

Policy Gradient Reinforcement Learning for Uncertain Polytopic LPV Systems based on MHE-MPC

Hossein Nejatbakhsh Esfahani * Sébastien Gros *

* Department of Engineering Cybernetics, Norwegian University of Science and Technology (NTNU), Trondheim, Norway (e-mail: {hossein.n.esfahani, sebastien.gros}@ntnu.no).

Abstract: In this paper, we propose a learning-based Model Predictive Control (MPC) approach for the polytopic Linear Parameter-Varying (LPV) systems with inexact scheduling parameters (as exogenous signals with inexact bounds), where the Linear Time Invariant (LTI) models (vertices) captured by combinations of the scheduling parameters becomes wrong. We first propose to adopt a Moving Horizon Estimation (MHE) scheme to simultaneously estimate the convex combination vector and unmeasured states based on the observations and model matching error. To tackle the wrong LTI models used in both the MPC and MHE schemes, we then adopt a Policy Gradient (PG) Reinforcement Learning (RL) to learn both the estimator (MHE) and controller (MPC) so that the best closed-loop performance is achieved. The effectiveness of the proposed RL-based MHE/MPC design is demonstrated using an illustrative example.

Copyright © 2022 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: Model Predictive Control, Moving Horizon Estimation, Reinforcement Learning, Polytopic LPV, Multi-Model Linear System

1. INTRODUCTION

Model predictive control (MPC) is an advanced control approach and widely used to many control problems due to its capability to handle both input and state constraints by making explicit use of the process model to predict the future evolution of the system (Rawlings et al., 2017). However, in most applications some of the model parameters are uncertain at design time so that the prediction model may not be accurate enough to capture the true dynamics of a real system that this issue can affect the closed-loop performance (Cairano, 2012).

In many real applications of MPC, an accurate model of the (possibly nonlinear) real system may not be available so that the simplified linear models are commonly adopted as the prediction models, which are captured at a specific operation point of the real system. However, these identified models may not work for some uncertain systems where some parameters of the dynamics vary or not be known perfectly at different operation points. Therefore, the Linear Parameter-Varying (LPV) framework is usually adopted to address the issue above in which the linear models with parametric uncertainties (i.e, exogenous scheduling parameters) are represented by the LPV models (Leith and Leithead, 2000).

The use of polytopic LPV prediction models has been proven to be an effective approach to develop the MPC schemes for both the linear and nonlinear systems, e.g., see (Pour et al., 2017; Ding and Pan, 2016; Pipino and Adam, 2019). In most of the MPC-based polytopic LPVs, a multi-model approach is used to represent a convex combination of the vertices as Linear Time Invariant (LTI) models in which the uncertainty (exogenous/endogenous scheduling parameter with known bounds) is associated to the convex combination vector. However, if the

scheduling parameters are not known and constantly change, the robust MPC schemes, i.e, a tube MPC can be used (Langson et al., 2004). As an alternative approach, as we consider in this paper, this combination vector is assumed to be constant or changing slowly, which can be estimated by a parameter estimation scheme (Di Cairano and Danielson, 2021; Zhou et al., 2017). As an interesting method to concurrently estimate this vector and some unmeasured states, a Moving Horizon Estimation (MHE) was adopted in (Pipino et al., 2020). However, the LTI models used as vertices of the polytopic LPV may not be exact due to the inexact selection of the scheduling parameter bounds and then, these wrong models used in both the MHE and MPC schemes can decrease the closed-loop performance. As a solution to address this problem, a learning mechanism, i.e, a Reinforcement Learning (RL) can be leveraged to learn these schemes with the aim of increasing the closed-loop performance.

Reinforcement Learning (RL) is a powerful tool for solving Markov Decision Processes (MDP) problems (Sutton and Barto, 2018), which has been recently combined with the MPC and MHE schemes in order to improve their performance where the adopted prediction/estimation models may not capture the real system dynamics perfectly, e.g., see (Esfahani et al., 2021b; Kordabad et al., 2021; Cai et al., 2021).

