Proceedings of Machine Learning Research vol 144:1–14, 2021

Contraction $\mathcal{L}_1$-Adaptive Control using Gaussian Processes

Aditya Gahlawat * † GAHLAWAT@ILLINOIS.EDU
Arun Lakshmanan* † LAKSHMA2@ILLINOIS.EDU
Lin Song * LINSONG2@ILLINOIS.EDU
Andrew Patterson * APPATTE2@ILLINOIS.EDU
Zhuohuan Wu * ZW24@ILLINOIS.EDU
Naira Hovakimyan * NHOVAKIM@ILLINOIS.EDU
Evangelos A. Theodorou ‡ EVANGELOS.THEODOROU@GATECH.EDU

Abstract

We present a control framework that enables safe *simultaneous* learning and control for systems subject to uncertainties. The two main constituents are contraction theory-based $\mathcal{L}_1$-adaptive ($\mathcal{CL}_1$) control and Bayesian learning in the form of Gaussian process (GP) regression. The $\mathcal{CL}_1$ controller ensures that control objectives are met while providing safety certificates. Furthermore, the controller incorporates any available data into GP models of uncertainties, which improves performance and enables the motion planner to achieve optimality *safely*. This way, the safe operation of the system is *always* guaranteed, even during the learning transients.

Keywords: Safe Learning, Planning, Adaptive Control, Gaussian Process Regression

1. Introduction

Machine learning (ML) algorithms are potent tools for producing complex and accurate models of uncertain systems. The accurate representations help model-based reinforcement learning (MBRL) algorithms achieve performance and optimality (Recht, 2019). However, model uncertainties can make the system unstable during learning transients, which can have serious consequences, especially for safety-critical systems (Knight, 2002). Control-theoretic approaches based on Lyapunov functions and control invariant sets can offer safety certificates (Perkins and Barto, 2002; Berkenkamp et al., 2017; Chow et al., 2018). For instance, control-theoretic notions like asymptotic stability (Khalil, 2014, Chapter 3) are useful to guarantee system behavior in the limit. However, it is equally important to quantify the system's behavior during the complete operation and not just in the limit. More importantly, for learning-based control, the question of how to quantify and guarantee the system's safety during learning-transients is crucial.

Statement of Contributions: We propose a learning-based control framework using robust adaptive control theory for nonlinear systems that allows *improvement of optimality and performance while simultaneously guaranteeing safety*. The framework is illustrated in Fig. 1. Let the *actual* (uncertain) dynamics of the system be given by $\dot{x} = F(x, u)$, where x and u are the system state and input, respectively. Additionally, let $\dot{x} = \bar{F}(x, u)$ denote the *known* (learned or nominal) model. Given any *realizable* desired state-input pair (x_d, u_d) with respect to the known dynamics, i.e., $\dot{x}_d = \bar{F}(x_d, u_d)$, we design the input u to ensure that the state x of the actual system $\dot{x} = F(x, u)$ follows x_d safely and with quantifiable performance.

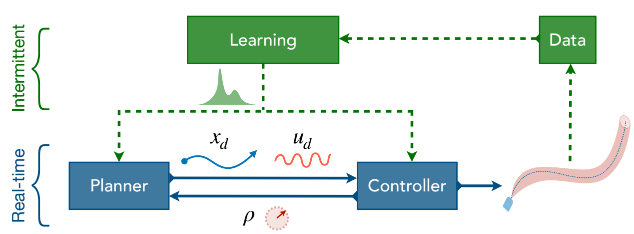

Figure 1: Illustration of our control framework.

[*] Mechanical Science and Engineering, University of Illinois at Urbana Champaign, Urbana, IL-61801
[†] Equal contribution
[‡] School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA-30332

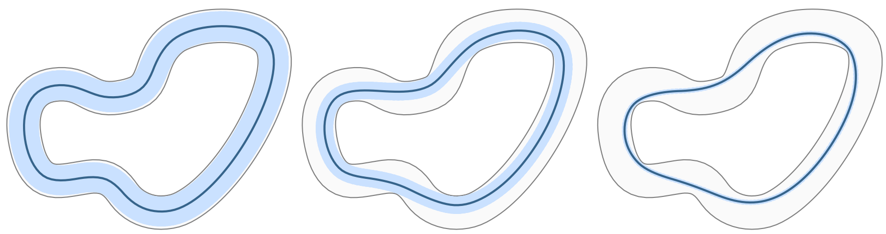

Figure 2: Consider a vehicle traversing a race track with some nominal model knowledge. Depending on the uncertainty and robustness requirements the control framework guarantees the safety bounds (blue tubes). As the learning improves, so does the performance and optimality.

Safety: We provide a systematic design of the feedback law $u = \pi(x, x_d, u_d)$, such that we can *apriori* compute a positive scalar ρ to ensure $\|x_d(t) - x(t)\| \leq \rho$, for all $t \in [0, T_f]$, with high-probability, where $T_f \leq \infty$ is the planning horizon. This guarantee implies the existence of a *tube* $\mathcal{O}_{x_d}(\rho)$ of radius ρ centered around the trajectory x_d, in which the actual state of the system x is guaranteed to lie in. The control design is based on our recent work in Lakshmanan et al. (2020). We then define the notion of *safety* as the existence of the apriori quantifiable tubes $\mathcal{O}_{x_d}(\rho)$. *Any* planning algorithm which produces (x_d, u_d) w.r.t. $\bar{F}$ by incorporating the additional constraint $\mathcal{O}_{x_d}(\rho) \notin \mathcal{X}_{\text{obs}}$ will thus ensure that the actual state $x \notin \mathcal{X}_{\text{obs}}$ (obstacle set). This minimal requirement enables the framework to be used in conjunction with many popular planning algorithms like Tassa et al. (2012); Williams et al. (2018); Wagener et al. (2019); LaValle and Kuffner Jr (2001); Cichella et al. (2017); Howell et al. (2019). Note the communication of ρ from the controller to the planner in Fig. 1. The safety is guaranteed regardless of the quality of learning. Thus, model learning can be performed whenever possible, and not at a high-rate (note the 'intermittent' and 'real-time' distinction in Fig. 1).

Performance: We define *performance* as the radius ρ of the tube since a smaller radius ρ implies better tracking performance and vice-versa. To improve performance, we rely on Bayesian learning (GP regression) to learn the model uncertainties. We use GP learning's predictive distribution to compute high-probability error bounds for the estimated uncertainties. These estimates are then incorporated within the feedback law $u = \pi(x, x_d, u_d)$ to handle the uncertainties as represented by the variance. The planner can thus operate by *only* incorporating the mean dynamics. The notion of performance also includes the desired robustness margins of the closed-loop system.

