from an expert and evaluating $f(x, u)$ on the trajectories and on random state-control perturbations locally around the trajectories. We select $a = 5$ in Alg. 2. 275 more samples are used in Alg. 1 to estimate L_{f-g} , validated with a KS test with a p value of $0.58 \gg 0.05$. We obtain $\hat{\gamma} = 0.087$ and $c = 0.001$, leading to $\epsilon = 0.111$. In Table III, we provide statistics on maximum tracking error and distance to goal under plans with LMTD-RRT and the naïve RRT baseline (with a model trained on the full dataset of 2750 points), averaged over 25 runs of each method. Notably, closed-loop tracking of plans found with LMTD-RRT have lowest error, with a worst case error much smaller than $\epsilon = 0.111$. Planning takes on average 1.552 and 0.167 sec. for LMTD-RRT and naïve RRT, respectively. We suspect the naïve RRT exploits poor dynamics outside of D , expediting planning.

For the obstacle avoidance example (Fig. 8), we are provided 15266 datapoints, which again take the form of continuous trajectories plus perturbations. We select $a = 8$ in Alg. 2, and use 1696 more points to estimate L_{f-g} using Alg. 1, which we validate with a KS test with a p value of $0.37 \gg 0.05$. We obtain $\hat{\gamma} = 0.156$ and $c = 0.010$, leading to $\epsilon = 0.111$. In planning, as described in Sec. IV-C.3, we perform collision checking by randomly sampling configurations in an ϵ -ball about each point along the trajectory. Though this collision checker is not guaranteed to detect collision, in running LMTD-RRT over 20 random seeds, we did not observe collisions in execution for any of the 20 plans, and the arm safely reaches the goal without collision. Over these trajectories, the worst case tracking error is 0.107, which remains within $\epsilon = 0.111$. One such plan computed by LMTD-RRT and the corresponding open-loop and closed-loop tracking trajectories, is shown in the top row of Fig. 8. The three trajectories nearly overlap exactly due to small tracking error. In contrast, the naïve RRT plan cannot be accurately tracked, even with closed-loop control, due to planning outside of the trusted domain, causing the executed trajectories to diverge and collide with the table.

VI. DISCUSSION AND CONCLUSION

We present a method to bound the difference between learned and true dynamics in a given domain and derive




Fig. 8. Planning to move a 7DOF arm from below to above a table. Trajectory-tracking time-lapse (time increases from left to right). Red (nominal), green (closed loop), blue (open loop). **Top:** LMTD-RRT (red, green, blue overlap due to tight tracking). **Bottom:** Naïve RRT (poor tracking causes collision).

conditions that guarantee a one-step feedback law exists. We combine these two properties to design a planner that can guarantee safety, goal reachability, and that the closed-loop system remains in a small region about the goal.

While the method presented has strong guarantees, it also has limitations which are interesting targets for future work. First, the true dynamics are assumed to be deterministic. Stochastic dynamics may be possible by estimating the Lipschitz constant of the mean dynamics while also appropriately modeling the noise. Second, the actuation requirement limits the systems that this method can be applied to. For systems with $\dim(\mathcal{U}) < \dim(\mathcal{X})$, it may be possible to construct a similar feedback law that guarantees the learned dynamics will lie within a tolerance of planned states which, in turn, could still give strong guarantees on safety and reachability.