Motivated by this observation in the context of the MPC-based RL, in this paper, we propose to use an MHE scheme to deliver the convex combination vector needed in the MPC scheme and deal with the polytopic LPV systems with possibly inexact vertices (wrong LTI models) by adopting a Deterministic Policy Gradient (DPG) RL to adjust a certain number of the parameters in the parameterized MHE and MPC schemes. It is worth highlighting that the proposed *compatible deterministic policy gra-*

Reinforcement learning based on a combined MHE-MPC scheme is a central contribution in this paper. Although this learning-based (state-parameter) estimator/controller scheme is developed and employed in the context of the LPV systems in this paper, the proposed algorithm can be readily used to control many other applications, i.e., partially observable systems in which a combined estimator/observer-MPC is needed.

In what follows, the polytopic LPV model is introduced in Section 2, and the parameterized MHE and MPC schemes are detailed in Section 3. The MHE/MPC-based policy gradient reinforcement learning is then developed in Section 4. Simulations on a mass-spring-damper case study are reported in Section 5, and finally, conclusions and future work are given in Section 6.

2. UNCERTAIN POLYTOPIC LPV

In this paper, we focus on the classic (pure) LPV systems, in which the uncertain dynamical model of the system can be described as follows:

$$\mathbf{x}_{k+1} = A(\boldsymbol{\rho}_k) \mathbf{x}_k + B(\boldsymbol{\rho}_k) \mathbf{u}_k \quad (1a)$$

$$\mathbf{y}_k = C(\boldsymbol{\rho}_k) \mathbf{x}_k \quad (1b)$$

where the scheduling parameters $\boldsymbol{\rho} \in \mathcal{P} \in \mathbb{R}^p$ are considered as some exogenous signals applied to the real system. The vectors $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{u} \in \mathbb{R}^m$ are states and control inputs, respectively. The system measurements is labeled by $\mathbf{y}_k \in \mathbb{R}^e$. Hence, the system (1) can be represented by a polytopic approximation $[A(\boldsymbol{\rho}_k), B(\boldsymbol{\rho}_k), C(\boldsymbol{\rho}_k)] \in \Omega$, where a polytope Ω is defined as a convex hull upon the Linear Time Invariant (LTI) models as follows:

$$\Omega = Co\{[A_1, B_1, C_1], [A_2, B_2, C_2], \dots, [A_\ell, B_\ell, C_\ell]\} \quad (2)$$

where $\ell = 2^p$ is the number of vertices and $Co\{\cdot\}$ reads a convex hull. Indeed, the LTI matrices are delivered by combinations of the scheduling parameters at their bounds $\underline{\boldsymbol{\rho}}, \bar{\boldsymbol{\rho}}$. We then evaluate the variable matrices in (1) as a convex combination of the LTI models (a.k.a vertices) as follows:

$$A(\boldsymbol{\rho}_k) = \sum_{i=1}^{\ell} \beta_i A_i, \quad B(\boldsymbol{\rho}_k) = \sum_{i=1}^{\ell} \beta_i B_i, \quad (3a)$$

$$C(\boldsymbol{\rho}_k) = \sum_{i=1}^{\ell} \beta_i C_i \quad (3b)$$

where A_i, B_i, C_i are known matrices of appropriate size and $\beta_i, i = 1 \dots \ell$ is the unknown weight of each vertex (a.k.a convex combination vector) such that:

$$\chi = \left\{ \boldsymbol{\beta} \in \mathbb{R}^\ell : \sum_{i=1}^{\ell} \beta_i = 1, \quad 0 \leq \beta_i \leq 1 \right\} \quad (4)$$

Assumption 1. By (Di Cairano, 2016), the convex combination vector $\boldsymbol{\beta}_k$ is assumed to be constant or changing slowly with respect to the dynamics and there exists an estimator to compute the estimate of this vector such that $\hat{\boldsymbol{\beta}}_k \in \chi$ for all $k \in \mathbb{Z}_{0+}$.

