

Dual adaptive MPC using an exact set-membership reformulation ^{*}

Anilkumar Parsi ^{*} Diyou Liu ^{*} Andrea Iannelli ^{**}
Roy S. Smith ^{*}

^{*} *Automatic Control Laboratory, ETH Zurich, Zürich 8092 (e-mail: aparsi,rsmith@control.ee.ethz.ch, diyliu@student.ethz.ch)*

^{**} *University of Stuttgart, Institute for Systems Theory and Automatic Control, Stuttgart 70569 (e-mail: andrea.iannelli@ist.uni-stuttgart.de)*

Abstract: Adaptive model predictive control (MPC) methods using set-membership identification to reduce parameter uncertainty are considered in this work. Strong duality is used to reformulate the set-membership equations exactly within the MPC optimization. A predicted worst-case cost is then used to enable performance-oriented exploration. The proposed approach guarantees robust constraint satisfaction and recursive feasibility. It is shown that method can be implemented using homothetic tube and flexible tube parameterizations of state tubes, and a simulation study demonstrates performance improvement over state-of-the-art controllers.

Keywords: Adaptive MPC, dual control, set-membership, uncertain linear systems

1. INTRODUCTION

In order to optimally control systems affected by model uncertainty, control actions must have two features, (i) excite the system to obtain new information (called exploration); (ii) use available information to minimize a control cost (called exploitation). Achieving the optimal exploration-exploitation trade-off is the dual control problem, which is a challenging problem in control theory, studied for over six decades (Feldbaum (1960-61); Mesbah (2018)). Recent advances in computing capabilities have motivated research on controllers which can learn model uncertainties while performing control tasks (Klenske and Hennig (2016)), raising the relevance of the dual control problem.

One popular class of controllers which has been studied in this regard is model predictive control (MPC). The ability of MPC to ensure state and input constraint satisfaction in the face of uncertainty has been viewed as an important feature for learning-based control, because it can guarantee safety of the system (Hewing et al. (2020)). Robust adaptive MPC is one such method, where features from robust MPC (such as tube MPC (Kouvaritakis and Cannon (2015))) are combined with online identification techniques (such as set-membership (Milanese and Vicino (1991))) to ensure safe learning of the model uncertainty. Using various combinations of these two features, multiple robust adaptive MPC algorithms have been proposed in the recent past, e.g., Lorenzen et al. (2019); Bujarbaruah et al. (2020); Soloperto et al. (2020); Köhler et al. (2021).

However, most existing robust adaptive MPC algorithms use *passive* exploration. That is, the controller consists of two separately designed components: the MPC optimizer which selects inputs based only on exploitation; and the online identification algorithm which uses the collected

measurements to update model uncertainty. This approach can result in suboptimality, because the dual control effect is unmodeled. One way to partially address this issue is to enforce persistence of excitation on the control inputs (Lorenzen et al. (2019)). However, optimal dual control trade-off would require extensive tuning of these methods, and tuning methods have been proposed in Soloperto et al. (2020). An alternative strategy is to use performance-based exploration, where exploration is induced by modeling the expected performance improvement from reduction in uncertainty. This approach was considered in Parsi et al. (2022), where set-membership identification is used to reduce model uncertainty, and the key novelty consisted of including the effect of identification in the MPC optimizer. Nevertheless, this required approximations which result in overestimating the capabilities of the identification algorithm, and can lead to suboptimal exploratory actions.

In this work, we propose a novel way to exactly reformulate set-membership identification inside MPC optimization problems. The proposed method has two advantages. First, it results in an improved navigation of the dual control trade-off compared to existing methods because the identification is exactly modeled within the MPC controller. Second, the method can be used to generate dual control algorithms starting from multiple existing formulations of passive adaptive MPC. We highlight this feature by implementing dual control using homothetic and flexible state tube parameterizations, and demonstrate performance improvement using a simulation study.

1.1 Notation

The sets of real numbers and non-negative real numbers are denoted by \mathbb{R} and $\mathbb{R}_{\geq 0}$ respectively. The sequence of integers from n_1 to n_2 is represented by $\mathbb{N}_{n_1}^{n_2}$. For a vector b , $\|b\|_{\infty}$ represents its ∞ -norm. The i^{th} row of a matrix A is denoted by $[A]_i$, and $\mathbf{1}$ denotes a column

^{*} This work was supported by the Swiss National Science Foundation under Grant 200021_178890.

vector of appropriate length with all entries 1. The convex hull operator is represented by $\text{co}\{\cdot\}$. The notation $a_{l|k}$ denotes the value of a at time $k+l$ computed at time k , I_n denotes the identity matrix with n rows, and \oplus denotes the Minkowski sum of two sets.

2. PROBLEM FORMULATION

Consider an uncertain, linear system described by

$$x_{t+1} = A(\theta)x_t + B(\theta)u_t + w_t, \quad (1)$$

where $x_t \in \mathbb{R}^n$ represents the state, $u_t \in \mathbb{R}^m$ the control input and $w_t \in \mathbb{R}^n$ the additive disturbance at time t . The state space matrices have the parametric description

$$A(\theta) = A_0 + \sum_{i=1}^p A_i[\theta]_i, \quad B(\theta) = B_0 + \sum_{i=1}^p B_i[\theta]_i, \quad (2)$$

where $\theta \in \mathbb{R}^p$ is a constant but unknown parameter, and $\{A_i, B_i\}_{i=0}^p$ are known matrices which model structured uncertainty. The true value of $\theta = \theta^*$ is known to lie inside a compact set

$$\Theta := \{\theta \in \mathbb{R}^p \mid H_\theta \theta \leq h_\theta\}, \quad (3)$$

where $H_\theta \in \mathbb{R}^{n_\theta \times p}$ and $h_\theta \in \mathbb{R}^{n_\theta}$. The disturbance w_t is unknown, but always lies within the compact set

$$\mathbb{W} := \{w \in \mathbb{R}^n \mid H_w w \leq h_w\}, \quad (4)$$

where $H_w \in \mathbb{R}^{n_w \times n}$, $h_w \in \mathbb{R}^{n_w}$ and \mathbb{W} contains the origin. The state of the system is perfectly measured at each time step. Given an initial state x_0 , the objective is to track a reference trajectory given by $r_t \in \mathbb{R}^n$ for time steps t in $[0, T]$, where T is the simulation length. This objective is specified as a minimization of the worst-case deviation of the states and inputs from the optimal reference, i.e.,

$$J = \min_{u_t} \max_{w_t \in \mathbb{W}} \sum_{t=0}^T \|Q(x_t - r_t)\|_\infty + \|R(u_t - u_t^*)\|_\infty, \quad (5)$$

where Q and R are positive definite matrices, and $\{u_t^*\}_{t=0}^T$ is a sequence of (possibly) unknown input setpoints which track the reference trajectory while satisfying the dynamics of the true system. The states and inputs of the system are constrained to lie inside the compact set

$$\mathbb{Z} = \{(x, u) \in \mathbb{R}^n \times \mathbb{R}^m \mid Fx + Gu \leq \mathbf{1}\}, \quad (6)$$

where $F \in \mathbb{R}^{n_c \times n}$ and $G \in \mathbb{R}^{n_c \times m}$ are known matrices.