Optimality: The improved performance, as defined above, implies an improvement in optimality. Using improved performance given guaranteed safety (in the form of reduced radii tubes), an planner can produce the desired trajectory x_d that is optimal in the sense of the total path's length and the time taken to traverse the path. We refer to this as *performance-dependent optimality*. However, there is an additional notion of optimality w.r.t. to the learned models. MBRL algorithms rely on the known/learned model $\bar{F}$ to generate pairs (x_d, u_d) optimal for $\bar{F}$. Therefore, as learning improves, and thus $\bar{F} \to F$, the underlying MBRL should produce desired trajectories approaching optimality w.r.t. the actual dynamics. We refer to this as *model-based optimality*. Both performance-based and model-based optimality constitute the overall optimality. It is important to highlight that, as aforementioned, our control framework is planner agnostic: it *enables* the improvement of optimality via any planner capable of doing so, rather than guaranteeing it. The proposed framework provides the planner with improved performance guarantees and learned models; it is up to the planner to use these to improve optimality. The proposed framework enables the behavior in Fig. 2. Note that this figure does not show the distinction of the aforementioned optimality types but rather an improvement in the overall cost of the planned trajectory.

Related Work: Tracking controllers in the form of min-max MPC (Magni et al., 2001; Raimondo et al., 2009) handle uncertainties using worst-case bounds, rendering them overly conservative.

Tube-based MPC (Raković et al., 2016; Lopez et al., 2019) addresses this issue by using an ancillary controller to attenuate disturbances. Such methods typically place the limiting assumption on ancillary controllers' existence or instead have further assumptions on the dynamics like feedback linearizability and strict feedback forms. Another class of corrective methods that ensure safety is based on control barrier functions (CBFs) (Ames et al., 2019, 2016; Xu et al., 2015; Lopez et al., 2020). CBFs rely on specialized functions that ensure set invariance which prevent the system states from reaching unsafe regions. However, CBFs do not provide tracking error bounds with respect to a desired trajectory which is critical in evaluating the robot's performance. Moreover, these approaches require discovering a CBF and an ancillary controller which are non-trivial. In our approach, we provide an *explicit* design for the feedback controller with stability and performance guarantees. Learning-based tracking control (Aswani et al., 2013; Wabersich and Zeilinger, 2018; Soloperto et al., 2018; Rosolia and Borrelli, 2019) reduces conservatism by using measured data to improve models. However, several frameworks (Beckers et al., 2019; Capone and Hirche, 2019; Umlauft and Hirche, 2019; Helwa et al., 2019; Greeff and Schoellig, 2020) in this domain also require similar restrictions on the structure of the dynamics (e.g. strict feedback form, differential flatness, etc.) to ensure safety in the presence of uncertainties. The authors in Berkenkamp et al. (2017) use the regularity of the uncertainty and the sufficient statistics of the learned GP models to safely expand the region of attraction, and improve control performance using Lyapunov functions. In contrast, our control architecture actively compensates for the model uncertainties allowing the system states to reach any part of the operating region safely even when the quality of the learned model is poor. Probabilistic chance constraint methods, which use uncertainty propagation, have been shown to provide both asymptotic and transient bounds on the tracking performance (Koller et al., 2018). The implementations that rely on approximate uncertainty propagation offer excellent empirical performance without theoretical guarantees, shown in Hewing et al. (2019); Ostafew et al. (2016). However, uncertainty propagation methods sacrifice long-term accuracy for computational efficiency, for example by linearization, in order to be more tractable for real-time applications. Our proposed method avoids uncertainty propagation completely when considering nonlinear dynamics. Instead, we rely on uniform error bounds for GP predictions to apriori guarantee tracking performance with respect to the desired trajectory. This controller is capable of incorporating the learned dynamics while ensuring safety. This incorporation is based on both contraction theory (Manchester and Slotine, 2017; Singh et al., 2017; Lopez and Slotine, 2020) and the $\mathcal{L}_1$-adaptive control theory (Hovakimyan and Cao, 2010). Safe planning and control using $\mathcal{L}_1$-adaptive control theory can be found in Pereida and Schoellig (2018); Pravitra et al. (2020); Lakshmanan et al. (2020).

2. Problem Statement and Preliminaries

We consider the following uncertain *control-affine* nonlinear dynamics of the form

$$\dot{x}(t) = f(x(t)) + B(x(t))(u(t) + h(x(t))), \quad x(0) = x_0, \tag{1}$$

where $x(t) \in \mathbb{R}^n$, $u(t) \in \mathbb{R}^m$, for $t \in \mathbb{R}_{\geq 0}$, represent the system state and control input, respectively. The functions $f(x) \in \mathbb{R}^n$ and $B(x) \in \mathbb{R}^{n \times m}$ are the *known* components of the dynamics, whereas, $h(x) \in \mathbb{R}^m$ denotes the *model uncertainties*. The control-affine systems presented in (1) cover a wide range of physical control systems including, for e.g., nonlinear aircraft models (Garrard and Jordan, 1977), and quadrotor models (Singh et al., 2019; Mokhtari et al., 2006). Note that in (1), for the clarity of exposition, we place the assumption that the uncertainties are matched, i.e., $g(x) = B(x)h(x) \in \text{span}\{B\}$. The proposed method can be extended for uncertainties $\notin \text{span}\{B\}$ following the work laid out in Singh et al. (2019) and Manchester and Slotine (2018), with the expectation that the unmatched uncertainties cannot be fully compensated but instead only attenuated.

Definition 1 *We define the known and uncertain model parameter sets as* $\bar{\mathcal{M}} = \{f, B\}$ *and* $\hat{\mathcal{M}} = \{h\}$, *respectively. These sets induce the vector fields* $F(\bar{\mathcal{M}}; x, u) = f(x) + B(x)u$ *and*

$F(\mathcal{M}; x, u) = f(x) + B(x)(u + h)$, *where* $\mathcal{M} = \bar{\mathcal{M}} \cup \hat{\mathcal{M}}$, *which define the known and actual (uncertain) dynamics respectively as*[1]

$$\dot{x} = F(\bar{\mathcal{M}}; x, u) = f(x) + B(x)u, \quad \dot{x} = F(\mathcal{M}; x, u) = f(x) + B(x)(u + h(x)). \tag{2}$$

The following definition is placed for planning.

Definition 2 *Over a planning horizon* $[0, T_f]$, $0 < T_f \leq \infty$, $(x_d(t), u_d(t))$ *is a desired state-input pair if* $\dot{x}_d(t) = F(\bar{\mathcal{M}}; x_d(t), u_d(t))$ *and* $x_d(t) \in \mathcal{X}$, *for all* $t \in [0, T_f]$, *where* $\mathcal{X} \subset \mathbb{R}^n$ *is any compact convex set. Given any* $\rho > 0$, *we define*

$$\Omega(\rho, x_d(t)) := \{y \in \mathbb{R}^n \mid \|y - x_d(t)\| \leq \rho\}, \quad \mathcal{O}_{x_d}(\rho) = \cup_{t \in [0, T_f]} \Omega(\rho, x_d(t)). \tag{3}$$

We refer to $\mathcal{O}_{x_d}(\rho)$ *as the tube. Here* $\|\cdot\|$ *denotes the Euclidean norm.*

Since $F(\bar{\mathcal{M}}; x, u)$ is known, any model-based planner can generate the desired pair $(x_d(t), u_d(t))$ satisfying the state-constraints. The following ensures the generation of safe desired trajectories.