Note that the stability and constraint satisfaction of the MPC schemes for the polytopic LPV systems under Assumption 1 was detailed and established in (Di Cairano and Danielson, 2021) and by Lemma 2 and Theorem 1 in (Di Cairano, 2016). We then focus on this kind of the LPV system where its vertices may not be known and exact. As discussed, the vertices in a polytopic LPV may have a wrong model because of the inexact bounds of the scheduling parameters, some parametric

uncertainties and the linearization approximation effects. Then, the wrong LTI models in the polytope Ω can be introduced by the inexact vertices as follows:

$$\hat{\Omega} = Co\left\{[\hat{A}_1, \hat{B}_1, \hat{C}_1], [\hat{A}_2, \hat{B}_2, \hat{C}_2], \dots, [\hat{A}_\ell, \hat{B}_\ell, \hat{C}_\ell]\right\} \quad (5)$$

In the context of MPC schemes for the polytopic LPVs, under Assumption 1, one needs to evaluate an estimation of $\beta_i, i = 1 \dots \ell$ (Di Cairano and Danielson, 2021; Pipino et al., 2020). In this paper we propose to adopt an MHE scheme to deliver these unknown weights while some unmeasured states are estimated. However, the inexact vertices in $\hat{\Omega}$ can affect parameter/state estimations, the policy captured from the MPC and the closed-loop performance, subsequently. In addition to this inexact polytope issue, the partially measurable states can be challenging since the MHE scheme as an estimator of $\beta_{i,k}$ needs to capture all states as measurements and minimize a distance as a model matching error between these measurements and those states delivered by the uncertain LPV model.

To cope with these problems, we propose to parameterize both the MHE and MPC schemes and adopt a reinforcement learning to adjust them in order to achieve the best closed-loop performance even when using some wrong LTI models in a polytopic LPV system. More precisely, it is possible to generate the optimal policy based on a wrong model by learning the MHE and MPC schemes.

3. ADJUSTABLE MHE-MPC IN A LPV FRAMEWORK

Under Assumption 1, an obvious choice for the combination vectors used in the MPC scheme would be $\beta_{j|k} = \beta_k$, for all $j \in \mathbb{Z}_{k, \dots, k+N}$. As discussed in (Di Cairano, 2016), this choice could not be proper because if the entire β -parameter prediction vector suddenly changes, the MPC value function may not be decreasing. Hence, under Assumption 1, one can select a N -step delay in this parameter prediction as follows:

$$\beta_{j|k} = \hat{\beta}_{k-N+j}, \quad \forall j \in \mathbb{Z}_{0, \dots, N} \quad (6)$$

We then adopt an MHE-based estimator to deliver these weights needed in an MPC scheme. Let us formulate a parameterized MHE scheme as β -parameter and state estimator, where the vertices of the polytopic LPV are inexact and the real system states are not fully measurable.

$$\left\{ \hat{\boldsymbol{\beta}}_{k-N \dots k}, \hat{\mathbf{x}}_{k-N \dots k} \right\} = \arg \min_{\boldsymbol{\beta}, \mathbf{x}} \gamma^N \|\mathbf{x}_{k-N} - \tilde{\mathbf{x}}_{k-N}\|_{A_\theta}^2$$

$$+ \sum_{j=k-N+1}^k \gamma^{k-j} \left(\|\boldsymbol{\mu}_j\|_{R_\theta}^2 + \|\Delta\beta_{j|k}\|^2 \right)$$

$$\text{s.t. } \mathbf{x}_{j+1} = \hat{A}(\boldsymbol{\rho}_j) \mathbf{x}_j + \hat{B}(\boldsymbol{\rho}_j) \bar{\mathbf{u}}_j, \quad (7a)$$

$$\hat{A}(\boldsymbol{\rho}_j) = \sum_{i=1}^{\ell} \beta_{j,i} \hat{A}_i, \quad \hat{B}(\boldsymbol{\rho}_j) = \sum_{i=1}^{\ell} \beta_{j,i} \hat{B}_i, \quad (7b)$$

