

Robust Adaptive Model Predictive Control for High-Accuracy Trajectory Tracking in Changing Conditions

Karime Pereida and Angela P. Schoellig

Abstract—Robots are being deployed in unknown and dynamic environments where they are required to handle disturbances, unmodeled dynamics, and parametric uncertainties. Sophisticated control strategies can guarantee high performance in these changing environments. In this work, we propose a novel robust adaptive model predictive controller that combines robust model predictive control (MPC) with an underlying \mathcal{L}_1 adaptive controller to improve trajectory tracking of a system subject to unknown and changing disturbances. The \mathcal{L}_1 adaptive controller forces the system to behave close to a specified linear reference model. The controlled system may still deviate from the reference model, but this deviation is shown to be upper bounded. An outer-loop robust MPC uses this upper bound, the linear reference model and system constraints to calculate the optimal reference input that minimizes the given cost function. The proposed robust adaptive MPC is able to achieve high-accuracy trajectory tracking even in the presence of unknown disturbances. We show preliminary experimental results of an adaptive MPC on a quadrotor. The adaptive MPC has a lower trajectory tracking error compared to a predictive, non-adaptive approach, even when wind disturbances are applied.

I. INTRODUCTION

Robots and automated systems are being deployed in unstructured and changing environments. Small changes in the environmental conditions may significantly deteriorate the performance and cause instability in traditional, model-based controllers (see [1] and [2]). Control methods designed for robots deployed in changing environments must be robust against model uncertainties, unknown disturbances, and changing dynamics. Certain control methods can provide theoretical performance guarantees, where the overall control performance, even in the presence of unexpected changes and unmodeled disturbances, is specified. These performance guarantees can be leveraged by robust controllers to further improve the performance of the overall system.

In this work, we present a robust adaptive controller that achieves high-accuracy tracking performance and is robust to unknown disturbances and changing dynamics. We extend previous work [3] where we proposed an output feedback \mathcal{L}_1 adaptive control and *nominal* model predictive control (MPC) framework. In this work we robustify the previous approach by combining state feedback \mathcal{L}_1 adaptive control and *robust* MPC (see Fig. 1). The underlying \mathcal{L}_1 adaptive controller forces a system to behave close to a specified linear reference model, even in the presence of unknown disturbances. The controlled system may still deviate from the reference model behavior; however, this deviation is shown to have an upper bound. Unlike [3], we leverage this upper bound and the reference model with a robust MPC

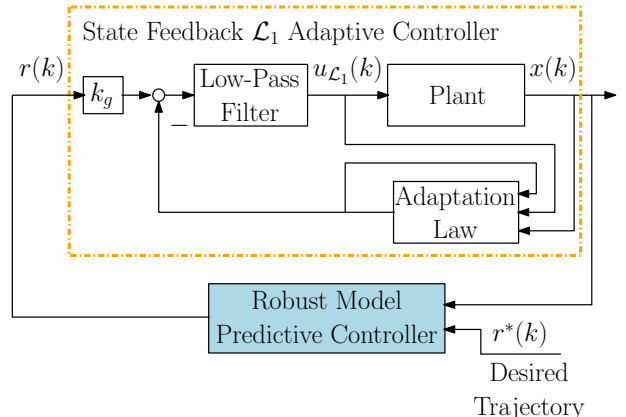


Fig. 1. The proposed robust adaptive model predictive control architecture. The underlying \mathcal{L}_1 adaptive controller forces the system to behave close to a specified linear reference model even in the presence of uncertainties and disturbances. Performance guarantees specifying how far apart the controlled system and the reference model can be are provided. A robust model predictive controller using the aforementioned performance guarantees and the specified linear reference model, calculates, at each time step, the optimal input for the plant controlled by the \mathcal{L}_1 adaptive controller (which now behaves as the specified reference model) such that the trajectory tracking error is minimized.

module to compute the optimal reference input for the \mathcal{L}_1 adaptive controlled system. The robust adaptive MPC uses performance guarantees in order to achieve high-accuracy trajectory tracking of a system subject to unknown disturbances. We highlight the theoretical contributions of this work and revisit a key experimental result presented in [3]. The preliminary results show that an adaptive MPC can improve the trajectory tracking performance of a quadrotor compared to a non-adaptive approach, even when unknown disturbances are present. The video with the preliminary results can be found here <http://tiny.cc/q60e4y>.

II. RELATED WORK

\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 model uncertainties, to behave as a specified linear reference 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 [8]. In [3], we proposed to replace ILC with a MPC, which enabled the system to achieve high-accuracy

trajectory tracking on the first iteration. In this work, we aim to improve the controller performance by adding robustness to the MPC to account for deviations between the real system and the reference model behavior.

