All Models Are Wrong: Robust Adaptive Model Predictive Control for Safe and High Accuracy Trajectory Tracking in the Presence of Model Errors

Karime Pereida and Angela P. Schoellig

***Abstract*— Robots deployed in unknown and dynamic environments require controllers that guarantee high performance and handle disturbances, unmodeled dynamics, and parametric uncertainties. Adaptive controllers enable systems to compensate for disturbances and unmodeled dynamics based on state measurements. In this paper, we propose a robust adaptive model predictive controller for safe and high accuracy trajectory tracking in the presence of model uncertainties. The proposed approach combines robust model predictive control (MPC) with an underlying discrete time ℓ_1 adaptive controller. The ℓ_1 adaptive controller forces the system to behave close to a linear reference model despite the presence of disturbances and unmodeled dynamics. The robust MPC uses the linear reference model to compute inputs that minimize tracking error. The true dynamics of the underlying adaptive system may deviate from the linear reference model. In this work we prove this deviation is bounded and account for the modeling error with the robust MPC. In experiments with a quadrotor we show that the proposed robust adaptive MPC enables the system to achieve high accuracy trajectory tracking and successfully complete challenging tasks (obstacle avoidance through aggressive maneuvers) that an adaptive MPC is not able to complete.**

I. INTRODUCTION

Robots are being deployed in challenging environments to carry out tasks with high overall performance. Operating in these environments requires control methods that guarantee high performance even in the presence of model uncertainties, unknown disturbances and changing dynamics. Robotic applications in these unstructured and changing environments include autonomous driving, assistive robotics, and unmanned aerial vehicle applications. In such robotic applications small changes in the environmental conditions may significantly deteriorate the performance and cause instability in traditional, model-based controllers (see [1] and [2]). Controllers for robots deployed in changing environments must adapt to and be robust against model uncertainties and unknown disturbances. Additionally, the controllers should have fast update rates to achieve better tracking performance.

In this paper, we extend previous work where we combined an underlying output feedback $\mathcal{L}_1$ adaptive controller with *nominal* model predictive control (MPC) [3]. The underlying $\mathcal{L}_1$ adaptive controller forces the system to behave close to a specified linear reference model, even in the presence of unknown disturbances. The MPC uses the linear reference model to compute an input that minimizes the tracking error of the $\mathcal{L}_1$ controlled system. This architecture enables us to achieve high tracking performance of a possibly nonlinear system at the computational cost of a linear MPC even when disturbances are applied; however, it assumes that

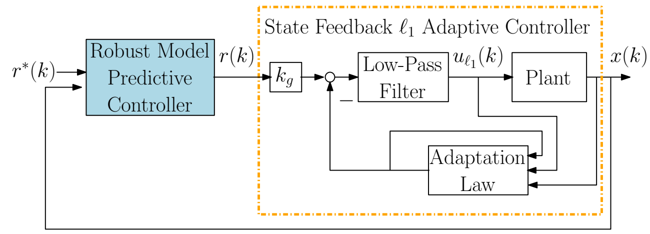

Fig. 1. The proposed robust adaptive model predictive control architecture. The underlying ℓ_1 adaptive controller forces the system to behave close to a specified linear model even in the presence of unknown disturbances. The ℓ_1 controlled system may deviate from the linear model, but this deviation is bounded. Using this bound, a robust model predictive controller calculates at each time step the input for the ℓ_1 controlled system that minimizes tracking error.

the linear model exactly represents the underlying adaptive system, which may not be entirely accurate.

In this work we propose a robust adaptive MPC that combines a discrete time state feedback ℓ_1 adaptive controller with *robust* MPC (see Fig. 1). Note that we use $\mathcal{L}_1$ to refer to the continuous time approach and ℓ_1 to refer to the discrete time approach. We formulate and provide performance guarantees for the ℓ_1 adaptive controller, which forces a system to behave close to a linear reference model. However, the ℓ_1 controlled system may deviate from the linear reference model. In this work we prove that this deviation has a uniform upper bound. Unlike [3], we use this upper bound and the linear reference model in a robust MPC module to compute an input that minimizes the tracking error of the ℓ_1 controlled system. Taking into account the modeling error enables the proposed robust adaptive MPC to successfully perform aggressive maneuvers where adaptive MPC would fail. We highlight theoretical contributions of the proposed approach and present experimental results on a quadrotor.

$\mathcal{L}_1$ adaptive control is based on the model reference adaptive control (MRAC) architecture but includes a low-pass filter that decouples robustness from adaptation [4]. It has successfully been applied to fixed-wing vehicles [5], [6], quadrotors [7], [8], a NASA AirSTAR flight test vehicle [9], a tailless unstable aircraft [10], and hexacopter and octocopter vehicles [11]. The $\mathcal{L}_1$ adaptive controller forces a system subject to uncertainties and disturbances, to behave as a specified linear model. Leveraging this characteristic, it was used in combination with iterative learning control (ILC), where ILC enabled the system to improve trajectory tracking over iterations even when external disturbances were applied to the system [8]. In [3], we proposed to replace ILC with

Sub version. Submitted to the *2019 IEEE International Conference on Robotics and Automation (ICRA).*
©2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

a MPC, which enabled the system to achieve high accuracy trajectory tracking on the first iteration by calculating the optimal input online. In this work, we improve the controller performance by adding robustness to the MPC to account for modeling errors.