$$\bar{\mathbf{y}}_j = \left(\sum_{i=1}^{\ell} \beta_{j,i} \hat{C}_i \right) \mathbf{x}_j + \boldsymbol{\mu}_j, \quad (7c)$$

$$\sum_{i=1}^{\ell} \beta_{j,i} = 1, \quad 0 \leq \beta_{j,i} \leq 1 \quad (7d)$$

$$\Delta\beta_{j|k} = \beta_{j|k} - \beta_{j|k-1}, \quad \beta_j = Col\{\beta_{j,1}, \dots, \beta_{j,\ell}\} \quad (7e)$$

where $\gamma \in \mathbb{Z}_{(0,1)}$ is a discount factor and $\bar{\mathbf{y}}_j, \bar{\mathbf{u}}_j$ are the measurements available at the physical time k . We label R_θ as

an adjustable weight matrix penalizing the deviation between some estimated states captured from the LPV model and their corresponding measurements. The first term in the above least squares problem is an arrival cost weighted with matrix A_θ , which aims at approximating the information prior to $k - N_{\text{MHE}}$, where $\tilde{\mathbf{x}}$ is the available estimation $\hat{\mathbf{x}}_{k-N}$ at time $k-1$. We then detail a policy gradient RL in order to adjust the combination vectors β_i , where the LTI models (vertices of the polytope) are wrong.

This paper proposes a learning-based estimation and control for a polytopic LPV system with possibly inexact LTI models at the vertices, where the control policy is delivered by a parameterized MPC scheme as follows:

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{u}, \boldsymbol{\sigma}} \quad & \gamma^{k+N} (T_\theta(\mathbf{x}_{k+N}, \boldsymbol{\beta}_{k+N}) + \mathbf{w}_f^\top \boldsymbol{\sigma}_{k+N}) \\ & + \sum_{j=k}^{k+N-1} \mathcal{G}_\theta(\mathbf{x}_j, \mathbf{u}_j) + \sum_{j=k}^{k+N-1} \gamma^j (l_\theta(\mathbf{x}_j, \mathbf{u}_j) + \mathbf{w}^\top \boldsymbol{\sigma}_j) \end{aligned} \quad (8a)$$

$$\text{s.t.} \quad \mathbf{x}_{j+1} = \hat{A}(\boldsymbol{\rho}_j) \mathbf{x}_j + \hat{B}(\boldsymbol{\rho}_j) \mathbf{u}_j, \quad (8b)$$

$$\mathbf{x}_k = \hat{\mathbf{x}}_k, \quad \beta_{j|k} = \hat{\beta}_{j-N} \quad (8c)$$

$$\hat{A}(\boldsymbol{\rho}_j) = \sum_{i=1}^{\ell} \beta_{j,i} \hat{A}_i, \quad (8d)$$

$$\hat{B}(\boldsymbol{\rho}_j) = \sum_{i=1}^{\ell} \beta_{j,i} \hat{B}_i, \quad (8e)$$

$$\mathbf{g}(\mathbf{u}_j) \leq 0, \quad (8f)$$

$$\mathbf{h}_\theta(\mathbf{x}_j, \mathbf{u}_j) \leq \boldsymbol{\sigma}_j, \quad \mathbf{h}_\theta^f(\mathbf{x}_{k+N}) \leq \boldsymbol{\sigma}_{k+N} \quad (8g)$$

$$\boldsymbol{\sigma}_{k, \dots, k+N} \geq 0 \quad (8h)$$

where T_θ and l_θ are the parameterized terminal and stage costs, respectively. The function \mathcal{G}_θ is labeled the cost modification term, which can be, i.e., in a form of gradient $\mathbf{f}^\top[\mathbf{x}_j, \mathbf{u}_j]^\top$. However, the polytopic LPV model is uncertain with wrong vertices that could cause some violations in its constraints and bring the MPC in an infeasible region. To avoid such an infeasibility, we introduce some slack variables $\boldsymbol{\sigma}$ penalized by enough large weights \mathbf{w}, \mathbf{w}_f . We also propose to have some adjustable parameters upon inequality constraints $\mathbf{h}_\theta, \mathbf{h}_\theta^f$. Then, we let RL to tune the MPC parameters (a.k.a policy parameters). Note that, the proposed parameterization of the MHE and MPC is a general form of a rich parameterization and one may take into account some of these parameters to be adjusted by RL.