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. In this way, 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 method 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. The algorithm is initially conservative due to large parameter uncertainties, but performance is improved over time as the parameter estimate's uncertainty is reduced and the estimate converges to the true value. However, the quality of the estimation depends on the excitation of the state. Finally, 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 this work we combine a state feedback \mathcal{L}_1 adaptive controller with a robust MPC and make the following key contributions:

- use the performance guarantees provided by the \mathcal{L}_1 adaptive controller in a robust MPC to control systems subject to disturbances and changing environments;
- remove the need for persistent excitation to achieve accurate adaptation as in existing robust adaptive MPC strategies [13]; and
- validate the proposed approach through extensive experimental results that include external disturbances, (e.g. wind).

III. PROBLEM STATEMENT

The objective of this work is to achieve high-accuracy trajectory tracking in the presence of uncertain, and possibly changing conditions on the first iteration. Consider a system whose dynamics ('Plant' block in Fig. 1) are 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:

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

where $x(k) \in \mathbb{R}^n$ is the system state vector (assumed to be measured); $u(k)$ 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 and ξ are the *unknown* parameters, which belong 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 a hypercube. The parameter ξ is bounded by $|\xi| \leq L_\xi$ where Ξ is the set that includes all the values less than or equal to L_ξ . The boundaries of Θ and Ξ can be represented as:*

$$e_\theta^T \theta = b_\theta, \quad e_\xi^T \xi = b_\xi,$$

where $e_\theta \in \mathbb{R}^{n \times 2n}$, $e_\xi \in \mathbb{R}^{1 \times 2}$ have exactly one value per column not equal to zero and $b_\theta \in \mathbb{R}^{2n}$ and $b_\xi \in \mathbb{R}^2$ contain the maximum and minimum values in a given dimension.

The system is required to accurately track a desired trajectory $x^*(k)$ defined over a finite number of steps $N < \infty$ and assumed to be feasible with respect to the true dynamics of the \mathcal{L}_1 -controlled system (orange dashed box in Fig. 1). The desired trajectory can be written as $\mathbf{x}^* = (x^*(1), \dots, x^*(N))$, and the state of the plant as $\mathbf{x} = (x(1), \dots, x(N))$. The goal is to take into account system disturbances when minimizing the tracking performance criterion J defined as:

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

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

IV. METHODOLOGY

We propose a robust adaptive control framework to achieve high-accuracy trajectory tracking in the first iteration for a system subject to unknown disturbances. We consider two subsystems: (1) the discrete-time, state feedback \mathcal{L}_1 adaptive controller (orange dashed box in Fig. 1) that makes the system behave close to a predefined linear system, even in the presence of disturbances; and (2) the robust MPC (blue box in Fig. 1) that includes disturbances in the calculation of the optimal input that minimizes the tracking error.

A. Discrete-Time, State Feedback \mathcal{L}_1 Adaptive Control

\mathcal{L}_1 adaptive control forces a system to behave close to a specified reference linear model and provides transient and steady-state performance guarantees. These performance guarantees are upper bounds that, given certain disturbance signals, specify how far the \mathcal{L}_1 controlled system can deviate from the reference linear system. These upper bounds need to be computed for the implementation of robust MPC, which requires information on the disturbance present in the system. In particular, 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*; therefore, the bounds cannot be calculated. For this reason, we develop the performance guarantees for a discrete-time, state feedback \mathcal{L}_1 adaptive controller. For this controller the

performance guarantees can be calculated as all the variables are known. Finally, the robust MPC framework used in this work [19] requires a discrete-time model of the system to be controlled. Therefore, we present the discrete-time, state feedback \mathcal{L}_1 adaptive controller.

The discrete-time, state feedback \mathcal{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):

$$u(k) = u_m(k) + u_{\mathcal{L}_1}(k), \quad u_m(k) = -k_m^T x(k),$$

where $k_m \in \mathbb{R}^n$ renders $A_m \triangleq A - bk_m^T$ Schur stable [20], while $u_{\mathcal{L}_1}(k)$ is the adaptive component, which will be defined shortly. The static feedback gain k_m leads to the following partially closed-loop system:

$$\begin{aligned} x(k+1) &= A_m x(k) + b_m(u_{\mathcal{L}_1}(k) + \theta^T x(k) + \xi), \\ x(0) &= x_0, \quad y(k) = c_m^T x(k). \end{aligned} \quad (2)$$

The equations that describe the implementation of the discrete-time, state feedback \mathcal{L}_1 adaptive controller are:

State Predictor: We use the following state predictor:

$$\begin{aligned} \hat{x}(k+1) &= A_m \hat{x}(k) + b_m \left(u_{\mathcal{L}_1}(k) + \hat{\theta}^T(k)x(k) + \hat{\xi}(k) \right), \\ \hat{x}(0) &= x_0, \quad \hat{y}(k) = c_m^T \hat{x}(k), \end{aligned} \quad (3)$$

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:

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

where $x_{ext}(k) = [x^T(k), 1]^T$, $\hat{\rho}(k) = [\hat{\theta}^T, \hat{\xi}^T]^T$, $\hat{\rho}(0) = \hat{\rho}_0 \in [\Theta, \Xi]^T$, and $b_0 \triangleq \frac{b_m}{b_m^T b_m}$ is a constant vector. State measurements and the linear model (2) are used to calculate $\theta^T x(k) + \xi = x(k+1) - A_m x(k) - b_m u_{\mathcal{L}_1}(k)$.

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

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

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

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 DC gain $C(1) = 1$, and its state-space realization assumes zero initialization.

The discrete-time state feedback \mathcal{L}_1 adaptive controller is defined via (3) – (5) with $C(z)$ verifying the following \mathcal{L}_1 -norm condition:

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

where

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

Ideally, the uncertainties are within the bandwidth of the low-pass filter and the controller is able to cancel the uncertainties in the system exactly. In this ideal scenario, the system response is the following:

$$x_{id}(k+1) = A_m x_{id}(k) + b_m k_g r(k). \quad (8)$$

The system described above by A_m, b_m can be used as a model in a robust MPC scenario. In reality not all uncertainties are canceled and $x(k)$ could deviate from $x_{id}(k)$. It can be shown that for the system (2) and the controller defined via (3) – (5) subject to the \mathcal{L}_1 -norm condition in (6), there exists an upper bound to this deviation, i.e., $\|x - x_{id}\|_{\ell_\infty} \leq \Gamma$. In the discrete-time, state feedback case, Γ depends on known system variables such as A_m, b_m ; known disturbance bounds L_θ, L_ξ ; and design parameters $C(z)$. Hence, this bound can be calculated.

B. Robust Model Predictive Control

Robust MPC is able to account for disturbances in the calculation of the optimal input for a given system. In the previous section, we showed that the proposed \mathcal{L}_1 adaptive controller is able to make the real system $x(k)$ behave close to the ideal $x_{id}(k)$ system described in (8). However, there could be a deviation from the ideal system and this deviation is upper bounded by Γ . For this reason, we could rewrite the behavior of the system as:

$$x_{MPC}(k+1) = A_m x_{MPC}(k) + b_m k_g r(k) + \gamma, \quad (9)$$

where $\gamma \in \Gamma$ is an additive disturbance affecting the ideal system and Γ is the upper bound of the difference between real and ideal system state. In this work, we propose a robust MPC approach as described in [19]. This robust MPC approach is for systems with bounded additive disturbance as (9). This system is subject to state and input constraints

$$u \in \mathbb{U}, \quad x \in \mathbb{X}, \quad (10)$$

where $\mathbb{U} \subset \mathbb{R}^m$ is compact, $\mathbb{X} \subset \mathbb{R}^n$ is closed, and each set contains the origin in its interior. The optimal control problem for robust MPC modifies the conventional optimal control problem by including the initial state x_0 as a parameter of the control law, instead of setting the initial state as the current state $x(k)$. Hence, the cost function can be written as:

$$V_N(x(k), \mathbf{u}) \triangleq \sum_{i=0}^{N-1} l(x(i), u(i)) + V_f(x(N)), \quad (11)$$

where $l(x, u) \triangleq \frac{1}{2}[x^T Q x + u^T R u]$, $V_f(x) \triangleq \frac{1}{2}x^T P x$, and Q, R and P are positive definite. We propose the following control law, yielded by the minimization of the

cost function (11):

$$r(x(k)) \triangleq u_0^*(x(k)) + K(x(0) - x_0^*(x(k))),$$

where $K^T \in \mathbb{R}^n$ is such that $A_m + b_m K$ is stable, $x(0)$ is the measured state of the real system at the current time step, $u_0^*(x(k))$ is the first input in the optimal control sequence and $x_0^*(x(k))$ is the initial state calculated by the robust MPC. Intuitively, this robust MPC implementation finds an initial state that reduces the overall cost of (11) and devises a control law that makes the system get closer to the calculated optimal initial state. The calculated input $r(x(k))$ is the reference signal $r(k)$ for the underlying \mathcal{L}_1 adaptive controller.

Lessons learned: the bound Γ (from the \mathcal{L}_1 adaptive controller) is too conservative to be used in a robust MPC. However, an experimental bound can be obtained by applying a step function to the \mathcal{L}_1 controlled framework and measuring the deviation from the ideal system.