Model predictive control solves a finite horizon optimal control problem at each time step to calculate a control sequence that minimizes a given objective function. Model uncertainties may significantly degrade the performance of MPC implementations. Instead of relying on the inherent robustness properties of standard MPC implementations, the work in [12] combines a parameter update mechanism with robust MPC algorithms. The optimization process accounts for parameter adaptation until it converges to the true system over time. Parameter uncertainty is decreased at every time step to reduce conservativeness of the algorithm. In [13], a robust adaptive MPC for a class of constrained linear, time-invariant systems is proposed. This approach proposed a novel method to estimate the parameters of a model suitable for MPC. Initially conservative, performance is improved over time as the parameter estimate's uncertainty is reduced and the estimate converges. However, the quality of the estimation depends on the excitation of the state. Moreover, learning-based MPC approaches have used neural networks (see [14], [15] and [16]) or Gaussian processes (see [17] and [18]) to learn the dynamics of the system used in the controller. These approaches require a significant amount of data in order to build an accurate model and often do not adapt to changes in the environment in real time and at high update rates. In these approaches, the model in the MPC is adapted, which makes it difficult to guarantee feasibility and stability theoretically [13]. In this work we tackle this problem by separating adaptation from MPC.

We combine a discrete time state feedback ℓ_1 adaptive controller with a robust MPC and make the following key contributions: (i) formulate a discrete time state feedback ℓ_1 adaptive controller, prove stability and provide performance guarantees even in the presence of disturbances; (ii) first to prove the upper bound on the difference between linear model and ℓ_1 controlled system; (iii) use the performance guarantees in a robust MPC framework to account for modeling errors; and (iv) validate the proposed approach through experiments on a quadrotor.

II. PROBLEM STATEMENT

The objective of this work is to achieve high accuracy trajectory tracking in the presence of uncertain, and changing conditions on the first iteration. Consider a system whose dynamics ('Plant' block in Fig. 1) are partially unknown and can be described by a single-input, single-output (SISO) system (this approach can later extended to multi-input, multi-output (MIMO) systems) for state feedback:

$$\begin{aligned} &x(k+1) = Ax(k) + b_m(u(k) + \theta^T x(k) + \xi)\,,\ x(0) = x_0 \\ &y(k) = c_m^T x(k)\,. \end{aligned} \tag{1}$$

where $x(k) \in \mathbb{R}^n$ is the system state vector (assumed to be measured); $u(k) \in \mathbb{R}$ is the control signal; b_m, $c_m \in \mathbb{R}^n$ are known constant vectors; A is the known $n \times n$ matrix, with $(A\,, b_m)$ controllable; θ is a vector of *unknown* parameters, and ξ is a disturbance both belonging to the compact convex sets $\theta \in \Theta \subset \mathbb{R}^n$ and $\xi \in \Xi \subset \mathbb{R}$; and $y(k) \in \mathbb{R}$ is the regulated output. We introduce an assumption on the sets Θ and Ξ.

Assumption 1. *The set Θ is compact and convex. Each dimension i is bounded by $L_{\Theta_{i,\min}} \le \theta_i \le L_{\Theta_{i,\max}}$ and the set Θ includes all the values that satisfy this criterion.*

The parameter ξ is bounded by $L_{\xi,\min} \le \xi \le L_{\xi,\max}$ where Ξ is the set that includes all the values that satisfy this criterion.

The system is required to accurately track a desired trajectory $r^*(k)$ defined over a finite number of steps $N < \infty$ and is assumed to be feasible with respect to the true dynamics of the ℓ_1-controlled system (orange dashed box in Fig. 1). The desired trajectory can be written as $\mathbf{r}^* = (r^*(1), \ldots, r^*(N))$, and the state of the plant as $\mathbf{x} = (x(1), \ldots, x(N))$. The goal is to take into account model uncertainties when minimizing the tracking performance criterion J defined as:

$$J \triangleq \min_{\mathbf{e}} \mathbf{e}^T \mathbf{Q} \mathbf{e}$$

where $\mathbf{e} = \mathbf{x} - \mathbf{r}^*$ is the tracking error and $\mathbf{Q}$ is a positive definite matrix.

III. METHODOLOGY

We propose a robust adaptive MPC to achieve safe and high accuracy trajectory tracking in the presence of modeling errors. We consider two subsystems: (1) a discrete time state feedback ℓ_1 adaptive controller (orange dashed box in Fig. 1) that makes the system behave close to a specified linear model, even in the presence of unknown disturbances; and (2) a robust MPC (blue box in Fig. 1) that uses the ℓ_1 linear model and modeling error bound in the calculation of the input that minimizes tracking error. The ℓ_1 adaptive controller provides transient and steady-state performance guarantees that specify how far the ℓ_1 controlled system can deviate from the specified linear model, given certain disturbance signals. The robust MPC improves tracking performance while accounting for the modeling error.

A. Discrete Time, State Feedback ℓ_1 Adaptive Control

The discrete time state feedback ℓ_1 adaptive controller is inspired by [19]; however, in this work we (i) extend the system to include a disturbance ξ (see (1)), (ii) modify the adaptation law to avoid division by zero, and (iii) prove the uniform upper bound on the difference between the ℓ_1 controlled system and the linear reference model.