Assumption 1 *Given any tube width* $\rho > 0$ *and planning horizon* $[0, T_f]$, $0 < T_f \leq \infty$, *the planner produces a state-input pair* $(x_d(t), u_d(t))$ *(as in Def. 2) such that the induced tube* $\mathcal{O}_{x_d}(\rho)$ *satisfies* $\mathcal{O}_{x_d}(\rho) \in \mathcal{X}/\mathcal{X}_{obs}$, *for all* $t \in [0, T_f]$, *where* $\mathcal{X}_{obs}$ *represents the obstacles. The desired control input* $u_d(t)$ *satisfies* $\|u_d(t)\| \leq \Delta_{u_d}$, *for all* $t \in [0, T_f]$, *with the upper bound known.*

Problem Statement: Given the learned probabilistic estimates of the uncertainty $h(x)$, any desired state-input pair $(x_d(t), u_d(t))$, $t \in [0, T_f]$, designed by a planner using the known dynamics $F(\bar{\mathcal{M}}; x, u)$ (Defs. 1 and 2), and the desired robustness margins, the goal is to design the control input $u(t)$ that guarantees the existence of an apriori computable tube-width ρ so that the state of the uncertain dynamics in (1) ($F(\mathcal{M}; x, u)$ in Def. 1) satisfies $x(t) \in \Omega(\rho, x_d(t)) \subset \mathcal{O}_{x_d}(\rho)$ with high probability, for all $t \geq 0$, from all initial conditions $x_0 \in \mathcal{X}$, while satisfying the robustness requirements. Importantly, the existence of the pre-computable tubes should not depend on the quality of the learned estimates, thus ensuring that *safety remains decoupled from learning*. The learning should only affect the performance bounds and the optimality of the planned trajectory. We now place the following assumptions on the known and uncertain model parameters.

Assumption 2 *The functions* $f(x)$, $B(x)$, $h(x)$ *are continuous, bounded, and Lipschitz, for all* $x \in D \subset \mathbb{R}^n$, *where* D *is any compact set which can be arbitrarily large. Moreover, the matrix* $B(x)$ *has full column rank for all* $x \in D$, *thus guaranteeing the existence of the Moore-Penrose inverse* $B^{\dagger}(x) = \left(B^{\top}(x)B(x)\right)^{-1} B^{\top}(x)$.

We now discuss the constituent components of the proposed framework.

2.1. Control Contraction Metric Theory

Contraction theory analyzes the stability of trajectories of nonlinear system by considering the dynamics on the tangent space of the underlying manifold (Lohmiller and Slotine, 1998). This helps establish the notion of incremental exponential stability between any pair of trajectories. Analogous to a control Lyapunov function (CLF), controller synthesis based on contraction theory uses a *control contraction metric* (CCM) (Manchester and Slotine, 2017). In particular, a smooth function $M(x)$ is defined to be a CCM for the known dynamics $\dot{x} = F(\bar{\mathcal{M}}; x, u)$ in (2), if it satisfies the following conditions for all $(x, \delta_x) \in T\mathcal{X}$ (the tangent bundle of $\mathcal{X}$) (Manchester and Slotine, 2017):

1. We suppress the temporal dependencies for brevity.

$$\underline{\alpha}\mathbb{I}_n \preceq M(x) \preceq \bar{\alpha}\mathbb{I}_n, \quad \partial_{[b]_{\cdot,j}} M(x) + \left[M(x)\frac{\partial [b]_{\cdot,j}(x)}{\partial_x}\right]_{\mathbb{S}} = 0, \quad j \in \{1,\dots,m\}, \tag{4a}$$

$$\delta_x^\top M(x)B(x) = 0 \Rightarrow \delta_x^\top \left(\partial_f M(x) + \left[M(x)\frac{\partial f(x)}{\partial x}\right]_{\mathbb{S}} + 2\lambda M(x)\right)\delta_x \leq 0, \tag{4b}$$

for some scalars $\lambda > 0$, $0 < \underline{\alpha} < \bar{\alpha} < \infty$. Here $[b]_{\cdot,j}$ denotes the j^{th} column of $B(x)$, $\mathbb{I}_n$ is the identity matrix of dimension n, and $\partial_f M(x)$ denotes the directional derivative of $M(x)$ with respect to $f(x)$. The same holds for $\partial_{[b]_{\cdot,j}} M(x)$. Moreover, $[A]_{\mathbb{S}}$ denotes the symmetric part of the matrix A. Further details are presented in Manchester and Slotine (2017) and Lakshmanan et al. (2020). Crucially, note that the synthesis of the CCM $M(x)$ depends only on the known vector field $F(\bar{\mathcal{M}}; x, u)$ and can be computed offline.

Assumption 3 *The known vector field $F(\bar{\mathcal{M}}; x, u)$ in (2) admits a CCM $M(x)$, for all $x \in \mathcal{X}$, and for some positive constants λ, $\underline{\alpha}$, and $\bar{\alpha}$ as in (4).*

Under this assumption, one can construct a feedback law u_c, which renders the known dynamics $\dot{x} = F(\bar{\mathcal{M}}; x, u_c)$ incrementally exponential stable, see Singh et al. (2019) for details.

2.2. Bayesian Learning

The probabilistic estimates of the uncertainty $h(x)$ in (1) are learned using GP regression.

Assumption 4 *We assume that each of the elements $[h]_i(x)$, $i \in \{1,\dots,m\}$, are independent. Moreover, we assume that each element is a sample from a GP $[h]_i(x) \sim \mathcal{GP}(0, K_i(x,x'))$, where the kernel functions $K_i : \mathbb{R}^{n\times n} \to \mathbb{R}$ are known. Moreover, the kernels are twice-continuously differentiable with known constants L_{K_i}, $\nabla_x L_{K_i}$, such that $L_{K_i} = \max_{x,x'\in\mathcal{X}} \|\nabla_x K_i(x,x')\|$, and $\nabla_x L_{K_i} = \max_{x,x'\in\mathcal{X}} \left\|\nabla_x^2 K_i(x,x')\right\|$, for $i \in \{1,\dots,m\}$.*

The assumption is less conservative than requiring the uncertainty to be a member of the reproducing kernel Hilbert space (RKHS) associated with the kernel. For example, sample functions of GPs with squared-exponential (SE) kernels correspond to continuous functions, whereas the associated RKHS space contains only analytic functions (Van Der Vaart and Van Zanten, 2011). Moreover, the constants assumed to exist in Assumption 4 are easily computable, for example, for the squared-exponential (SE) kernel. However, it is important to note that the element-wise independence assumption on the uncertainty might be restrictive in certain scenarios and we are investigating relaxing this condition in future work.