In the following sections, a dual control algorithm will be described to perform the above task. The proposed method will be implemented using two popular parameterizations of state tubes from literature, namely homothetic tubes (HT) (Lorenzen et al. (2019)) and flexible tubes (FT) (Lu and Cannon (2019)).

3. IDENTIFICATION AND ROBUST STATE TUBE

3.1 Set-membership identification and parameter estimates

In order to reduce the uncertainty in θ , set-membership identification (Milanese and Vicino (1991)) is used to construct a sequence of sets $\{\Theta_k\}_{k=0}^T$ satisfying

$$\Theta_{k-1} \cap \Delta_k \subseteq \Theta_k \subseteq \Theta_{k-1}, \quad \Theta_0 = \Theta, \quad (7)$$

$$\Delta_k = \{\theta \mid x_{i+1} - A(\theta)x_i - B(\theta)u_i \in \mathbb{W}, \forall i \in \mathbb{N}_{k-\tau}^{k-1}\} \\ = \{\theta \mid -H_w D_i \theta \leq h_w + H_w d_{i+1}, \forall i \in \mathbb{N}_{k-\tau}^{k-1}\},$$

where Δ_k is the set of all parameters θ that could have generated the last τ measurements $\{x_i\}_{i=k-\tau+1}^k$, and

$$D_i := D(x_i, u_i) = [A_1 x_i + B_1 u_i, \dots, A_p x_i + B_p u_i], \quad (8) \\ d_{i+1} := A_0 x_i + B_0 u_i - x_{i+1}, \quad \forall i \in \mathbb{N}_{k-\tau}^{k-1}.$$

The set Θ_k is described by a fixed number of hyperplanes for computational efficiency, i.e.,

$$\Theta_k := \{\theta \mid H_\theta \theta \leq h_{\theta_k}\}. \quad (9)$$

Here h_{θ_k} is initialized with the initial estimate of uncertainty h_θ , and is updated such that (7) is satisfied by solving a set of linear programs, as described in Liu (2022a). Whereas HT adaptive MPC methods allow H_θ in (9) to be arbitrarily chosen, FT methods restrict Θ_k to polytopes with hyperplanes parallel to coordinate axes (or hyperboxes). The numerical advantage is that the vertices of Θ_k , denoted by $\{\theta_k^j\}_{j=1}^{q_\theta}$, can be readily computed as a predefined linear combination of h_{θ_k} (Liu (2022a)).

Lemma 1. (Lorenzen et al., 2019, Theorem 14(ii)) Under the set-membership scheme (7)-(9), $\theta^* \in \Theta_k$ for all $k \geq 0$.

In addition to Θ_k , an estimate $\bar{\theta}_k \in \Theta_k$ of the parameter θ^* is computed using a least mean squares filter

$$\tilde{\theta}_{k+1} = \bar{\theta}_k + \mu D(x_k, u_k)^T (x_{k+1} - A(\bar{\theta}_k)x_k - B(\bar{\theta}_k)u_k), \quad (10) \\ \bar{\theta}_{k+1} = \Pi_{\Theta_k}(\tilde{\theta}_{k+1}),$$

where Π_{Θ_k} represents a projection operator, and $\bar{\theta}_0$ is an initial guess. Note that in (10), computing $\tilde{\theta}_{k+1}$ only requires matrix multiplications, and the projection onto Θ_k can be performed by solving a convex quadratic program. Alternatively, the Chebyshev center of the set Θ_k can also be used as a parameter estimate, which requires solving a linear program (Boyd and Vandenberghe (2004)).

3.2 Control policies and tracking formulation

The control input is parameterized as

$$u_{l|k}(x) = K(x - r_{k+l}) + v_{l|k}, \quad \forall l \in \mathbb{N}_0^{N-1}, \quad (11)$$

where $v_{l|k}$ are online optimization variables, K is a stabilizing feedback gain and N is the MPC prediction horizon.

Assumption 1. The parameter set Θ is such that there exists a feedback gain K which asymptotically stabilizes $A_{cl}(\theta) = A(\theta) + B(\theta)K$, $\forall \theta \in \Theta$.

The input setpoints to be tracked are unknown due to the uncertainty in the true parameter. Hence, the MPC problem estimates u_{k+l}^* using $\bar{\theta}_k$ and optimization variables $\bar{u}_{l|k}$ such that

$$[A(\bar{\theta}_k) \ B(\bar{\theta}_k)] \begin{bmatrix} r_{k+l} \\ \bar{u}_{l|k} \end{bmatrix} = [r_{k+l+1}], \quad \forall l \in \mathbb{N}_0^{N-1}, \quad (12) \\ [A(\bar{\theta}_k) \ B(\bar{\theta}_k)] \begin{bmatrix} r_{k+N} \\ \bar{u}_{N|k} \end{bmatrix} = [r_{k+N}].$$

The setpoint for timestep $k+N$ is chosen to be an equilibrium point, in order to define the terminal components of MPC (discussed in Section 3.4). Under assumptions on the rank of the system matrices and the number of available inputs (Parsi et al., 2022, Assumption 2), it can be guaranteed that there always exists a feasible input sequence satisfying (12).

3.3 State tube construction

Using the control policy (11), a sequence of sets $\{\mathbb{X}_{l|k}\}_{l=0}^N$ is constructed such that $\mathbb{X}_{0|k} \ni x_k$ and for all $x \in \mathbb{X}_{l|k}$, $\theta \in \Theta_k$ and $w \in \mathbb{W}$,

$$\mathbb{X}_{l+1|k} \ni A(\theta)x + B(\theta)u(x) + w, \quad l \in \mathbb{N}_0^{N-1}. \quad (13)$$

This sequence of sets is called the robust state tube (RST), as it contains all feasible state trajectories within the MPC prediction horizon. The RST is constructed by outer approximating the reachable states using parameterized sets for computational efficiency.