The output feedback $\mathcal{L}_1$ adaptive controller used in [3] has performance guarantees that depend on the system transfer function $A(s)$, which is strictly-proper but *unknown*; hence, the bounds are not realizable. We propose a discrete time, state feedback ℓ_1 adaptive controller with bounds that are realizable. We formulate the controller in discrete time because (i) the robust MPC requires a discrete time system, and (ii) a discrete time formulation is more robust to changes in the sampling time than discretized $\mathcal{L}_1$ controllers [20]. The discrete time state feedback ℓ_1 adaptive controller aims to design a control input $u(k)$ such that the output $y(k)$ tracks a bounded reference input $r(k)$. Consider the following control structure for the system in (1):

$$u(k) = u_m(k) + u_{\ell_1}(k)\,, \quad u_m(k) = -k_m^T x(k)\,,$$

where $k_m \in \mathbb{R}^n$ renders $A_m \triangleq A - b_m k_m^T$ stable, while $u_{\ell_1}(k)$ is the adaptive component, which will be defined shortly. The static feedback gain k_m leads to the following partially closed-loop system:

$$\begin{aligned} &x(k+1) = A_m x(k) + b_m(u_{\ell_1}(k) + \theta^T x(k) + \xi)\,, \\ &x(0) = x_0\,, \quad y(k) = c_m^T x(k)\,. \end{aligned} \tag{2}$$

The equations that describe the implementation of the discrete time state feedback ℓ_1 adaptive controller are:

State Predictor: We use the following state predictor:

$$\begin{aligned} &\hat{x}(k+1) = A_m \hat{x}(k) + b_m \left(u_{\ell_1}(k) + \hat{\theta}^T(k)x(k) + \hat{\xi}(k)\right)\,, \\ &\hat{x}(0) = x_0\,,\ \hat{y}(k) = c_m^T \hat{x}(k)\,, \end{aligned} \tag{3}$$

where $\hat{x}(k) \in \mathbb{R}^n$ is the predicted state and $\hat{\theta}(k)$ and $\hat{\xi}(k)$ are the adaptive estimate of parameters θ and ξ.

Adaptation Law: We use a projection algorithm estimator [21], [22] that avoids division by zero:

$$\hat{\rho}(k+1) = \hat{\rho}(k) + \frac{x_{ext}(k)\left[b_0^T(\theta^T x(k) + \xi) - \hat{\rho}^T(k)x_{ext}(k)\right]}{1 + x^T(k)x(k)}\,, \tag{4}$$

where $x_{ext}(k) = [x^T(k), 1]^T$, $\hat{\rho}(k) = [\hat{\theta}^T(k), \hat{\xi}(k)]^T$, $\hat{\rho}(0) = \hat{\rho}_0 \in \Theta \times \Xi$, and $b_0 \triangleq \dfrac{b_m}{b_m^T b_m}$ is a constant vector. State measurements and the linear model (2) are used to calculate $\theta^T x(k) + \xi = x(k+1) - A_m x(k) - b_m u_{\ell_1}(k)$.

If $\hat{\rho}(k+1)$ lies outside $\Theta \times \Xi$, then we need to orthogonally project $\hat{\rho}(k+1)$ on the boundary of $\Theta \times \Xi$ [22]. The latter guarantees that the estimate $\hat{\rho}(k)$ remains in the set $\Theta \times \Xi$, which is needed for the performance guarantees.

Control law: The z-transform of the control law is:

$$u(z) = -C(z)(\hat{\eta}(z) - k_g r(z))\,, \tag{5}$$

where $r(z)$ and $\hat{\eta}(z)$ are the z-transforms of command input $r(k)$ and $\hat{\eta}(k) = \hat{\rho}^T(k)x_{ext}(k)$, respectively, $k_g \triangleq (c_m^T(I_n - A_m)^{-1}b_m)^{-1}$, and $C(z)$ is a bounded-input, bounded-output (BIBO) stable, strictly-proper, discrete time transfer function with gain $C(1) = 1$, and its state-space realization assumes zero initialization.

The discrete time state feedback ℓ_1 adaptive controller is defined via (3) – (5) with $C(z)$ satisfying the following ℓ_1-norm condition:

$$\lambda_\theta \triangleq \|G(z)\|_{\ell_1} L_\theta < 1\,, \quad \lambda_\xi \triangleq \|G(z)\|_{\ell_1} L_\xi < \infty\,, \tag{6}$$

where

$$\begin{aligned} &G(z) \triangleq H(z)(1 - C(z))\,, \quad &&H(z) \triangleq (zI_n - A_m)^{-1} b_m\,, \\ &L_\theta \triangleq \max_{\theta \in \Theta} \|\theta\|_1\,, &&L_\xi \triangleq \max_{\xi \in \Xi} \|\xi\|_1\,. \end{aligned} \tag{7}$$

Consider the following nonadaptive version of the system (2) controlled by (5) which defines the following *closed-loop reference system*:

$$\begin{aligned} x_{ref}(k+1) =& A_m x_{ref}(k) \\ &+ b_m(u_{ref}(k) + \theta^T x_{ref}(k) + \xi)\,, \\ x_{ref}(0) =& x_0\,, \\ u_{ref}(z) =& -C(z)(\theta^T x_{ref}(z) + \xi - k_g r(z))\,, \\ y_{ref}(k) =& c_m^T x_{ref}(k)\,. \end{aligned} \tag{8}$$

To show stability of the ℓ_1 controlled system, we first show that the closed-loop reference system is stable. Then we show the ℓ_1 controlled system stays close to the closed-loop reference system.