Assume that we have N measurements of the form $y_k = h(x_k) + \kappa = B^\dagger(x_k)\left(\dot{x}_k - f(x_k)\right) - u_k + \kappa \in \mathbb{R}^m$, $k \in \{1,\dots,N\}$, where $\kappa \sim \mathcal{N}(0_m, \sigma^2\mathbb{I}_m)$ is the measurement noise and $0_m \in \mathbb{R}^m$ is a vector of zeros. We set up the data as $\mathbf{D} = \{\mathbf{Y},\mathbf{X}\}$, $\mathbf{Y} = \begin{bmatrix} y_1 & \cdots & y_N\end{bmatrix} \in \mathbb{R}^{m\times N}$, $\mathbf{X} = \begin{bmatrix} x_1 & \cdots & x_N\end{bmatrix} \in \mathbb{R}^{n\times N}$. Thus, for each of the constituent functions $[h]_i$, $i \in \{1,\dots,m\}$, we have the data as $\mathbf{D}_i = \{[\mathbf{Y}]_{i,\cdot}, \mathbf{X}\}$, where $[\mathbf{Y}]_{i,\cdot}$ denotes the i^{th} row of $\mathbf{Y}$. Conditioning the prior in Assumption 4 on the measured data $\mathbf{D}$, we obtain the posterior distribution at any test point $x^\star \in \mathcal{X}$ as (Williams and Rasmussen, 2006)

$$\mathbb{R} \ni [h]_i(x^\star) \sim \mathcal{N}\left(\nu_{i,N}(x^\star), \sigma_{i,N}^2(x^\star)\right), \quad i \in \{1,\dots,m\}, \tag{5}$$

with mean $\nu_{i,N}(x^\star) = K_i(x^\star,\mathbf{X})^\top \left[K_i(\mathbf{X},\mathbf{X}) + \sigma^2\mathbb{I}_N\right]^{-1} \left([\mathbf{Y}]_{i,\cdot}\right)^\top$, and variance $\sigma_{i,N}^2(x^\star) = K_i(x^\star,x^\star) - K_i(x^\star,\mathbf{X})^\top \left[K_i(\mathbf{X},\mathbf{X}) + \sigma^2\mathbb{I}_N\right]^{-1} K_i(x^\star,\mathbf{X})$. Using the linearity of the differential operator, we also compute the posterior distributions of the partial derivatives of $h(x)$ as

$$(\nabla_x [h]_i(x^\star))^\top \sim \mathcal{N}\left(\nabla_x \nu_{i,N}(x^\star)^\top, \nabla_x \sigma^2_{i,N}(x^\star)\right), \tag{6}$$

with mean $\nabla_x \nu_{i,N}(x^\star)^\top = (\nabla_x K_i(x^\star, \mathbf{X}))^\top \left[K_i(\mathbf{X},\mathbf{X}) + \sigma^2 \mathbb{I}_N\right]^{-1} ([\mathbf{Y}]_{i,\cdot})^\top \in \mathbb{R}^n$ and variance $\nabla_x \sigma^2_{i,N}(x^\star) = \nabla^2_{x,x'} K_i(x^\star, x^\star) - (\nabla_x K_i(x^\star, \mathbf{X}))^\top \left[K_i(\mathbf{X},\mathbf{X}) + \sigma^2 \mathbb{I}_N\right]^{-1} \nabla_x K_i(x^\star, \mathbf{X}) \in \mathbb{S}^n$.

3. $\mathcal{CL}_1$ Control with Gaussian Process Learning

We now present our control framework which brings together contraction theory-based $\mathcal{L}_1$-adaptive ($\mathcal{CL}_1$) control with Bayesian learning (GP regression). The detailed description of the framework is illustrated in Fig. 3. Given the posterior distribution of the uncertainty in (5), we may update the known and uncertain model parameter sets in (2), as $\bar{\mathcal{M}} = \{f + B\nu_N, B\}$, $\hat{M} = \{h - \nu_N\}$, $\mathcal{M} = \bar{\mathcal{M}} \cup \hat{\mathcal{M}}$ which induce the *learned representations* of the known and actual dynamics as

$$\dot{x} = F(\bar{\mathcal{M}}; x, u) = f(x) + B(x)\nu_N(x) + B(x)u, \tag{7a}$$
$$\dot{x} = F(\mathcal{M}; x, u) = f(x) + B(x)\nu_N(x) + B(x)(u + h(x) - \nu_N(x)). \tag{7b}$$

Note that this step simply entails adding and subtracting the mean ν_N in the control channel. The control design is driven by the philosophy that the input u compensates for the uncertainty $h(x) - \nu_N(x)$ as quantified by the variance of the posterior distribution in (5). The uncertainty is compensated so that the actual system $\dot{x} = F(\mathcal{M}; x, u)$ behaves like the known/learned $\dot{x} = F(\bar{\mathcal{M}}; x, u)$ within guaranteed tube bounds presented in Definition 2. Then, any underlying planner can generate the desired pair $(x_d.u_d)$ satisfying the deterministic and uncertainty-free dynamics $\dot{x}_d = F(\bar{\mathcal{M}}; x_d, u_d)$ safe in the knowledge that the state x of the actual uncertain system $\dot{x} = F(\mathcal{M}; x, u)$ will remain in the tube $\mathcal{O}_{x_d}(\rho)$ centered on x_d. Therefore the following sequential tasks need to be performed: i) quantification of the uncertainty $h - \nu_N$, and ii) design of the input u to compensate the quantified uncertainty.

To quantify the uncertainty, we use the posterior distributions in (5) and (6) to show that there exist bounds Δ_{Ξ_u}, such that for all $x \in \mathcal{X}$, with high probability,

$$\|\Xi_u(x)\| \leq \Delta_{\Xi_u}, \tag{8}$$

where $\Xi_u = \left\{h - \nu_N, \frac{\partial h - \nu_N}{\partial x}\right\}$. The bounds are presented in the following theorem, which is presented in a highly condensed form due to space considerations. The expanded version of the theorem and its proof can be found in the extended version at Gahlawat et al. (2020a, Thm. 3.1).

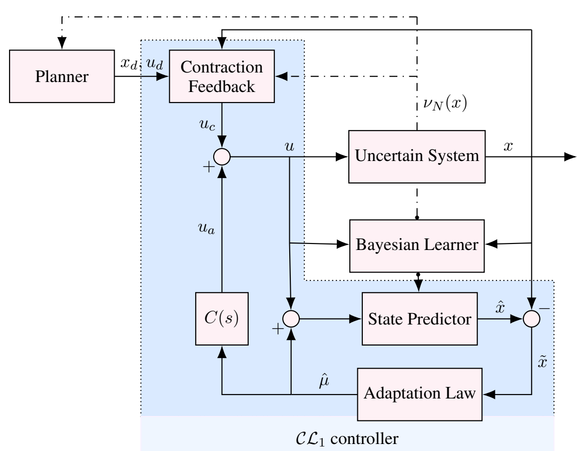

Figure 3: The control framework.