Homothetic tubes. In an HT formulation, the RST is parameterized as

$$\mathbb{X}_{l|k}^H = \{z_{l|k}\} \oplus \alpha_{l|k} \mathbb{X}_0, \quad (14)$$

where $z_{l|k} \in \mathbb{R}^n$ and $\alpha_{l|k} \in \mathbb{R}_{\geq 0}$ are translation and scaling variables and \mathbb{X}_0 is a predefined polytope described in both hyperplane and vertex forms

$$\mathbb{X}_0^H := \{x | H_x x \leq \mathbf{1}\} = \text{co}\{x^1, x^2, \dots, x^q\}, \quad (15)$$

where $H_x \in \mathbb{R}^{n_x \times n}$. The following notation is introduced

$$\begin{aligned} x_{l|k}^j &= z_{l|k} + \alpha_{l|k} x^j, & u_{l|k}^j &= K(x_{l|k}^j - r_{k+l}) + v_{l|k}, \\ D_{l|k}^j &= D(x_{l|k}^j, u_{l|k}^j), & d_{l|k}^j &= A_0 x_{l|k}^j + B_0 u_{l|k}^j - z_{l+1|k}, \\ [\bar{f}]_i &= \max_{x \in \mathbb{X}_0} [F + GK]_i x, & [\bar{w}]_i &= \max_{w \in \mathbb{W}} [H_x]_i w, \end{aligned} \quad (16)$$

where $j \in \mathbb{N}_1^q$, and \bar{f} and \bar{w} are computed offline.

Proposition 1. (Lorenzen et al., 2019, Proposition 9) Let the RST be parameterized according to (14). Then, (6) and (13) are satisfied if, $\forall j \in \mathbb{N}_1^q$ and $l \in \mathbb{N}_0^{N-1}$, there exist $\Lambda_{l|k}^{H,j} \in \mathbb{R}_{\geq 0}^{n_x \times n_\theta}$ such that

$$(F + GK)z_{l|k} + G(v_{l|k} - Kr_{k+l}) + \alpha_{l|k} \bar{f} \leq \mathbf{1}, \quad (17a)$$

$$-H_x z_{0|k} - \alpha_{0|k} \mathbf{1} \leq -H_x x_k, \quad (17b)$$

$$\Lambda_{l|k}^{H,j} h_{\theta_k} + H_x d_{l|k}^j - \alpha_{l+1|k} \mathbf{1} \leq -\bar{w}, \quad (17c)$$

$$H_x D_{l|k}^j = \Lambda_{l|k}^{H,j} H_\theta. \quad (17d)$$

The reformulation (17c)-(17d) is obtained by imposing (13) at each vertex of $\mathbb{X}_{l|k}$, and using strong duality (Boyd and Vandenberghe (2004)) to ensure (13) for all $\theta \in \Theta_k$.

Flexible tubes. In an FT formulation, the state tube is parameterized as

$$\mathbb{X}_{l|k}^F = \{x | H_x(x - r_{k+l}) \leq h_{x_{l|k}}\}, \quad \forall l \in \mathbb{N}_0^N, \quad (18)$$

where H_x is a matrix chosen offline, and $h_{x_{l|k}}$ are online optimization variables. The following proposition can then be used to reformulate the tube propagation constraints.

Proposition 2. (Lu and Cannon, 2019, Lemmas 8 and 9) Let the RST be parameterized according to (18). Then, (6) and (13) are satisfied if, $\forall i \in \mathbb{N}_1^{n_x}$, $j \in \mathbb{N}_1^{n_\theta}$ and $l \in \mathbb{N}_0^{N-1}$,

$$\Lambda_{1,l|k}^F (h_{x_{l|k}} + H_x r_{k+l}) + G(v_{l|k} - Kr_{k+l}) \leq \mathbf{1},$$

$$H_x(x_k - r_k) \leq h_{x_{0|k}},$$

$$[1 \ \theta_k^j]^T \Lambda_{2,l|k}^{F,i} (h_{x_{l|k}} + H_x r_{k+l}) - [H_x]_i r_{k+l+1} -$$

$$[H_x]_i (B(\theta_k^j)(Kr_{k+l} - v_{l|k})) + [\bar{w}]_i \leq [h_{x_{l+1|k}}]_i, \quad (19)$$

where $\Lambda_{1,l|k}^F \in \mathbb{R}_{\geq 0}^{n_x \times n_x}$ and $\Lambda_{2,l|k}^{F,i} \in \mathbb{R}^{(p+1) \times n_x}$ are Lagrange multipliers and are computed offline by solving $\forall l \in \mathbb{N}_0^{N-1}$,

$$\begin{aligned} \forall s \in \mathbb{N}_1^{n_c}, \quad & [\Lambda_{1,l|k}^F]_s = \arg \min_{\lambda} \lambda(\mathbf{1} + \mu H_x r_{k+l}) \\ \text{s.t.} \quad & \lambda \geq 0, \quad \lambda H_x = [F + GK]_s, \end{aligned} \quad (20)$$

$$\Lambda_{2,l|k}^{F,i} = \arg \min_{\Lambda} \max_{j \in \mathbb{N}_1^{n_\theta}} [1 \ \theta_k^j]^T \Lambda (\mathbf{1} + \mu H_x r_{k+l})$$

$$\text{s.t.} \quad [1 \ \theta_k^j]^T \Lambda \geq 0, \quad \forall j \in \mathbb{N}_1^{n_\theta},$$

$$\Lambda H_x = [[H_x]_i (A_0 + B_0 K)^T \dots [H_x]_i (A_p + B_p K)^T]^T, \quad (21)$$

where $\mu \in \mathbb{R}_{\geq 0}$.

The reformulation (19) is obtained by imposing (13) at each vertex of Θ_k , and using duality to ensure (13) for all $x \in \mathbb{X}_{l|k}$. In addition, duality is also used to reformulate (6). The tuning factor μ facilitates the reduction of conservatism by weighing the effects of both $h_{x_{l|k}}$ and r_{k+l} on (6) and (13) while choosing the Lagrange multipliers offline. Note that Proposition 2 is a slight modification of the result in Lu and Cannon (2019), where the reference was the origin.

Remark 1. The FT formulation requires that both hyperplanes and vertices of Θ_k are available, and HT requires the same for \mathbb{X}_0 . Moreover, the FT formulation computes the Lagrange multipliers offline for faster online solve times, which results in conservatism (Köhler et al. (2019)).