Lemma 1. *If* $\|G(z)\|_{\ell_1} L_\theta < 1$*, and* $\|G(z)\|_{\ell_1} L_\xi < \infty$*, then* (8) *is bounded-input, bounded-state (BIBS) stable [23] with respect to* $r(z)$ *and* x_0*.*

Proof. From the definition in (8), it follows that:

$$\begin{aligned} x_{ref}(z) = \;& H(z)k_g C(z) r(z) + G(z)\theta^T x_{ref}(z) + G(z)\xi \\ & + x_{in}(z) \end{aligned}$$

where $x_{in}(z) \triangleq (zI_n - A_m)^{-1} x_0$. Since $H(z)$, $C(z)$ and $G(z)$ are proper BIBO-stable discrete-time transfer functions, it follows from (8) that for all $i \in \mathbb{N} \cup \{0\}$ the following bound holds:

$$\begin{aligned} \|x_{ref}|_i\|_{\ell_\infty} \;\leq\; & \|H(z)k_g C(z)\|_{\ell_1} \|r|_i\|_{\ell_\infty} + \|x_{in}|_i\|_{\ell_\infty} \\ & + \|G(z)\theta^T\|_{\ell_1} \|x_{ref}|_i\|_{\ell_\infty} + \|G(z)\xi\|_{\ell_1}\,. \end{aligned}$$

Since A_m is stable, $x_{in}(k)$ is uniformly bounded. Then, we have the following relationship

$$\begin{aligned} \|G(z)\theta^T\|_{\ell_1} = \;& \max_{m=1,\ldots,n} \|G_m(z)\|_{\ell_1} \textstyle\sum_{p=1}^n |\theta_p| \\ & \leq \|G(z)\|_{\ell_1} L_\theta < 1\,, \end{aligned} \tag{9}$$

where $\|G_m(z)\|_{\ell_1}$ is the ℓ_1-norm for the impulse response for each output $g_m(k)$ where $m = 1, \ldots, n$. Recall from Assumption 1 that $\xi \in \Xi$, which leads to

$$\|G(z)\xi\|_{\ell_1} = \max_{m=1,\ldots,n} \|G_m(z)\|_{\ell_1} |\xi| \leq \|G(z)\|_{\ell_1} L_\xi < \infty\,. \tag{10}$$

Consequently,

$$\|x_{ref}|_i\|_{\ell_\infty} \leq \frac{\alpha}{1 - \|G(z)\theta^T\|_{\ell_1}}\,. \tag{11}$$

where $\alpha \triangleq \|H(z)k_g C(z)\|_{\ell_1} \|r|_i\|_{\ell_\infty} + \|G(z)\|_{\ell_1} L_\xi + \|x_{in}|_i\|_{\ell_\infty}$. Since $r(k)$ and $x_{in}(k)$ are uniformly bounded, $\|G(z)\|_{\ell_1} L_\xi$ is bounded, and (11) holds uniformly for all $i \in \mathbb{N} \cup \{0\}$, $x_{ref}(k)$ is uniformly bounded. Boundedness of $y_{ref}(k)$ follows from its definition. □

The closed-loop reference system is BIBS stable. It remains to show that the ℓ_1 controlled system stays close to the closed-loop reference system. To do this we start by defining the predictor error dynamics as $\tilde{x}(k+1) = A_m \tilde{x}(k) + b_m(\tilde{\theta}^T(k)x(k) + \tilde{\xi}(k))$, $\tilde{x}(0) = 0$, where $\tilde{x}(k) \triangleq \hat{x}(k) - x(k)$. Following a procedure similar to the proof of Lemma 4 in [19], it can be shown that the prediction error is uniformly bounded by

$$\|\tilde{x}\|_{\ell_\infty} \le \sqrt{\frac{e^{w\rho_{\max}}-1}{\mu\lambda_{\min}(P)}}\,, \tag{12}$$

where e is the Euler number, $w > 0$, $\mu > 0$, $\rho_{\max} \triangleq 4\max_{\rho\in\Theta\times\Xi}\|\rho\|^2$, P is the unique positive definite matrix satisfying the Lyapunov equation $P = A_m^T P A_m + R + I_n$, $R \in \mathbb{R}^{n\times n}$ is a positive definite matrix and $\lambda_{\min}(P)$ is the minimum eigenvalue of P.

Theorem 1. *For the system in* (2) *and the controller defined via* (3) – (5) *subject to the ℓ_1-norm condition in* (6)*, we have*

$$\|x_{ref} - x\|_{\ell_\infty} \le \frac{\|C(z)\|_{\ell_1}}{1-\|G(z)\|_{\ell_1}L_\theta}\|\tilde{x}\|_{\ell_\infty}\,, \tag{13}$$

where $\|\tilde{x}\|_{\ell_\infty}$ is defined in (12)*.*

Proof. The response of the closed-loop system (2) with the ℓ_1 adaptive controller (5) can be written in the z domain as:

$$\begin{aligned} x(z) \;&=\; H(z)C(z)k_g r(z) - H(z)C(z)\tilde{\eta}(z) \\ &\quad +G(z)(\theta^T x(z)+\xi) + x_{in}(z)\,. \end{aligned} \tag{14}$$

The above equality holds because $\tilde{\eta}(z) = \hat{\eta}(z) - \rho^T x_{ext}(z)$. From the definition of the closed-loop reference system (8), it follows that

$$x_{ref}(z) = H(z)k_g C(z)r(z) + G(z)\rho^T x_{ext}(z) + x_{in}(z)\,.$$

The predictor error dynamics let us show that $\tilde{x}(z) = H(z)\tilde{\eta}(z)$, which implies that

$$\begin{aligned} x_{ref}(z) - x(z) &= G(z)\theta^T(x_{ref}(z)-x(z)) + H(z)C(z)\tilde{\eta}(z) \\ &= G(z)\theta^T(x_{ref}(z)-x(z)) + C(z)\tilde{x}(z)\,, \end{aligned}$$

hence, the following uniform upper bound holds

$$\|(x_{ref}-x)\|_{\ell_\infty} \le \frac{\|C(z)\|_{\ell_1}\|\tilde{x}\|_{\ell_\infty}}{1-\|G(z)\|_{\ell_1}L_\theta}\,.$$