Theorem 3 *Let Assumptions 2 and 4 hold. Consider the posterior distributions in (5) and (6) and any scalars $\delta \in (0,1)$, $\tau > 0$. Then there exist computable functions $\Delta h(x,\tau)$ and $\nabla_x \Delta h(x,\tau)$, so that with $\Delta_{h-\nu_N} = \sup_{x\in\mathcal{X}} \Delta h(x,\tau)$ and $\Delta_{\frac{\partial h - \nu_N}{\partial x}} = \sup_{x\in\mathcal{X}} \nabla_x \Delta h(x,\tau)$, and the bounds in Eqn. (8) hold with probability at least $1-\delta$.*

With the bounds on the learned representation of the uncertainty $h(x) - \nu_N(x)$ established in Theorem 3, we proceed with the control design. To reiterate, for any (x_d, u_d) satisfying the deterministic learned dynamics $\dot{x}_d = F(\bar{\mathcal{M}}; x_d, u_d)$ in (7a), we design the input u so that the state of the actual dynamics $\dot{x} = F(\mathcal{M}; x, u)$ in (7b) satisfies $\|x - x_d\| \leq \rho \Rightarrow x \in \mathcal{O}_{x_d}(\rho)$ uniformly in time, for some tube-width $\rho > 0$. The control input, as in Fig. 3, is computed as

$$u(t) = u_c(\bar{\mathcal{M}}; t) + u_a(\bar{\mathcal{M}}; t), \tag{9}$$

where $u_c(\bar{\mathcal{M}}; t)$ is the contraction-theory based input and $u_a(\bar{\mathcal{M}}; t)$ is the $\mathcal{L}_1$ -adaptive input. By Assumption 3 in Sec. 2.1, there exists a CCM $M(x)$ for the original known dynamics $\dot{x} = F(\bar{\mathcal{M}}; x, u)$ (Eqn. (2)). More importantly, by Lopez and Slotine (2020, Lemma 1), $M(x)$ is a CCM for both the learned representations of the known and actual (uncertain) dynamics in (7). This allows us to seamlessly incorporate the learned mean ν_N and straightforwardly design the input $u_c(\bar{\mathcal{M}}, t)$ for the learned representation in (7a) using the CCM $M(x)$ as in Singh et al. (2019, Sec. 5.1). The design of the $\mathcal{L}_1$ -adaptive controller consists of a state-predictor, adaptation-law, and a low-pass filter $C(s)$ as illustrated in Fig. 3. Jointly, the $\mathcal{L}_1$ input $u_a(\bar{\mathcal{M}}; t)$ can be represented as

$$\dot{\hat{x}} = F(\bar{\mathcal{M}}; x, u + \hat{\mu}) + A_m \tilde{x}, \quad \dot{\hat{\mu}} = \Gamma \text{Proj}_{\hat{\mathcal{H}}}\left(\hat{\mu}, -B^\top(x) P \tilde{x}\right), \quad u_a(\bar{\mathcal{M}}; s) = -C(s)\hat{\mu}(s), \tag{10}$$

with $\hat{x}(0) = x_0$, $\hat{\mu}(0) \in \hat{\mathcal{H}}$, and s represents the Laplace variable. Here, $\hat{x}$ is the state of the predictor, $\tilde{x} = \hat{x} - x$, and $A_m \in \mathbb{R}^{n \times n}$ is an arbitrary Hurwitz matrix. The uncertainty estimate $\hat{\mu}$ is generated by the adaptation law for adaptation rate $\Gamma > 0$, with the projection operator $\text{Proj}_{\hat{\mathcal{H}}}(\cdot, \cdot)$, which ensures that $\hat{\mu} \in \hat{\mathcal{H}}$, $\hat{\mathcal{H}} = \{y \in \mathbb{R}^m \mid \|y\| \leq \Delta_{h-\nu_N}\}$, and $\mathbb{S}^n \ni P \succ 0$, which is the solution to the Lyapunov equation $A_m^\top P + P A_m = -Q$, for any $\mathbb{S}^n \ni Q \succ 0$. Finally, the low-pass filter $C(s)$ has a bandwidth ω and satisfies $C(0) = \mathbb{I}_m$. The high-level design idea is that the input u_a compensates for the uncertainty $h - \nu_N$ via the estimate $\hat{\mu}$ and within the bandwidth of $C(s)$.

Next, we analyze the uncertain system $\dot{x} = F(\mathcal{M}; x, u)$ in (7b) driven by the input (9). The complete details of the analysis can be found in the extended version of the manuscript in Gahlawat et al. (2020a). Given a desired trajectory x_d and arbitrarily chosen positive scalars ρ_a and ϵ, define

$$\rho_r = \sqrt{\frac{\bar{\alpha}}{\underline{\alpha}}}\, \|x_d(0) - x_0\| + \epsilon, \;\; \rho = \rho_r + \rho_a, \tag{11}$$

where $\bar{\alpha}, \underline{\alpha}$ are defined in (4a). Under Assumptions 1-3, we obtain conditions on the magnitude of the rate of adaptation Γ and the bandwidth ω of the low-pass filter $C(s)$ in (10) so that we are guaranteed stability and can quantify the performance. Once again, for clarity we choose not to present the complete definitions of the conditions and the reader is directed to Gahlawat et al. (2020a) for details. It is important to note that there always exists an adaptation rate Γ and a bandwidth ω that satisfy these conditions, see Lakshmanan et al. (2020) for further discussions. The following theorem establishes the performance of the closed-loop system and its proof can be found in the extended version at Gahlawat et al. (2020a, Thm. 3.2).

Theorem 4 *Let Assumptions 1-4 hold and let the bounds in Theorem 3 be computed for some $\delta \in (0, 1)$ and $\tau > 0$. Suppose the control input in (9) is designed so that the conditions on the rate of adaptation and filter bandwidth as given in Gahlawat et al. (2020a) are satisfied. Then, given any desired pair (x_d, u_d) satisfying the deterministic known dynamics $\dot{x}_d = F(\bar{\mathcal{M}}; x_d, u_d)$ in (7a), the state x of the actual (uncertain) dynamics $\dot{x} = F(\mathcal{M}; x, u)$ in (7b) driven by the input u from (9) satisfies with probability at least $1 - \delta$*

$$x(t) \in \Omega(\rho, x_d(t)) \subset \mathcal{O}_{x_d}(\rho), \quad \forall t \geq 0, \tag{12}$$

where ρ is defined in (11). Furthermore, the actual state x is uniformly ultimately bounded, with probability at least $1-\delta$, as

$$x(t) \subset \Omega(\hat{\delta}(\omega,T), x_d(t)) \subset \mathcal{O}_{x_d}(\rho), \quad \forall t \geq T > 0, \tag{13}$$

where the uniform ultimate bound (UUB) is defined as $\hat{\delta}(\omega,T) = \mu(\omega,T) + \rho_a$ with $\mu(\omega,T) = \sqrt{e^{-2\lambda T}\mathcal{E}(x_{d,0},x_0)/\underline{\alpha} + \zeta_1\left(\Xi_{\{u,k,c\}},\omega\right)}$.