REFERENCES

- [1] K. Zhou and J. C. Doyle, *Essentials of robust control*, 1998.
- [2] G. R. Mode and K. A. Hoque, “Adversarial examples in deep learning for multivariate time series regression,” 2020.
- [3] L. De Haan and A. Ferreira, *Extreme value theory: an introduction*. Springer Science & Business Media, 2007.
- [4] M. DeGroot and M. Schervish, *Probability & Statistics*. Pearson, 2013.
- [5] 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.
- [6] F. Berkenkamp, M. Turchetta, A. Schoellig, and A. Krause, “Safe model-based reinforcement learning with stability guarantees,” in *NeurIPS*, 2017, pp. 908–918.
- [7] J. Van Den Berg, P. Abbeel, and K. Goldberg, “Lqg-mp: Optimized path planning for robots with motion uncertainty and imperfect state information,” *IJRR*, vol. 30, no. 7, pp. 895–913, 2011.
- [8] S. Dean, N. Matni, B. Recht, and V. Ye, “Robust guarantees for perception-based control,” in *LADC*, 2020, pp. 350–360.
- [9] A. Robey, H. Hu, L. Lindemann, H. Zhang, D. V. Dimarogonas, S. Tu, and N. Matni, “Learning control barrier functions from expert demonstrations,” 2020.
- [10] A. Aswani, H. Gonzalez, S. S. Sastry, and C. Tomlin, “Provably safe and robust learning-based model predictive control,” *Automatica*, vol. 49, no. 5, pp. 1216–1226, 2013.
- [11] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin, “A time-dependent hamilton-jacobi formulation of reachable sets for continuous dynamic games,” *TAC*, vol. 50, no. 7, pp. 947–957, 2005.
- [12] T. Koller, F. Berkenkamp, M. Turchetta, and A. Krause, “Learning-based model predictive control for safe exploration,” in *CDC*, 2018.
- [13] 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.
- [14] B. Ichter and M. Pavone, “Robot motion planning in learned latent spaces,” *IEEE RA-L*, vol. 4, no. 3, pp. 2407–2414, 2019.
- [15] J. Guzzi, R. Chavez-Garcia, M. Nava, L. Gambardella, and A. Giusti, “Path planning with local motion estimations,” *RA-L*, vol. 5, no. 2, 2020.
- [16] D. McConachie, T. Power, P. Mitrano, and D. Berenson, “Learning when to trust a dynamics model for planning in reduced state spaces,” *IEEE RA-L*, vol. 5, no. 2, pp. 3540–3547, 2020.
- [17] R. Tedrake, “Lqr-trees: Feedback motion planning on sparse randomized trees,” *Robotics: Science and Systems V*, 2009.
- [18] A. Majumdar and R. Tedrake, “Funnel libraries for real-time robust feedback motion planning,” *IJRR*, vol. 36, no. 8, pp. 947–982, 2017.
- [19] S. LaValle, *Planning algorithms*. Cambridge university press, 2006.
- [20] G. R. Wood and B. P. Zhang, “Estimation of the lipschitz constant of a function,” *J. Glob. Optim.*, vol. 8, no. 1, pp. 91–103, 1996.
- [21] T.-W. Weng, H. Zhang, P.-Y. Chen, J. Yi, D. Su, Y. Gao, C.-J. Hsieh, and L. Daniel, “Evaluating the robustness of neural networks: An extreme value theory approach,” *ICLR*, 2018.
- [22] M. Fazlyab, A. Robey, H. Hassani, M. Morari, and G. Pappas, “Efficient and accurate estimation of lipschitz constants for deep neural networks,” in *NeurIPS*, 2019, pp. 11 427–11 438.
- [23] M. Jordan and A. G. Dimakis, “Exactly computing the local lipschitz constant of relu networks,” in *NeurIPS*, 2020.
- [24] S. M. LaValle and J. J. Kuffner Jr, “Randomized kinodynamic planning,” *IJRR*, vol. 20, no. 5, pp. 378–400, 2001.
- [25] P. Lötstedt, “Perturbation bounds for the linear least squares problem subject to linear inequality constraints,” *BIT Numerical Mathematics*, vol. 23, no. 4, pp. 500–519, 1983.
- [26] S. Quinlan, *Real-time modification of collision-free paths*. Stanford, 1994, no. 1537.
- [27] E. Todorov, T. Erez, and Y. Tassa, “Mujoco: A physics engine for model-based control,” in *IRROS*, 2012, pp. 5026–5033.
- [28] F. Sabatino, “Quadrotor control: modeling, nonlinear control design, and simulation,” 2015.
- [29] “Harvardagileroboticslab/gym-kuka-mujoco.” [Online]. Available: <https://github.com/HarvardAgileRoboticsLab/gym-kuka-mujoco>