□

Theorem 1 shows that the ℓ_1 contolled system stays close to the BIBS stable closed-loop reference system (8); hence, the ℓ_1 controlled system is stable. Ideally, the uncertainties in the system would be in the bandwidth of the low-pass filter and the controller would cancel them exactly. In this scenario, the *ideal* system response would be:

$$x_{id}(k+1) = A_m x_{id}(k) + b_m k_g r(k)\,. \tag{15}$$

In this work the ideal model (15) is used as the model of the robust MPC. In reality not all uncertainties are canceled and $x(k)$ could deviate from $x_{id}(k)$. We prove that there is a uniform upper bound on the difference between the ideal and the ℓ_1 controlled system state. We first use Corollary 1 in [19], to show that $C(z)\tilde{\eta}(z) = H_1(z)\tilde{x}(z)$, where $H_1(z)$ is proper and BIBO stable. We also note that it is possible to show that $\|\hat{x}\|_{\ell_\infty}$ is uniformly bounded in a similar fashion as the proof of Lemma 3 in [19].

Theorem 2. *For the system* (2) *and the controller* (3) – (5) *subject to the ℓ_1-norm condition in* (6)*, the uniform upper bound of the difference between ideal and real state is*

$$\begin{aligned} \|x_{id} - x\|_{\ell_\infty} \le \;& \|H(z)(1-C(z)k_g)\|_{\ell_1}\|r\|_{\ell_\infty} \\ &+\|H(z)H_1(z)\|_{\ell_1}\|\tilde{x}\|_{\ell_\infty} \\ &+\lambda_\theta\left(\|\hat{x}\|_{\ell_\infty} + \|\tilde{x}\|_{\ell_\infty}\right) + \lambda_\xi\,. \end{aligned} \tag{16}$$

Proof. The response of the ideal system in the z domain can be written as:

$$x_{id}(z) = H(z)k_g r(z) + x_{in}(z)\,.$$

The response of the closed-loop system (2) with the ℓ_1 adaptive controller can be written in the z domain as (14). The latter implies that:

$$\begin{aligned} x(z) - x_{id}(z) &= H(z)k_g(C(z)-1)r(z) - H(z)C(z)\tilde{\eta}(z) \\ &\quad +G(z)(\theta^T x(z)+\xi) \\ &= H(z)k_g(C(z)-1)r(z) - H(z)H_1(z)\tilde{x}(z) \\ &\quad +G(z)(\theta^T x(z)+\xi)\,. \end{aligned}$$

Using the above, the following uniform bound can be derived

$$\begin{aligned} \|(x-x_{id})|_i\|_{\ell_\infty} \;\le\; & \|H(z)k_g(C(z)-1)\|_{\ell_1}\|r|_i\|_{\ell_\infty} \\ &+\|H(z)H_1(z)\|_{\ell_1}\|\tilde{x}\|_{\ell_\infty} \\ &+\lambda_\theta\left(\|\hat{x}\|_{\ell_\infty} + \|\tilde{x}\|_{\ell_\infty}\right) + \lambda_\xi\,. \end{aligned}$$

□

We have shown that the ℓ_1 adaptive controller makes a system state $x(k)$ remain close to an ideal system state $x_{id}(k)$. The actual system state may deviate from the ideal system state, but this deviation is uniformly bounded (16).

B. Robust Model Predictive Control

Robust MPC solves an optimization problem at each time step to minimize a given cost function based on (i) a model of the system, (ii) a measure of the disturbance on the system, and (iii) input and state constraints. In this work we use the robust MPC proposed by [24]. For convenience and completeness, we present a brief summary of the robust MPC. Consider a linear system with additive disturbance:

$$x_{\text{MPC}}(k+1) = A_m x_{\text{MPC}}(k) + b_m k_g u(k) + \gamma\,, \tag{17}$$

where $\gamma \in \Gamma$ is an additive disturbance affecting the system. In the proposed framework, A_m, b_m, k_g are provided by the ℓ_1 adaptive controller, and Γ is the set containing all the values that satisfy the uniform bound (16). This system is subject to state and input constraints

$$u \in \bar{\mathbb{U}}\,, \quad x_{\text{MPC}} \in \bar{\mathbb{X}}\,, \tag{18}$$

where $\bar{\mathbb{U}} \subset \mathbb{R}^m$ is compact, $\bar{\mathbb{X}} \subset \mathbb{R}^n$ is closed, and each set contains the origin in its interior. The robust MPC assumes there is an ancillary controller with gain K that makes the disturbed system (17) stay in a disturbance invariant set Z. Hence, despite the presence of disturbances the system (17) stays in an area defined by Z around the nominal system

$$x_{\text{MPC}}(k+1) = A_m x_{\text{MPC}}(k) + b_m k_g u(k)\,, \tag{19}$$

which does not have a disturbance. In nominal MPC implementations the initial state of the optimal control problem is the current state $x_{\text{MPC}}(k)$. In the robust MPC framework the initial state x_0 is included as a decision variable and a parameter of the ancillary controller. The calculated initial state x_0 must lie in the invariant set Z around the state $x_{\text{MPC}}(k)$. The cost function to be optimized by the robust MPC is:

$$V_N(x_{\text{MPC}}(k), \mathbf{u}) \triangleq \sum_{i=0}^{N_H-1} l(x_{\text{MPC}}(i), u(i)) + V_f(x_{\text{MPC}}(N_H)), \quad (20)$$