Discussion: As the learning improves, the variance of the predictive Gaussian distribution collapses, and thus the constants in (8) decrease. Therefore, without changing the filter bandwidth ω and adaptation rate Γ, the UUB in Theorem 4 decreases. The decrease in the UUB, and the lack of a requirement for the re-tuning of the control parameters, is due to the monotonic dependence of the constants ζ_i on Δ_{Ξ_u} in (8). Furthermore, as aforementioned, the CCM $M(x)$ does not need to be re-synthesized as the model is updated using learning. Thus, without re-tuning the parameters of the control input (M, Γ, and ω), with the control designed using only Assumption 2, the performance improves as a function of learning. Of course the learning is not guaranteed to improve always, in which case, it will be reflected in the bounds Δ_{Ξ_u}. In this scenario, we are in no compulsion to incorporate the learned estimates, since the controller guarantees safety with the previously learned, or no, estimates. This is the exact reason that the proposed method does not require a high-rate, or any fixed rate, of model updates. Whenever it is provided with an improved model, it will be incorporated. The uniform bound in (12) is lower-bounded by the initialization error in (11). The size of the terms ϵ and ρ_a depends on the value of the adaptation rate Γ and filter-bandwidth ω. Thus, while in theory we can achieve the lowest-possible tube width, the size of Γ is limited by the available computation, and ω is limited by the desired robustness margins. Alternatively, instead of the uniform tube, a planner can use the UUB in (13), which induces tubes that exponentially collapse to a fixed radius dependent on ζ_1, a term that decreases as the learning improves. Compared to our initial work in Gahlawat et al. (2020b), the presented work is much more applicable to real-world problems. In particular, in Gahlawat et al. (2020b) we could only consider linear known systems and did not provide any theoretical guarantees. In the presented work, we are able to explicitly consider nonlinear systems because of bringing contraction theory within $\mathcal{L}_1$ control with theoretical guarantees. This further enables the use of learning for performance and optimality improvement with persistent safety as presented. Also note that since the adaptive control directly compensates for the uncertainty as quantified by the variance of the posterior distribution, any underlying planner need only incorporate the deterministic mean and not perform any uncertainty propagation, which is both approximate and computationally expensive. Finally, note the semi-global nature of the $\mathcal{L}_1$ augmentation. For a given CCM controller u_c that renders the known dynamics incrementally stable, the $\mathcal{L}_1$ augmentation can make any tube, no matter how large, forward invariant for the actual (uncertain) dynamics. The semi-global nature comes from the fact that, as is evident in this section, the control design explicitly depends on the size of the set/tube.

4. Simulation Results

We demonstrate our approach using two illustrative simulations. In the first example, we consider a modified Dubin's car system from Sun et al. (2020) and show how our control framework is used to ensure safety guarantees during the learning process, while outperforming a purely CCM-based approach. In the second example, we consider a planar quadrotor model from Singh et al. (2019) and show the usefulness of our control framework in feedback motion planning applications. The CCMs were discovered using DNNs from Sun et al. (2020) in the first example, and using the sum-of-square programming approach from Manchester and Slotine (2017) in the second example. In both scenarios, the dataset is generated by randomly sampling the state space, but one could also use more sophisticated exploration techniques to safely gather data based on our framework.

Dubin's Car: The vehicle dynamics are described by (p_x, p_y, θ, v), where p_x and p_y are positions, θ is the heading angle, and v is the velocity. The system has two control inputs that act on $(\dot{\theta}, \dot{v})$. The vehicle is tasked with traversing an obstacle forest from positions $(0, 0)$ to $(12, 0)$, and a desired trajectory is planned using the ALTRO solver presented in Howell et al. (2019) while minimizing an LQR objective. Let the system be randomly initialized around the origin and experience an unknown parasitic drag force given by $0.1v^2$. In Fig. 4a, a CCM-based feedback strategy is applied without concern for the uncertainty affecting the system. Out of the ten random initial conditions only two trajectories successfully reach the goal position, whereas in the majority of the simulations the vehicle collides with one of the obstacles before completing the task. With a conservative knowledge of the bounds on the uncertainty and its growth, a $\mathcal{CL}_1$ control is designed so that the system trajectories can be guaranteed to remain inside of a tube computed using the UUB in (13), as shown in Fig. 4b. As the uncertainty is learned following our approach, the bounds on the remainder uncertainty collapse with high probability as given in Theorem 3, Fig. 4c shows improved performance certificates in the form of the tightened tubes. Furthermore, the learned estimates are incorporated into the planner and $\mathcal{CL}_1$ architecture through the learned dynamics as (7). This example shows the clear improvement in the *performance-dependent optimality* and enables *model-based optimality* by incorporating $\bar{F}$ into the planner, while ensuring safety.

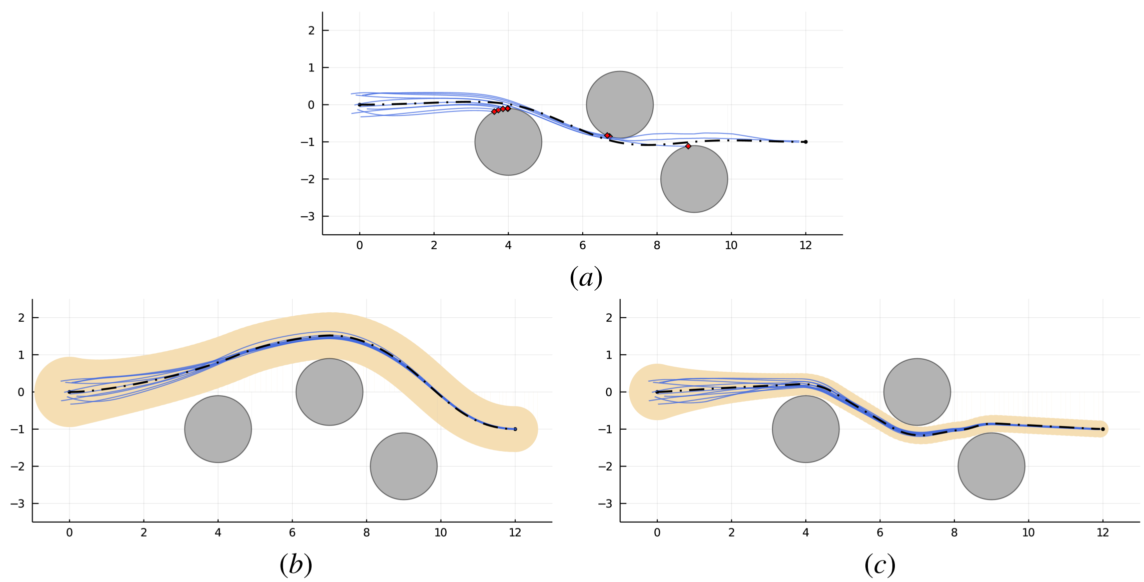

Figure 4: A Dubin's car traversing an obstacle forest using (a) CCM-based feedback (b) $\mathcal{CL}_1$ control, and (c) Our framework. The system trajectories are denoted in blue, the obstacles in gray, the planned trajectory in dashed black, and the tubes are shown as shaded orange regions. Any collisions with an obstacle terminate the trajectory immediately and are indicated as red diamonds.