3.4 Terminal Sets

In order to ensure recursive feasibility, terminal sets are used. That is, the last set in the RST must be an invariant set under a terminal control law of the form

$$u_{N|k}(x) = K(x - \tilde{x}_{N|k}) + v_{N|k}, \quad (22)$$

where $v_{N|k}$ are MPC optimization variables and $\tilde{x}_{N|k}$ is chosen as specified later. Then, the invariance conditions can be written as, for all $x \in \mathbb{X}_{N|k}$, $\theta \in \Theta_k$ and $w \in \mathbb{W}$,

$$\mathbb{X}_{N|k} \ni A(\theta)x + B(\theta)u_{N|k}(x) + w, \quad (x, u_{N|k}(x)) \in \mathbb{Z}. \quad (23)$$

Proposition 3. The constraints (23) are satisfied

- a) under the HT parameterization (14) and $\tilde{x}_{N|k} = z_{N|k}$, if $\exists \Lambda_{N|k}^{H,j} \in \mathbb{R}_{\geq 0}^{n_x \times n_\theta}$ such that $\forall j \in \mathbb{N}_1^q$

$$Fz_{N|k} + Gv_{N|k} + \alpha_{N|k} \bar{f} \leq \mathbf{1},$$

$$\Lambda_{N|k}^{H,j} h_{\theta_k} + H_x d_{N|k}^j - \alpha_{N|k} \mathbf{1} \leq -\bar{w}, \quad (24)$$

$$\Lambda_{N|k}^{H,j} \geq 0, \quad H_x D_{N|k}^j = \Lambda_{N|k}^{H,j} H_\theta.$$

- b) under the FT parameterization (18) and $\tilde{x}_{N|k} = r_{k+N}$, if $\forall i \in \mathbb{N}_1^{n_x}$, $j \in \mathbb{N}_1^{n_\theta}$,

$$\Lambda_{1,N|k}^F (h_{x_{N|k}} + H_x r_{k+N}) + G(v_{N|k} - Kr_{k+N}) \leq \mathbf{1},$$

$$\begin{aligned} [1 \ \theta_k^j]^T \Lambda_{2,N|k}^{F,i} (h_{x_{N|k}} + H_x r_{k+N}) - [H_x]_i r_{k+N} - \\ [H_x]_i (B(\theta_k^j)(Kr_{k+N} - v_{N|k})) + \bar{w} \leq [h_{x_{N|k}}]_i. \end{aligned} \quad (25)$$

The proof is similar to that of Propositions 1 and 2, and is omitted. A detailed proof can be found in Liu (2022a).

Assumption 2. In both HT and FT methods, the predefined state tube matrix H_x is chosen such that for all $\theta \in \Theta$ and $x \in \{x | H_x x \leq \mathbf{1}\}$, $\exists \lambda_c \in [0, 1)$ satisfying

$$H_x(A(\theta)x + B(\theta)Kx) \leq \lambda_c \mathbf{1}. \quad (26)$$

Assumption 2 imposes contractivity on the state tube shape when w_k is zero, and is common in literature (Lu and Cannon (2019)). In this work, λ_c will be used to define the terminal cost in Section 4.3.

4. PREDICTED STATE TUBE FOR EXPLORATION

In this section, set-membership identification is formulated as a function of MPC control input variables. Then, a predicted state tube (PST) is constructed to contain state trajectories as a function of predicted uncertainty sets. Finally, the MPC cost function is defined using the PST such that performance-based exploration is induced.

4.1 Predicted set-membership identification

The set-membership identification described in Section 3.1 uses measurements to construct the set Δ_k , and solves linear programs to compute h_{θ_k} . To predict identification that occurs in the future, the sets Δ_k are first predicted as a function of the control inputs. Then, the solutions of the linear programs to compute h_{θ_k} , and thereby the future parameter sets, are predicted as a function of the control inputs using strong duality (Boyd and Vandenberghe (2004)).

Define a sequence of predicted measurements $\{\hat{x}_{l|k}\}_{l=0}^{N_p}$

$$\begin{aligned} \hat{x}_{0|k} &= x_k, & \hat{x}_{l+1|k} &= A(\bar{\theta}_k)\hat{x}_{l|k} + B(\bar{\theta}_k)\hat{u}_{l|k}, \\ \hat{u}_{l|k} &= K(\hat{x}_{l|k} - r_{k+l}) + v_{l|k}, & \forall l \in \mathbb{N}_0^{N_p-1}, \end{aligned} \quad (27)$$

where $N_p \in \mathbb{N}_2^N$ is the user-defined look-ahead horizon. A large N_p improves the modeling of the dual effect inside MPC, but also increases the computational complexity. A prediction of Δ_{k+l} can then be constructed using (27) as

$$\begin{aligned} \hat{\Delta}_{l|k} &:= \{\theta \mid \hat{x}_{i+1|k} - A(\theta)\hat{x}_{i|k} - B(\theta)\hat{u}_{i|k} \in \mathbb{W} \forall i \in \mathbb{N}_{l-\tau}^{l-1}\} \\ &= \{\theta \mid \hat{H}_{\Delta_{l|k}}\theta \leq \hat{h}_{\Delta_{l|k}}\}, \end{aligned} \quad (28)$$

where $\hat{x}_{i|k} = x_{k+i}$, $\hat{u}_{i|k} = u_{k+i}$ when $i < 0$. Note that $\hat{H}_{\Delta_{l|k}}$ and $\hat{h}_{\Delta_{l|k}}$ depend on future control inputs, but the dependence is dropped for clarity of presentation. Using (28), the predicted parameter sets can be defined as

$$\hat{\Theta}_{0|k} := \Theta_k, \quad \hat{\Theta}_{l|k} := \{\theta \mid H_\theta \theta \leq \hat{h}_{\theta_{l|k}}\}, \quad l \in \mathbb{N}_0^N. \quad (29)$$

In (29), $\hat{h}_{\theta_{0|k}} = h_{\theta_k}$, and for $l \in \mathbb{N}_1^{N_p}$, $\hat{h}_{\theta_{l|k}}$ is computed as

$$\begin{aligned} \forall i \in \mathbb{N}_1^{n_\theta}, \quad [\hat{h}_{\theta_{l|k}}]_i &= \max_{\theta} [H_\theta]_i \theta \\ \text{s.t.} \quad H_\theta \theta &\leq \hat{h}_{\theta_{l-1|k}}, \quad \hat{H}_{\Delta_{l|k}}\theta \leq \hat{h}_{\Delta_{l|k}}. \end{aligned} \quad (30)$$

For $l \in \mathbb{N}_{N_p+1}^N$, $\hat{h}_{\theta_{l|k}} = \hat{h}_{\theta_{N_p|k}}$, i.e., no further identification is modeled. Consider the dual reformulation of (30),

$$\hat{h}_{\theta_{l|k}} = \hat{\Psi}_{l|k} \begin{bmatrix} \hat{h}_{\theta_{l-1|k}} \\ \hat{h}_{\Delta_{l|k}} \end{bmatrix}, \quad \hat{\Psi}_{l|k} \begin{bmatrix} H_\theta \\ \hat{H}_{\Delta_{l|k}} \end{bmatrix} = H_\theta, \quad \hat{\Psi}_{l|k} \geq 0, \quad (31)$$

where $\hat{\Psi}_{l|k}$ represent Lagrange multipliers. Using strong duality, it can be seen that there exist $\hat{\Psi}_{l|k}$ such that (31) holds. Thus, using $\hat{\Psi}_{l|k}$ as MPC optimization variables, (30) can be replaced by (31). A similar dual adaptive MPC method has been proposed for HT in Parsi et al. (2022), but it approximates set-membership online using an intersection of Θ_k and $\hat{\Delta}_{l|k}$. This results in MPC overestimating the capabilities of identification.