where N_H is the time horizon, $\mathbf{u} = [u(0), \ldots, u(N_H-1)]$ is the control sequence, $x_{\text{MPC}}(i)$ satisfies (19), $l(x,u) \triangleq \frac{1}{2}[x^T Q_{\text{MPC}} x + u^T R_{\text{MPC}} u]$, $V_f(x) \triangleq \frac{1}{2} x^T P_{\text{MPC}} x$, and Q_{MPC}, R_{MPC} and P_{MPC} are positive definite matrices of appropriate size. The robust MPC is constrained by (18) and $x_{\text{MPC}}(k) \in x_0 \oplus Z$, where $\oplus$ is the Minkowski set addition, i.e, $A \oplus B \triangleq \{a + b | a \in A, b \in B\}$. In order to account for the tube, in the robust MPC the constraints have to be tightened

$$u(i) \in \mathbb{U} \triangleq \bar{\mathbb{U}} \ominus KZ, \quad x_{\text{MPC}}(i) \in \mathbb{X} \triangleq \bar{\mathbb{X}} \ominus Z, \quad (21)$$

where $i = 0, \ldots, N_H - 1$, $\ominus$ denotes set subtraction, i.e., $A \ominus B \triangleq \{a | a \oplus B \subseteq A\}$. The robust MPC control law includes the ancillary controller:

$$r(x_{\text{MPC}}(k)) \triangleq u_0^*(x_{\text{MPC}}(k)) + K(x_{\text{MPC}}(k) - x_0^*(x_{\text{MPC}}(k))),$$

where $K^T \in \mathbb{R}^n$ is such that $A_m + b_m K$ is stable, $x_{\text{MPC}}(k)$ is the state of the real system at the current time step, $u_0^*(x_{\text{MPC}}(k))$ is the first input of the optimal control sequence and $x_0^*(x_{\text{MPC}}(k))$ is the initial state calculated by the robust MPC minimizing the cost function (20). Intuitively, the robust MPC finds an optimal state sequence with initial state $x_0^*(x_{\text{MPC}}(k))$ that minimizes the cost function (20) and uses an ancillary controller that makes the system stay in a disturbance invariant set Z around the optimal state sequence, despite the disturbances. In the robust adaptive MPC framework the control input $r(x_{\text{MPC}}(k))$ is the reference $r(k)$ for the underlying ℓ_1 adaptive controller.

Lessons learned: the bound Γ (from the ℓ_1 adaptive controller) may be too conservative to be used in the robust MPC. However, it can be approximated more tightly experimentally by applying a step function to the ℓ_1 controlled system and measuring the deviation from the ideal system.

IV. EXPERIMENTAL RESULTS

This section presents experimental results of the proposed robust adaptive MPC applied to a quadrotor. We assess two aspects of the proposed framework: (i) trajectory tracking performance, and (ii) obstacle avoidance with trajectory update. In [25] we showed an adaptive MPC (underlying $\mathcal{L}_1$ adaptive controller and nominal MPC) with tracking capabilities that outperform non adaptive and non predictive approaches even in the presence of external disturbances. To assess the performance of the proposed robust adaptive MPC, we compare it to adaptive MPC. We show that unlike adaptive MPC, the robust adaptive MPC enables the quadrotor to successfully complete challenging tasks. The results can be seen in a video at `http://tiny.cc/89lscz`.

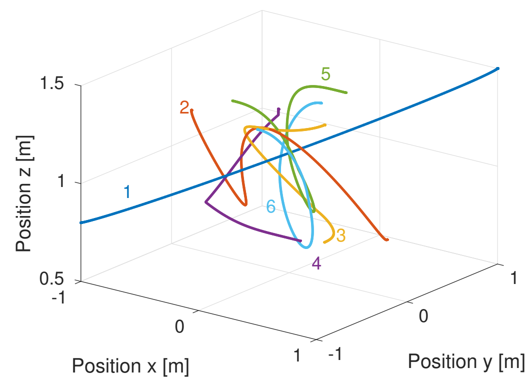

Fig. 2. The six different trajectories that are used to test trajectory tracking of the robust adaptive MPC.

We implement a discrete time state feedback ℓ_1 adaptive controller for position control of each axis, as described in Section III-A. We assume that the x, y, and z axes are decoupled. We also implement the robust MPC described in Section III-B. The signal $x(k)$ is the state and $r^*(k)$ the desired state, which in this case is composed of position and velocity in each axis. The nominal (for adaptive MPC) and robust MPC (for robust adaptive MPC) minimize the cost function (20) with the same weight gains. The differences are that robust MPC (i) includes the initial state x_0 as a decision variable, (ii) uses it for the ancillary controller, and (iii) tightens the input and state constraints (21). We define $Q_{\text{MPC}} = qI$, $R_{\text{MPC}} = rI$, and $P_{\text{MPC}} = pI$ and constrain the input to guarantee that the quadrotor remains in a given area at all times. We use IBM CPLEX optimizer to solve the above optimization problem.

The vehicle used in the experiments is the Parrot Bebop 2. A central overhead motion capture camera system provides position, roll-pitch-yaw Euler angles and rotational velocity measurements, and through numerical differentiation translational velocities of the quadrotor. The output of the ℓ_1 adaptive controller is $u_{i,\ell_1}(k)$ with $i = \mathsf{x}, \mathsf{y}, \mathsf{z}$, where x and y are commanded translational acceleration and z is commanded velocity specified in the global coordinate frame.

To quantify the performance of the controllers the average position error along the trajectory is defined:

$$e = \frac{\sum_{i=1}^{N} \sqrt{e_{\mathsf{x}}(i)^2 + e_{\mathsf{y}}(i)^2 + e_{\mathsf{z}}(i)^2}}{N},$$

where $e_j(i) = r^*_{\mathrm{j},pos}(i) - x_{j,pos}(i)$, $r^*_{j,pos}(i)$ are the desired positions and $x_{j,pos}(i)$ are the measured positions with $j = \mathsf{x}, \mathsf{y}, \mathsf{z}$. We propose six different trajectories to test our approach, as shown in Fig. 2.