Planar Quadrotor: The vehicle dynamics are described by $(p_x, p_y, v_x, v_y, \theta, \dot{\theta})$, where p_x and p_y are positions, v_x and v_y are velocities, and θ and $\dot{\theta}$ are the pitch angle and rate respectively. The system has two control inputs: the thrust and pitch moment commands that act on $(\dot{v}_y, \dot{\theta})$. The vehicle starts in a room at position $(0, 0)$ and is tasked with planning a trajectory that takes it to $(2, 0)$ safely. For such problems, complete or probabilistically complete planners are the algorithms of choice, since other methods typically get stuck at a local minimum and never reach the goal. We use the popular sampling-based planner BIT* (Gammell et al., 2015), where the two-point boundary value problem is solved using ALTRO. Similar to the previous example, consider that system experiences an unknown parasitic drag force given by $-0.1(v_x^2 + v_y^2)$ and a constant unknown offset in the thrust command. Planning without taking into account these uncertainties might generate trajectories that drive the system into regions that are unsafe. However in our framework, the uniform performance guarantees (12), provided by $\mathcal{CL}_1$ control, ensure that BIT* only samples states that lead to provably safe trajectories, Fig. 5a. Figures 5b and 5c show the tightening of the tubes as the uncertainty is learned batch-wise following our approach. This allows BIT* to construct a graph of safe trajectories with improved performance guarantees.

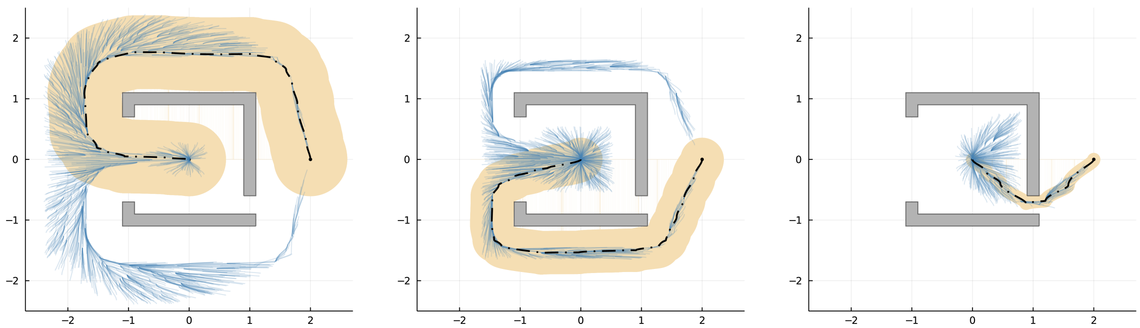

(*a*) (*b*) (*c*)

Figure 5: A planar quadrotor escaping a bug trap using (a) only a deterministic knowledge of the uncertainty, (b) model learned with N = 25 dataset, (c) model learned with N = 100 dataset. The blue lines indicate the edges of the random geometric graph constructed by BIT*. The dashed-black line indicates the lowest cost trajectory found by BIT*.

5. Conclusion

In this work, we presented a control framework, which enables safe simultaneous learning and control. The safety of the method is certified by the tracking error bounds produced by the ancillary contraction $\mathcal{L}_1$ controller. The learning is performed using Gaussian process regression. The learned Gaussian process model can be used to generate high probability uniform error bounds, which are incorporated into the controller to improve the tracking error bounds. Future work will extend the architecture to leverage the tracking error bounds in the path planning phase. The bounds are used to ensure safety, but can also be extended to provide worst case estimates for both the uncertainty reduction and cost associated with a desired trajectory. Finally, the guarantees will be extended to a larger class of nonlinear systems, explored in output feedback formulation, and other possible generalizations.

Acknowledgments

This work is financially supported by the National Aeronautics and Space Administration (NASA), Air Force Office of Scientific Research (AFOSR), National Science Foundation (NSF) Cyber Physical Systems (CPS) award # 1932529, and NSF National Robotics Initiative 2.0 (NRI-2.0) award # 1830639.

References

Aaron D Ames, Xiangru Xu, Jessy W Grizzle, and Paulo Tabuada. Control barrier function based quadratic programs for safety critical systems. *IEEE Transactions on Automatic Control*, 62(8): 3861–3876, 2016.

Aaron D Ames, Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. Control barrier functions: Theory and applications. In *Proceedings of 18th European Control Conference*, pages 3420–3431, Naples, Italy, 2019.

Anil Aswani, Humberto Gonzalez, S Shankar Sastry, and Claire Tomlin. Provably safe and robust learning-based model predictive control. *Automatica*, 49(5):1216–1226, 2013.

Thomas Beckers, Dana Kulić, and Sandra Hirche. Stable gaussian process based tracking control of euler–lagrange systems. *Automatica*, 103:390–397, 2019.

Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, and Andreas Krause. Safe model-based reinforcement learning with stability guarantees. In *Proceedings of 31st Conference on Neural Information Processing Systems*, pages 908–918, Long Beach, CA, USA, 2017.

Alexandre Capone and Sandra Hirche. Backstepping for partially unknown nonlinear systems using gaussian processes. *IEEE Control Systems Letters*, 3(2):416–421, 2019.

Yinlam Chow, Ofir Nachum, Edgar Duenez-Guzman, and Mohammad Ghavamzadeh. A Lyapunov-based approach to safe reinforcement learning. In *Proceedings of 32nd Conference on Neural Information Processing Systems*, pages 8092–8101, Quebec, Canada, 2018.

Venanzio Cichella, Isaac Kaminer, Claire Walton, and Naira Hovakimyan. Optimal motion planning for differentially flat systems using Bernstein approximation. *IEEE Control Systems Letters*, 2 (1):181–186, 2017.

Aditya Gahlawat, Arun Lakshmanan, Lin Song, Andrew Patterson, Zhuohuan Wu, Naira Hovakimyan, and Evangelos Theodorou. $\mathcal{RL}_1$-$\mathcal{GP}$: Safe simultaneous learning and control. *arXiv:2009.03864*, 2020a.

Aditya Gahlawat, Pan Zhao, Andrew Patterson, Naira Hovakimyan, and Evangelos Theodorou. $\mathcal{L}_1$-$\mathcal{GP}$: $\mathcal{L}_1$ adaptive control with Bayesian learning. In *Proceedings of 2nd Learning for Dyanmics & Control*, volume 120, pages 1–12, Online, 2020b.

Jonathan D Gammell, Siddhartha S Srinivasa, and Timothy D Barfoot. Batch informed trees (BIT*): Sampling-based optimal planning via the heuristically guided search of implicit random geometric graphs. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 3067–3074, Seattle, WA, USA, 2015.

William L Garrard and John M Jordan. Design of nonlinear automatic flight control systems. *Automatica*, 13(5):497–505, 1977.

Melissa Greeff and Angela P Schoellig. Exploiting differential flatness for robust learning-based tracking control using gaussian processes. *IEEE Control Systems Letters*, 5(4):1121–1126, 2020.

Mohamed K Helwa, Adam Heins, and Angela P Schoellig. Provably robust learning-based approach for high-accuracy tracking control of lagrangian systems. *IEEE Robotics and Automation Letters*, 4(2):1587–1594, 2019.

Lukas Hewing, Juraj Kabzan, and Melanie N Zeilinger. Cautious model predictive control using Gaussian process regression. *IEEE Transactions on Control Systems Technology*, 28(6):2736–2743, 2019.