V. PRELIMINARY EXPERIMENTAL RESULTS

In this section we present preliminary experimental results. These results were presented in [3] where an adaptive MPC that combines output feedback \mathcal{L}_1 adaptive control and *conventional* MPC was used. The video with the results can be found here <http://tiny.cc/q60e4y>. Note that a conventional MPC instead of a robust MPC and an output feedback instead of state feedback \mathcal{L}_1 adaptive controller were implemented. However, we want to highlight that an adaptive MPC that combines MPC with an underlying \mathcal{L}_1 adaptive controller is able to achieve high-accuracy trajectory tracking even in the presence of disturbances.

The adaptive MPC framework is used to control a Parrot Bebop 2 flying five different three-dimensional trajectories where it may be subject to dynamic disturbances. We assess two aspects to verify the effectiveness of the adaptive MPC framework: (i) trajectory tracking capability, and (ii) robustness to external disturbances. To quantify the tracking performance, an average position error along the trajectory is defined by:

$$e = \frac{1}{N} \sum_{i=1}^N \sqrt{e_x^2(i) + e_y^2(i) + e_z^2(i)}, \quad (12)$$

where $e_j(i) = r_j(i) - y_j(i)$ and $j = x, y, z$, where r is the desired trajectory and y is the output.

To assess the trajectory tracking capability, the adaptive MPC framework MPC- \mathcal{L}_1 is compared to an MPC-PID framework. The MPC-PID framework replaces the \mathcal{L}_1 adaptive controller with a PID controller. The system model used in the MPC is obtained by applying a step input to the x , y , and z directions separately and characterizing the system response when the quadrotor is controlled by the PID controller. The system is assumed to be a second-order, linear system in each direction. The cost functions used in the MPC-PID and MPC- \mathcal{L}_1 frameworks are the same. The average position errors of both approaches are shown in Fig. 2 in dark blue and dark red, respectively.




Fig. 2. Average position tracking errors over five different trajectories when using the MPC-PID and MPC- \mathcal{L}_1 controllers. The cases when no disturbance is applied to the system are depicted in dark blue (MPC-PID) and dark red (MPC- \mathcal{L}_1). A fan is used to apply a wind disturbance to the system and the resulting errors are depicted in light blue (MPC-PID) and light red (MPC- \mathcal{L}_1). The proposed MPC- \mathcal{L}_1 framework is able to compensate for the disturbance applied and has a performance similar to the case when no disturbance is applied. From [3].

In order to obtain the best performance, the system model obtained through step response experiments for the MPC-PID framework would need to be tuned to each trajectory. However, in order to fairly compare both approaches, the model used in the MPC-PID framework is kept constant across trajectories. The MPC- \mathcal{L}_1 framework achieves a lower error in all trajectories.

To assess robustness to disturbances, we introduce a wind disturbance with a fan at different points in each trajectory. The resulting average errors when the fan disturbance is applied are shown in Fig. 2 in light blue and light red for the MPC-PID and MPC- \mathcal{L}_1 frameworks, respectively. The MPC-PID approach performs generally worse when wind is applied. The MPC- \mathcal{L}_1 framework is able to keep approximately the same performance when a disturbance is applied since the underlying \mathcal{L}_1 adaptive controller is able to compensate for it. However, we expect the tracking performance to further improve by taking into account the disturbance bound Γ in the optimal control problem.

VI. CONCLUSIONS

In this work, we presented a robust adaptive MPC framework that combines a discrete-time, state feedback \mathcal{L}_1 adaptive controller and robust MPC to improve trajectory tracking performance of a system subject to unknown disturbances. The \mathcal{L}_1 adaptive controller forces systems to behave close to a specified linear reference model. Performance guarantees show that deviations from the ideal behavior, albeit bounded, remain. A robust MPC takes into account the disturbance remaining in the system to solve an optimal control problem that improves trajectory tracking performance. Preliminary experimental results using an output feedback and conventional MPC show that an adaptive MPC is able to improve the trajectory tracking performance of a quadrotor even in the presence of unknown disturbances and unexpected changes.

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 multi-robot, 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 \mathcal{L}_1 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 \mathcal{L}_1 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 \mathcal{L}_1 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. Duivendoorn, 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, " \mathcal{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, " \mathcal{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.*, " \mathcal{L}_1 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] 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.
- [20] G. Halikias, L. Dritsas, A. Pantelous, and V. Tsoukias, "Strong stability of discrete-time systems," *Linear Algebra and its Applications*, vol. 436, no. 7, pp. 1890–1908, 2012.
- [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.