In Section III-A we showed that there is an upper bound on the difference between the real and ideal state (16). The theoretical bound is conservative and not suitable to be used in a robust MPC framework. In this work we experimentally characterized this bound through simple step response experiments where we compare the response of the actual system and the ideal system in the x, y, and z axes.

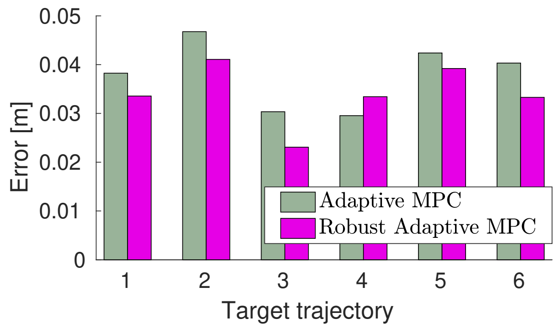

Fig. 3. Tracking error of six different trajectories with adaptive MPC and the proposed robust adaptive MPC, which achieves a lower trajectory tracking error in most trajectories.

A. Tracking Performance

We compare the adaptive MPC to the proposed robust adaptive MPC on the six test trajectories in Fig. 2. We want to show that the added robustness does not negatively affect the tracking performance. The average position tracking errors for each controller and trajectory are shown in Fig. 3. In the six trajectories both controllers have a similar performance. In five of the six trajectories, the robust adaptive MPC achieves a smaller tracking error, since it is able to account for modeling errors.

B. Obstacle Avoidance with Trajectory Update

We also assess the performance of the controllers when a cylinder obstacle is introduced in the environment and the desired trajectory avoids it, as shown in Fig. 4. The desired trajectory surrounds the obstacle and the state constraints are tightened to exclude the volume where the obstacle is. The drone must avoid the obstacle to successfully complete the task. Fig. 5 shows a top view into the x, y plane of the obstacle (in black), the desired trajectory (in red), and actual and predicted trajectories of both controllers at two different times. At 2.36 seconds (Fig. 5a) both controllers predict to go around the obstacle. Note that the position sequence predicted by the robust adaptive MPC does not begin at the current state, because it accounts for modeling errors and the ancillary controller. Since the adaptive MPC does not take into account the difference between ideal and actual system behavior, the quadrotor collides with the obstacle (red cross in Fig. 5b). The robust adaptive MPC is able to successfully avoid the obstacle and complete the task.

V. CONCLUSIONS

We presented a robust adaptive MPC framework that combines (i) a discrete time state feedback ℓ_1 adaptive controller that forces systems to behave close to a specified linear reference model despite the presence of disturbances and changes in the environment, and (ii) a robust MPC that improves tracking performance and is robust against model error. Key theoretical contributions include (i) formulating and showing performance guarantees for an ℓ_1 adaptive controller in discrete time, and (ii) proving that the difference between the linear model and the true underlying adaptive system is uniformly bounded.

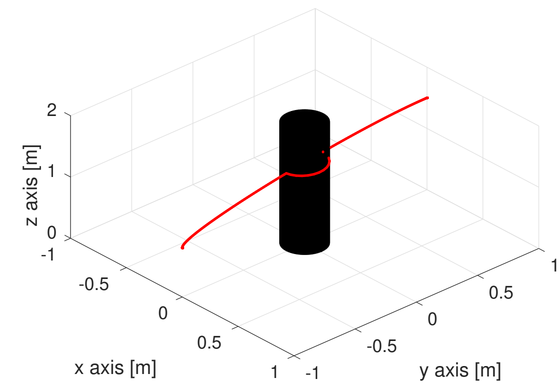

Fig. 4. Desired trajectory in red avoiding the obstacle. For successful task completion the drone must avoid the obstacle.

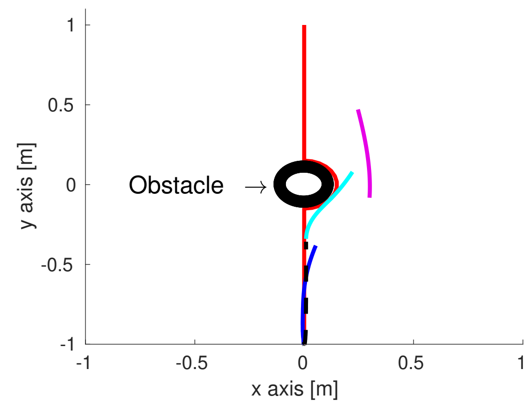

(a) Time = 2.36 [s]

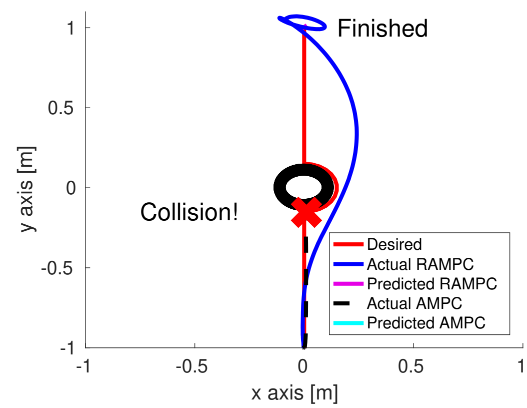

(b) Time = 6.82 [s]

Fig. 5. Actual and predicted paths of adaptive (AMPC) and robust adaptive MPC (RAMPC). RAMPC takes into account modeling errors which enables it to avoid the obstacle successfully completing the task. RAMP includes the initial state x_0 as a decision variable; hence, the predicted path (magenta) does not begin at the current position (end of solid line). AMPC does not include the potential disturbance in the system and collides with the obstacle.