Naira Hovakimyan and Chengyu Cao. *$\mathcal{L}_1$ Adaptive Control Theory: Guaranteed Robustness with Fast Adaptation*. SIAM, Philadelphia, PA, USA, 2010.

Taylor A Howell, Brian E Jackson, and Zachary Manchester. ALTRO: A fast solver for constrained trajectory optimization. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 7674–7679, Macau, China, 2019.

Hassan K Khalil. *Nonlinear control*. Pearson, London, UK, 2014.

John C Knight. Safety critical systems: Challenges and directions. In *Proceedings of the 24th International Conference on Software Engineering*, pages 547–550, Orlando, FL, USA, 2002.

Torsten Koller, Felix Berkenkamp, Matteo Turchetta, and Andreas Krause. Learning-based model predictive control for safe exploration. In *Proceedings of 57th IEEE Conference on Decision and Control*, pages 6059–6066, Miami Beach, FL, USA, 2018.

Arun Lakshmanan, Aditya Gahlawat, and Naira Hovakimyan. Safe feedback motion planning: A contraction theory and $\mathcal{L}_1$-adaptive control based approach. *arXiv:2004.01142*, 2020. Accepted for publication in IEEE Conference on Decision and Control 2020.

Steven M LaValle and James J Kuffner Jr. Randomized kinodynamic planning. *The International Journal of Robotics Research*, 20(5):378–400, 2001.

Winfried Lohmiller and Jean-Jacques E Slotine. On contraction analysis for non-linear systems. *Automatica*, 34(6):683–696, 1998.

Brett T Lopez and Jean-Jacques E Slotine. Adaptive nonlinear control with contraction metrics. *IEEE Control Systems Letters*, 5(1):205–210, 2020.

Brett T Lopez, Jonathan P Howl, and Jean-Jacques E Slotine. Dynamic tube MPC for nonlinear systems. In *Proceedings of American Control Conference*, pages 1655–1662, Philadelphia, PA, 2019.

Brett T Lopez, Jean-Jacques E Slotine, and Jonathan P How. Robust adaptive control barrier functions: An adaptive & data-driven approach to safety. *IEEE Control Systems Letters*, 5(3):1031–1036, 2020.

Lalo Magni, Henk Nijmeijer, and AJ Van Der Schaft. A receding–horizon approach to the nonlinear $\mathcal{H}_\infty$ control problem. *Automatica*, 37(3):429–435, 2001.

Ian R Manchester and Jean-Jacques E Slotine. Control contraction metrics: Convex and intrinsic criteria for nonlinear feedback design. *IEEE Transactions on Automatic Control*, 62(6):3046–3053, 2017.

Ian R Manchester and Jean-Jacques E Slotine. Robust control contraction metrics: A convex approach to nonlinear state-feedback $\mathcal{H}_\infty$ control. *IEEE Control Systems Letters*, 2(3):333–338, 2018.

Abdellah Mokhtari, Abdelaziz Benallegue, and Yuri Orlov. Exact linearization and sliding mode observer for a quadrotor unmanned aerial vehicle. *International Journal of Robotics and Automation*, 21(1):39–49, 2006.

Chris J Ostafew, Angela P Schoellig, and Timothy D Barfoot. Robust constrained learning-based NMPC enabling reliable mobile robot path tracking. *The International Journal of Robotics Research*, 35(13):1547–1563, 2016.

Karime Pereida and Angela P Schoellig. Adaptive model predictive control for high-accuracy trajectory tracking in changing conditions. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 7831–7837, Madrid, Spain, 2018.

Theodore J Perkins and Andrew G Barto. Lyapunov design for safe reinforcement learning. *Journal of Machine Learning Research*, 3(12):803–832, 2002.

Jintasit Pravitra, Kasey A Ackerman, Chengyu Cao, Naira Hovakimyan, and Evangelos A Theodorou. $\mathcal{L}_1$-adaptive MPPI architecture for robust and agile control of multirotors. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 7661–7666, Online, 2020.

Davide Martino Raimondo, Daniel Limon, Mircea Lazar, Lalo Magni, and Eduardo Ferná Camacho. Min-max model predictive control of nonlinear systems: A unifying overview on stability. *European Journal of Control*, 15(1):5–21, 2009.

Saša V Raković, William S Levine, and Behçet Açikmese. Elastic tube model predictive control. In *Proceedings of American Control Conference*, pages 3594–3599, Boston, MA, USA, 2016.

Benjamin Recht. A tour of reinforcement learning: The view from continuous control. *Annual Review of Control, Robotics, and Autonomous Systems*, 2:253–279, 2019.

Ugo Rosolia and Francesco Borrelli. Sample-based learning model predictive control for linear uncertain systems. In *Proceedings of 58th IEEE Conference on Decision and Control*, pages 2702–2707, Nice, France, 2019.

Sumeet Singh, Anirudha Majumdar, Jean-Jacques Slotine, and Marco Pavone. Robust online motion planning via contraction theory and convex optimization. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 5883–5890, Marina Bay Sands, Singapore, 2017.

Sumeet Singh, Benoit Landry, Anirudha Majumdar, Jean-Jacques Slotine, and Marco Pavone. Robust feedback motion planning via contraction theory. *The International Journal of Robotics Research*, 2019. Submitted.

Raffaele Soloperto, Matthias A Müller, Sebastian Trimpe, and Frank Allgöwer. Learning-based robust model predictive control with state-dependent uncertainty. *IFAC-PapersOnLine*, 51(20): 442–447, 2018.

Dawei Sun, Susmit Jha, and Chuchu Fan. Learning certified control using contraction metric. In *Conference on Robot Learning (CoRL)*, Virtual Conference, 2020.

Yuval Tassa, Tom Erez, and Emanuel Todorov. Synthesis and stabilization of complex behaviors through online trajectory optimization. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4906–4913, Vilamoura, Portugal, 2012.

Jonas Umlauft and Sandra Hirche. Feedback linearization based on gaussian processes with event-triggered online learning. *IEEE Transactions on Automatic Control*, 65(10):4154–4169, 2019.

Aad Van Der Vaart and Harry Van Zanten. Information rates of nonparametric Gaussian process methods. *Journal of Machine Learning Research*, 12(6), 2011.

Kim P Wabersich and Melanie N Zeilinger. Linear model predictive safety certification for learning-based control. In *Proceedings of 57th IEEE Conference on Decision and Control*, pages 7130–7135, Miami Beach, FL, USA, 2018.

Nolan Wagener, Ching an Cheng, Jacob Sacks, and Byron Boots. An online learning approach to model predictive control. In *Proceedings of Robotics: Science and Systems*, Freiburg im Breisgau, Germany, 2019.

Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*. MIT press, Cambridge, MA, USA, 2006.

Grady Williams, Paul Drews, Brian Goldfain, James M Rehg, and Evangelos A Theodorou. Information-theoretic model predictive control: Theory and applications to autonomous driving. *IEEE Transactions on Robotics*, 34(6):1603–1622, 2018.

Xiangru Xu, Paulo Tabuada, Jessy W Grizzle, and Aaron D Ames. Robustness of control barrier functions for safety critical control. *IFAC-PapersOnLine*, 48(27):54–61, 2015.