Note that (31) includes the product of $\hat{\Psi}_{l|k}$ with other optimization variables, resulting in bilinear constraints and a nonconvex MPC optimization problem.

4.2 Predicted state tubes

The effect of the future identification on the state trajectories can be estimated by defining a predicted state tube $\{\hat{\mathbb{X}}_{l|k}\}_{l=0}^N$ such that $x_k \in \hat{\mathbb{X}}_{0|k}$ and $\forall l \in \mathbb{N}_0^{N-1}$

$$\hat{\mathbb{X}}_{l+1|k} \ni A(\theta)x + B(\theta)u(x) + w, \quad \forall x \in \hat{\mathbb{X}}_{l|k}, \theta \in \hat{\Theta}_{l|k}, \forall w \in \mathbb{W}. \quad (32)$$

It can be seen from (13) that $\forall l \in \mathbb{N}_0^{N-1}$, $\hat{\mathbb{X}}_{l|k} \subseteq \mathbb{X}_{l|k}$ because $\hat{\Theta}_{l|k} \subseteq \Theta_k$ by construction (30).

Homothetic tubes. In the HT parameterization, the PST is represented by polytopes of the form

$$\hat{\mathbb{X}}_{l|k}^H = \{\hat{z}_{l|k}\} \oplus \hat{\alpha}_{l|k} \mathbb{X}_0, \quad \forall l \in \mathbb{N}_0^N, \quad (33)$$

where $\hat{z}_{l|k} \in \mathbb{R}^{n_x}$ and $\hat{\alpha}_{l|k} \in \mathbb{R}_{\geq 0}$ are translation and scaling variables. Then, the following is a direct extension of Proposition 1.

Corollary 1.1. The PST dynamics (32) are satisfied under the HT parameterization (33) if $\forall j \in \mathbb{N}_1^q$, $l \in \mathbb{N}_0^{N-1}$, there exist $\hat{\Lambda}_{l|k}^{H,j} \in \mathbb{R}^{n_x \times n_\theta}$ such that

$$\begin{aligned} -H_x \hat{z}_{0|k} - \hat{\alpha}_{0|k} \mathbf{1} &\leq -H_x x_k, & H_x \hat{D}_{l|k}^j &= \hat{\Lambda}_{l|k}^{H,j} H_\theta, \\ \hat{\Lambda}_{l|k}^{H,j} \hat{h}_{\theta_{l|k}} + H_x \hat{d}_{l|k}^j - \hat{\alpha}_{l+1|k} \mathbf{1} &\leq -\bar{w}, \end{aligned} \quad (34)$$

where $\hat{D}_{l|k}^j$ and $\hat{d}_{l|k}^j$ are predictions of $D_{l|k}^j$ and $d_{l|k}^j$ defined in (16).

Flexible tubes. In the FT parameterization the PST is represented by polytopes of the form

$$\hat{\mathbb{X}}_{l|k}^F = \{x \mid H_x(x - r_{k+l}) \leq \hat{h}_{x_{l|k}}\}, \quad \forall l \in \mathbb{N}_0^N, \quad (35)$$

where $\hat{h}_{x_{l|k}} \in \mathbb{R}^{n_x}$ are optimization variables. Then, the following is a direct extension of Proposition 2.

Corollary 2.1. The PST dynamics (32) are satisfied under the FT parameterization (35) if $\forall i \in \mathbb{N}_1^{n_x}$, $j \in \mathbb{N}_1^{n_\theta}$ and $l \in \mathbb{N}_0^{N-1}$

$$\begin{aligned} H_x(x_k - r_k) &\leq \hat{h}_{x_{0|k}}, \\ [1 \ \hat{\theta}_{l|k}^T] \Lambda_{2,l|k}^{F,i} (\hat{h}_{x_{l|k}} + H_x r_{k+l}) - [H_x]_i r_{k+l+1} - \\ &[H_x]_i (B(\hat{\theta}_{l|k}^j)(K r_{k+l} - v_{l|k})) + \bar{w} \leq [\hat{h}_{x_{l+1|k}}]_i, \end{aligned} \quad (36)$$

where $\Lambda_{2,l|k}^{F,i}$ are computed offline as defined in (21).

In (36), $\hat{\theta}_{l|k}^j$ can be formulated as linear combination of terms in $\hat{h}_{\theta_{l|k}}$ because $\hat{\Theta}_{l|k}$ is restricted to hyperboxes.

4.3 MPC Cost function

By design, the family of trajectories defined by the PST is a function of the input variables via the predicted parameter sets. That is, informativity of the input can be directly related to the size of the PST. This feature is now used to define the cost function, using the stage cost $J_S(\hat{\mathbb{X}}_{l|k}, v)$ and terminal cost $J_T(\hat{\mathbb{X}}_{N|k}, k)$

$$\begin{aligned} J_S(\hat{\mathbb{X}}_{l|k}, v_{l|k}) &= \\ \max_{x \in \hat{\mathbb{X}}_{l|k}} &\|Q(x - r_{k+l})\|_\infty + \|R(u_{l|k}(x) - \bar{u}_{l|k})\|_\infty, \\ J_T(\hat{\mathbb{X}}_{N|k}, k) &= \beta(k + N) \hat{J}_S(\hat{\mathbb{X}}_{N|k}, \bar{u}_{N|k}). \end{aligned} \quad (37)$$

where $\beta(j) = \frac{1 - \lambda_e^{T-j}}{1 - \lambda_e}$ is a time dependent scaling factor, which reduces the incentive for exploration towards the end of the control task. Overall, the terminal cost is estimating the worst-case cost after the prediction horizon assuming that the reference is fixed and no disturbances act on the system. The cost (37) induces performance-based exploration, as exploratory actions are used only when an improvement in performance is predicted.

Homothetic tubes. In the HT formulation, the vertices of the state tube are formulated as linear functions of MPC optimization variables. Thus (37) can be simplified to

$$J_S^H(\hat{\mathbb{X}}_{l|k}^H, v) = \max_{j \in \mathbb{N}_1^q} \left\| Q(\hat{x}_{l|k}^j - r_{k+l}) \right\|_\infty + \left\| R(\hat{u}_{l|k}^j - \bar{u}_{l|k}) \right\|_\infty, \quad (38)$$

which can be written as a linear cost function using an epigraph reformulation (Boyd and Vandenberghe (2004)). Such a reformulation is possible as the ∞ -norm is used to define the cost in (37). Instead, if the cost is defined using quadratic functions, second-order cone constraints can be used to reformulate the worst-case cost.