Experimental results on a quadrotor show that the robust adaptive MPC achieves high accuracy tracking performance and, unlike adaptive MPC, enables the system to complete the challenging task of introducing a obstacle in the quadrotor path.

REFERENCES

[1] R. Skelton, "Model error concepts in control design," *International Journal of Control*, vol. 49, no. 5, pp. 1725–1753, 1989.

[2] S. Skogestad and I. Postlethwaite, *Multivariable feedback control: analysis and design*. Wiley New York, 2007, vol. 2.

[3] K. Pereida, M. K. Helwa, and A. P. Schoellig, "Data-efficient multirobot, multitask transfer learning for trajectory tracking," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 1260–1267, 2018.

[4] N. Hovakimyan and C. Cao, $\mathcal{L}_1$ *Adaptive Control Theory: Guaranteed Robustness with Fast Adaptation*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2010.

[5] R. W. Beard, N. B. Knoebel, C. Cao, N. Hovakimyan, and J. S. Matthews, "An L1 adaptive pitch controller for miniature air vehicles," in *Proc. of the AIAA Guidance, Navigation, and Control Conference*, 2006, pp. 5788–5794.

[6] I. Kaminer, O. Yakimenko, V. Dobrokhodov, A. Pascoal, N. Hovakimyan, V. Patel, C. Cao, and A. Young, "Coordinated path following for time-critical missions of multiple UAVs via L1 adaptive output feedback controllers," in *Proc. of the AIAA Guidance, Navigation, and Control Conference*. AIAA 2007-6409, 2007, 2007.

[7] Z. Zuo and P. Ru, "Augmented L1 adaptive tracking control of quad-rotor unmanned aircrafts," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50, no. 4, pp. 3090–3101, 2014.

[8] K. Pereida, R. R. Duivenvoorden, and A. P. Schoellig, "High-precision trajectory tracking in changing environments through $\mathcal{L}_1$ adaptive feedback and iterative learning," in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 344–350.

[9] I. M. Gregory, C. Cao, E. Xargay, N. Hovakimyan, and X. Zou, "ℓ_1 adaptive control design for NASA AirSTAR flight test vehicle," in *Proc. of the AIAA Guidance, Navigation, and Control Conference*. AIAA 2009-5738, 2009.

[10] V. V. Patel, C. Cao, N. Hovakimyan, K. A. Wise, and E. Lavretsky, "ℓ_1 adaptive controller for tailless unstable aircraft," in *Proc. of the American Control Conference (ACC)*, 2007, pp. 5272–5277.

[11] S. Mallikarjunan, B. Nesbitt, E. Kharisov, E. Xargay, N. Hovakimyan, C. Cao *et al.*, "L1 adaptive controller for attitude control of multirotors," in *Proc. of the AIAA Guidance, Navigation and Control Conference*, 2012.

[12] V. Adetola, D. DeHaan, and M. Guay, "Adaptive model predictive control for constrained nonlinear systems," *Systems & Control Letters*, vol. 58, no. 5, pp. 320–326, 2009.

[13] H. Fukushima, T.-H. Kim, and T. Sugie, "Adaptive model predictive control for a class of constrained linear systems based on the comparison model," *Automatica*, vol. 43, no. 2, pp. 301–308, 2007.

[14] V. A. Akpan and G. D. Hassapis, "Nonlinear model identification and adaptive model predictive control using neural networks," *ISA transactions*, vol. 50, no. 2, pp. 177–194, 2011.

[15] T. Wang, H. Gao, and J. Qiu, "A combined adaptive neural network and nonlinear model predictive control for multirate networked industrial process control," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 27, no. 2, pp. 416–425, 2016.

[16] I. Lenz, R. A. Knepper, and A. Saxena, "DeepMPC: Learning deep latent features for model predictive control." in *Robotics: Science and Systems*, 2015.

[17] C. J. Ostafew, A. P. Schoellig, and T. D. Barfoot, "Robust constrained learning-based NMPC enabling reliable mobile robot path tracking," *International Journal of Robotics Research*, vol. 35, no. 13, pp. 1547–1563, 2016.

[18] C. D. McKinnon and A. P. Schoellig, "Experience-based model selection to enable long-term, safe control for repetitive tasks under changing conditions," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, accepted.

[19] H. Jafarnejadsani and N. Hovakimyan, "Optimal filter design for a discrete-time formulation of l1-adaptive control," in *AIAA Infotech@ Aerospace*, 2015, p. 0119.

[20] M. Elnaggar, M. S. Saad, H. A. Fattah, and A. L. Elshafei, "Discrete time l_1 adaptive control for systems with time-varying parameters and disturbances," in *Proc. of the IEEE Conference on Decision and Control (CDC)*, 2016, pp. 2115–2120.

[21] G. Goodwin, P. Ramadge, and P. Caines, "Discrete-time multivariable adaptive control," *IEEE Transactions on Automatic Control*, vol. 25, no. 3, pp. 449–456, 1980.

[22] G. C. Goodwin and K. S. Sin, *Adaptive Filtering Prediction and Control*. Prentice Hall, 1984.

[23] L. M. Silverman and B. D. Anderson, "Controllability, observability and stability of linear systems," *SIAM Journal on control*, vol. 6, no. 1, pp. 121–130, 1968.

[24] D. Q. Mayne, M. M. Seron, and S. Raković, "Robust model predictive control of constrained linear systems with bounded disturbances," *Automatica*, vol. 41, no. 2, pp. 219–224, 2005.

[25] K. Pereida and A. P. Schoellig, "Adaptive model predictive control for high-accuracy trajectory tracking in changing conditions," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 7831–7837.