Flexible tubes. In the FT formulation, (37) cannot be directly reformulated into a linear cost function in the optimization variables. Hence, the maximization over the state and input terms is separated, to obtain

$$J_S^F(\hat{\mathbb{X}}_{l|k}^F, v) = \max_{x \in \hat{\mathbb{X}}_{l|k}^F} \|Q(x - r_{k+l})\|_\infty + \max_{x \in \hat{\mathbb{X}}_{l|k}^F} \|R(u_{l|k}(x) - \bar{u}_{l|k})\|_\infty, \quad (39)$$

which can then be reformulated as a linear cost. Note that the (39) overestimates the stage cost defined in (37). This can potentially lead to suboptimal exploratory actions, but such approximations are also used in other FT methods in literature (Lu and Cannon (2019)) to simplify the resulting online optimization problem.

5. DUAL ADAPTIVE MPC

In this section, HT and FT dual adaptive MPC algorithms are presented using the ingredients described in Sections 3 and 4.

Homothetic tubes. In the HT formulation, the optimization variables are

$$\gamma_k^H = \left\{ \{z_{l|k}, \alpha_{l|k}, v_{l|k}, \bar{u}_{l|k}, \hat{z}_{l|k}, \hat{\alpha}_{l|k}, \{\Lambda_{l|k}^{H,j}\}_{j=1}^q\}_{l=0}^N, \right. \\ \left. \{ \{\hat{\Lambda}_{l|k}^{H,j}\}_{j=1}^q\}_{l=0}^{N-1}, \{ \hat{\Psi}_{l|k}, \hat{h}_{\theta_{l|k}}\}_{l=1}^{N_p} \right\}, \quad (40)$$

and the online optimization problem is

$$\min_{\gamma_k^H} J_{\mathcal{T}}^H(\hat{\mathbb{X}}_{N|k}^H, k) + \sum_{l=0}^{N-1} J_S^H(\hat{\mathbb{X}}_{l|k}^H, v_{l|k}) \quad (41) \\ \text{s.t.} \quad (12), (17), (24), (27), (31), (34).$$

The dual adaptive MPC algorithm using HT is described in Algorithm 1, where $\hat{u}_{k,0}^*$ is the optimal value of $\hat{u}_{k,0}$.

Algorithm 1 Homothetic dual adaptive MPC

Offline:

- 1: Choose $\tau, N, N_p, \bar{\theta}_0$
- 2: Design K, H_x satisfying Assumptions 1, 2

Online: At each time step $k \geq 0$:

- 1: Obtain the measurement x_k
 - 2: Compute Θ_k and $\bar{\theta}_k$ using set-membership and (10)
 - 3: Solve (41)
 - 4: Apply $\hat{u}_{k,0}^*$ to the system
-

Flexible tubes. The optimization variables in the FT formulation are

$$\gamma_k^F = \left\{ \{h_{x_{l|k}}, \hat{h}_{x_{l|k}}, \bar{u}_{l|k}, v_{l|k}\}_{l=0}^N, \{ \hat{\Psi}_{l|k}, \hat{h}_{\theta_{l|k}}\}_{l=1}^{N_p} \right\}, \quad (42)$$

and the online optimization problem is

Algorithm 2 Flexible dual adaptive MPC

Offline:

- 1: Choose $\tau, N, N_p, \bar{\theta}_0$ and μ
- 2: Design K, H_x satisfying Assumptions 1 and 2
- 3: Compute $\Lambda_{1,l|k}^{F,i}, \Lambda_{2,l|k}^{F,i}$ for $i \in \mathbb{N}_1^{n_x}, l \in \mathbb{N}_0^{N-1}$ from (20) and (21)

Online: At each time step $k \geq 0$:

- 1: Obtain the measurement x_k
 - 2: Compute Θ_k and $\bar{\theta}_k$ using set-membership and (10)
 - 3: Solve (43)
 - 4: Apply $\hat{u}_{k,0}^*$ to the system
-

$$\min_{\gamma_k^F} J_{\mathcal{T}}^F(\hat{\mathbb{X}}_{N|k}^F, k) + \sum_{l=0}^{N-1} J_S^F(\hat{\mathbb{X}}_{l|k}^F, v_{l|k}) \quad (43) \\ \text{s.t.} \quad (12), (19), (25), (27), (31), (36).$$

The dual adaptive MPC algorithm using flexible tubes is described in Algorithm 2.

Remark 2. In simulations, it was observed that (43) is difficult to solve. One possible reason for this was the recursive computation of $\hat{\Theta}_{l|k}$ in (31) resulting in a large number of bilinearities. A relaxation which worked well in practice was to compute $\hat{\Theta}_{l|k}$ using Θ_k and all the predicted measurements in $\{\hat{x}_{i|k}\}_{i=-\tau}^l$, resulting in the following approximation of (31)

$$\hat{h}_{\theta_{l|k}} = \hat{\Psi}_{l|k} \begin{bmatrix} h_{\theta_k} \\ \tilde{h}_{\Delta_{l|k}} \end{bmatrix}, \quad \hat{\Psi}_{l|k} \begin{bmatrix} H_\theta \\ \tilde{H}_{\Delta_{l|k}} \end{bmatrix} = H_\theta, \quad \hat{\Psi}_{l|k} \geq 0, \quad (44)$$

where $\tilde{h}_{\Delta_{l|k}}$ and $\tilde{H}_{\Delta_{l|k}}$ are defined according to

$$\{\theta \mid \hat{x}_{i+1|k} - A(\theta)\hat{x}_{i|k} - B(\theta)\hat{u}_{i|k} \in \mathbb{W}, \forall i \in \mathbb{N}_{-\tau}^{l-1}\} \\ = \{\theta \mid \tilde{H}_{\Delta_{l|k}}\theta \leq \tilde{h}_{\Delta_{l|k}}\}. \quad (45)$$

Although such a relaxation is an approximation of the set-membership in closed loop, it was required only for the FT case, for which active exploration methods have not been proposed in the literature.

5.1 Properties

Theorem 1. Let Assumptions 1, 2 hold. If the optimization problem (41) (or problem (43)) is feasible at time $k = 0$, then it remains feasible for all $k > 0$. Moreover, the closed-loop system formed by any dynamics (1)-(4) and the MPC controller satisfies (6) for all $k > 0$.

Proof. The proof for the FT case is given here. The proof for the HT case is similar, and can be found in Liu (2022a).

Let (43) be feasible at some time step k . Before (43) is solved at $k+1$, $\bar{\theta}_k$ and Θ_k are updated to $\bar{\theta}_{k+1}$ and Θ_{k+1} respectively, and the reference trajectory is shifted by one time step. It will now be shown that a feasible solution to (43) can be constructed using the solution at time k .

The modification of $\bar{\theta}_k$ affects only the cost function (through the PST and the input setpoints), and does not affect the feasibility of (43). Moreover, $\hat{\mathbb{X}}_{l|k}^F \subseteq \mathbb{X}_{l|k}^F$ and the RST also satisfies the PST dynamics (36). Therefore, one only needs to find a feasible solution for the RST (defined by the variables $\{h_{x_{l|k+1}}\}_{l=0}^{N-1}, \{v_{l|k+1}\}_{l=0}^{N-1}$) and the terminal set (defined by $v_{N|k+1}$ and $h_{x_{N|k+1}}$).

Consider the variables

$$v_{l-1|k+1} = v_{l|k}, \quad \forall l \in \mathbb{N}_1^{N-1}; \quad v_{N-1|k+1} = v_{N|k+1} = v_{N|k},$$

which ensure that $u_{l-1|k+1}(x) = u_{l|k}$ for $l \in \mathbb{N}_1^N$. Because the updated parameter set satisfies $\Theta_{k+1} \subseteq \Theta_k$ from (7), a shifted RST from time k also satisfies (19) at $k+1$.

Next, the invariance of the terminal set is used to define a feasible $h_{x_{N|k+1}}$, so that the terminal set remains the same. This is achieved by choosing

$$h_{x_{N|k+1}} = h_{x_{N|k}} + H_x r_{k+N} - H_x r_{k+N+1}, \quad (46)$$

which is a feasible solution for the terminal conditions (25).

Therefore, (43) is feasible at time $k+1$, and by induction, it is recursively feasible. Moreover, the constraints (6) are satisfied by design, because $\theta^* \in \Theta_k$.

To the best of the authors' knowledge, FT methods have been proposed only for regulation tasks (i.e., reference is the origin) in literature. Thus, in addition to the novelty of the dual control formulation, Algorithm 2 provides a novel design methodology for recursively feasible tracking MPC controllers using an FT formulation.

6. NUMERICAL EXAMPLE

In this section the performance of the proposed algorithms will be compared to each other, and existing methods from the literature. Four different controllers will be used in the simulations, as described in Table 1. The abbreviation ESM stands for exact set-membership, ASM stands for approximate set-membership, and P (in FT-P) is used to denote passive exploration. Note that Lu and Cannon (2019) design a controller for a regulation task, but the tracking formulation was added to FT-P for a meaningful comparison. The code to simulate the example is available in a public repository (Liu (2022b)).

The model of the system is parameterized by the matrices

$$A_0 = \begin{bmatrix} 0.9 & 0.5 \\ 0.2 & 0.8 \end{bmatrix}, \quad A_1 = \begin{bmatrix} 0.1 & 0 \\ 0 & 0.2 \end{bmatrix}, \quad A_2 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, \\ B_0 = \begin{bmatrix} 1 & 0.5 \\ 0.2 & 0.775 \end{bmatrix}, \quad B_1 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, \quad B_2 = \begin{bmatrix} 0 & 0.2 \\ 0 & 0.35 \end{bmatrix}.$$

The uncertainty set is $\Theta = \{\theta \in \mathbb{R}^2 \mid \|\theta\|_\infty \leq 1.2\}$. The constraint set \mathbb{Z} is defined as

$$\mathbb{Z} = \{(x, u) \in \mathbb{R}^{2 \times 2} \mid \|x\|_\infty \leq 3, \|u\|_\infty \leq 2\},$$

and the disturbance set is $\mathbb{W} := \{w \in \mathbb{R}^2 \mid \|w\|_\infty \leq 0.1\}$. The cost function is defined by $Q = 4I_2$ and $R = I_2$.

The parameters for all the controllers are chosen to be $N = 8$, $N_p = 5$, $\tau = 1$ and $\bar{\theta}_0 = [0.1, 0.1]^T$. A piecewise constant reference trajectory with five different setpoints is to be tracked, with a simulation length of $T=100$ time steps. The simulations are performed for 500 different realizations of the true parameter $\theta^* \in \Theta$ and disturbance sequences $w_k \in \mathbb{W}$. The parameters and disturbances are chosen randomly from a uniform distribution. All the optimization problems were implemented using YALMIP and solved using MOSEK and IPOPT on an Intel Xeon Gold 5118 processor with 2GB RAM. The FT-P method

Table 1. Controllers used for simulation.

Controller	Description
HT-ESM	Algorithm 1
FT-ESM	Algorithm 2
HT-ASM	Parsi et al. (2022)
FT-P	Lu and Cannon (2019)




Fig. 1. Distribution of closed loop costs over 500 random realizations of θ^* and w_k .

solves a linear program, whereas all the dual controllers solve nonconvex optimization problems. The average computation times to solve the online MPC optimization at each time step in FT-P and FT-ESM methods are 4.14s and 6.77s, and that in HT-ASM and HT-ESM methods are 6.05s and 10.54s respectively. Although FT-P and HT-ASM require lower computation times, they result in worse performance.

The performance of the controllers is compared in Figure 1, where the distribution of closed-loop costs achieved by each controller is shown. First, it can be seen that the FT-P controller results in higher costs compared to the dual controllers, because the exploration is passive. Such a result has been demonstrated for HT methods in Parsi et al. (2022), and now is replicated for FT methods. Second, the use of exact set-membership reformulation in HT-ESM results in lower costs compared to HT-ASM, because the latter overestimates the information that can be obtained from dual control. Third, it can be seen that both the HT methods perform better compared to the FT-ESM method. This is because the FT method computes Lagrange multipliers offline, uses the approximation in (44) and simplifies the cost function in (39) so that the optimization problem is tractable. Thus, although the FT parameterization allows for efficient computation, it results in a loss in performance.

It must be noted that the large variation in closed-loop costs for a given controller is because the true parameter affects the reference tracking performance. Moreover, the ESM methods are susceptible to high costs in cases when the initial guess $\bar{\theta}_0$ is far from the true parameter θ . It can be seen that the number of such cases is low, and future work will focus on reducing the dependence on $\bar{\theta}_0$.

Figure 2 shows the state trajectories, input trajectories and the bounds on the parameters for one specific realization of the true parameter and disturbance sequence. This simulation was selected among the 500 realizations to demonstrate the advantages of the proposed ESM methods. All the controllers used in the simulation ensure constraint satisfaction. The FT-P controller initially does not use the inputs due to the large uncertainty sets and lack of incentive to explore. This results in a large tracking error until passive exploration reduces the uncertainty. In contrast, exploratory actions are used by all the dual controllers, as seen in the deviation of the trajectories from the setpoint between time steps 12 and 20. This improves the overall closed loop costs. The FT-P controller

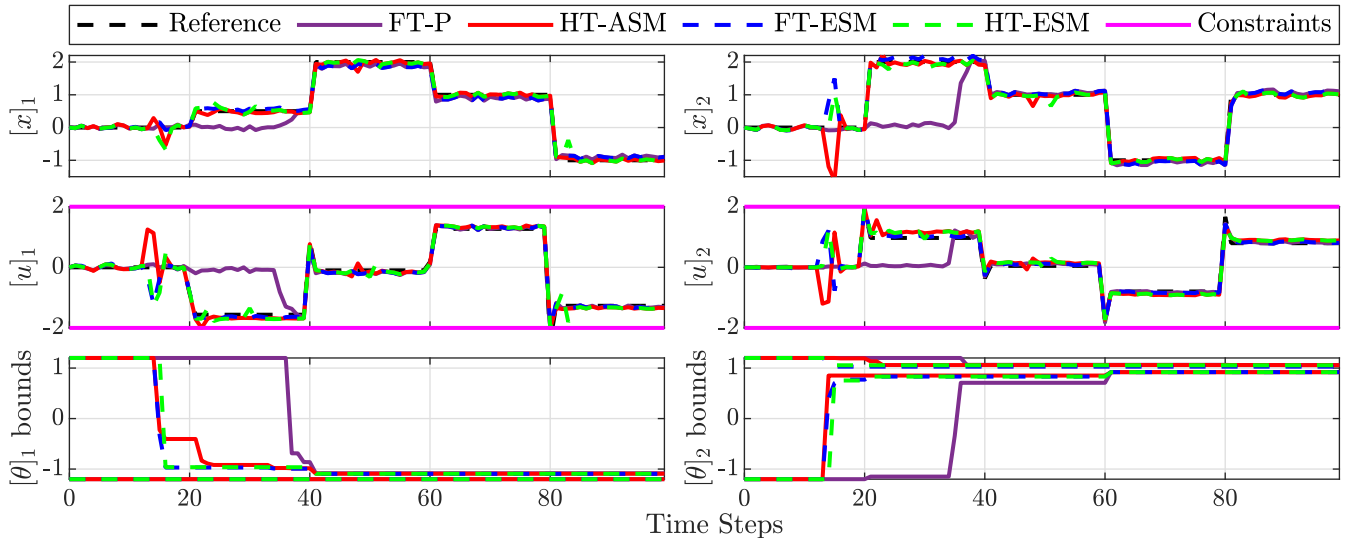


Fig. 2. State and input trajectories of the system under four different controllers, along with upper and lower bounds on parameters. The true parameter is $\theta^* = [-1.16, 0.96]$.

has a cost of 184.8, whereas the cost is 63.0 for FT-ESM, 47.3 for HT-ASM, and 43.5 for HT-ESM. Moreover, the HT-ASM controller has worse identification performance compared to HT-ESM (see time steps 10-25), owing to the approximation of set-membership equations.

7. CONCLUSION

In this work, a performance-based dual control framework is proposed to navigate the dual control trade-off in optimal tracking problems. This framework has been demonstrated with two popular tube MPC formulations used in the adaptive MPC literature. Dual control is achieved by an exact reformulation of set-membership identification within MPC optimization. The proposed method ensures recursive feasibility and robust constraint satisfaction, and numerical simulations demonstrate performance improvement compared to existing methods. Future work involves reducing the dependence of the dual effect on the initial parameter estimate, and extending the method to nonlinear adaptive MPC algorithms.

REFERENCES

- Boyd, S. and Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press.
- Bujarbaruah, M., Zhang, X., Tanaskovic, M., and Borrelli, F. (2020). Adaptive stochastic MPC under time-varying uncertainty. *IEEE Transactions on Automatic Control*, 66(6), 2840–2845.
- Feldbaum, A. (1960-61). Dual control theory. I-IV. *Avtomatika i Telemekhanika*, 21(9), 1240–1249.
- Hewing, L., Wabersich, K.P., Menner, M., and Zeilinger, M.N. (2020). Learning-based model predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems*, 3.
- Klenske, E.D. and Hennig, P. (2016). Dual control for approximate Bayesian reinforcement learning. *The Journal of Machine Learning Research*, 17(1), 4354–4383.
- Köhler, J., Andina, E., Soloperto, R., Müller, M.A., and Allgöwer, F. (2019). Linear robust adaptive model predictive control: Computational complexity and conservatism. In *58th Conference on Decision and Control (CDC)*, 1383–1388. IEEE.
- Köhler, J., Köttling, P., Soloperto, R., Allgöwer, F., and Müller, M.A. (2021). A robust adaptive model predictive control framework for nonlinear uncertain systems. *International Journal of Robust and Nonlinear Control*, 31(18), 8725–8749.
- Kouvaritakis, B. and Cannon, M. (2015). *Model Predictive Control: Classical, Robust and Stochastic*. Springer International Publishing.
- Liu, D. (2022a). *Dual adaptive MPC using an exact reformulation of set-membership identification*. Master’s thesis, ETH Zurich, Switzerland. doi:10.3929/ethz-b-000578724.
- Liu, D. (2022b). Supplementary material for the conference paper titled “Dual adaptive MPC using an exact set-membership reformulation”. doi:10.3929/ethz-b-000578725.
- Lorenzen, M., Cannon, M., and Allgöwer, F. (2019). Robust MPC with recursive model update. *Automatica*, 103, 461–471.
- Lu, X. and Cannon, M. (2019). Robust adaptive tube model predictive control. In *2019 American Control Conference (ACC)*, 3695–3701. IEEE.
- Mesbah, A. (2018). Stochastic model predictive control with active uncertainty learning: A survey on dual control. *Annual Reviews in Control*, 45, 107–117.
- Milanese, M. and Vicino, A. (1991). Optimal estimation theory for dynamic systems with set membership uncertainty: an overview. *Automatica*, 27(6), 997–1009.
- Parsi, A., Iannelli, A., and Smith, R.S. (2022). An explicit dual control approach for constrained reference tracking of uncertain linear systems. *IEEE Transactions on Automatic Control*. doi:10.1109/TAC.2022.3176800.
- Soloperto, R., Köhler, J., and Allgöwer, F. (2020). Augmenting MPC schemes with active learning: Intuitive tuning and guaranteed performance. *IEEE Control Systems Letters*, 4(3), 713–718.