Note that the parameterized terminal cost in the MPC scheme (8) is defined as a function of the terminal states and the convex combination vector as follows:

$$T_\theta = \mathbf{x}_{k+N}^\top \left(\sum_{i=1}^{\ell} \beta_{k+N,i} P_i \right) \mathbf{x}_{k+N} \quad (9)$$

where the matrices $P_i > 0, i \in \mathbb{Z}_{1, \dots, \ell}$ are adjusted by RL. However, the initial values for these matrices can be computed offline via a Linear Matrix Inequality (LMI) based on a parameter-dependent Lyapunov function and a parameter-dependent linear control law detailed in (Di Cairano and Danielson, 2021).

4. MHE/MPC-BASED POLICY GRADIENT RL

Let us define the closed-loop performance of a parameterized policy π_θ for a given stage cost $L(\mathbf{x}, \mathbf{a})$ as the following total expected cost:

$$J(\pi_\theta) = \mathbb{E}_{\rho_s} \left[\sum_{k=0}^{\infty} \gamma^k L(\mathbf{x}_k, \mathbf{a}_k) \mid \mathbf{a}_k = \pi_\theta(\mathbf{x}_k) \right] \quad (10)$$

where the expectation $\mathbb{E}_{\rho_s}[\cdot]$ is taken over the distribution of the Markov chain ρ_s in the closed-loop under policy π_θ . We then seek the optimal policy parameters as follows:

$$\boldsymbol{\theta}_* = \arg \min_{\boldsymbol{\theta}} J(\pi_\theta) \quad (11)$$

As a learning-based control approach in this paper, we propose to use a parameterized MPC scheme as a policy approximation in order to deliver π_θ . As an advantage to use a MPC-based policy approximation instead of DNNs, it can offer a learning-based controller so that all the state-input constraints are satisfied and the closed-loop performance is improved by adjusting the policy parameters.

For a given estimated state $\hat{\mathbf{x}}_k$ and parameter $\hat{\beta}_k$ delivered by the MHE scheme, the policy captured by a parameterized MPC scheme is

$$\pi_\theta(\hat{\mathbf{x}}_k, \hat{\beta}_k) = \mathbf{u}_0^*(\hat{\mathbf{x}}_k, \hat{\beta}_k, \boldsymbol{\theta}) \quad (12)$$

where \mathbf{u}_0^* is the first element of the control input sequence \mathbf{u}^* delivered by (8).

In this section, we propose a policy gradient RL framework based on MPC and MHE to deal with the polytopic LPV systems with wrong LTI models.

4.1 Compatible Deterministic Actor-Critic

In the context of DPG-based RL algorithms, the policy parameters $\boldsymbol{\theta}$ can be directly optimized by the gradient descent steps such that the best expected closed-loop cost (a.k.a policy performance index J) can be captured by applying the policy π_θ .

$$\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \alpha \nabla_{\boldsymbol{\theta}} J(\pi_\theta) \quad (13)$$

for some $\alpha > 0$ small enough as the step size. In the context of hybrid controller/observer scheme MHE-MPC, the input signal can be interpreted as a sequence of measurements $\mathbf{y} = [y_{0, \dots, k}^1, \dots, y_{0, \dots, k}^{n_y}]^\top$ at the physical time k , where n_y is the number of system states selected as measurable outputs. Then, the intermediate variables $\hat{\mathbf{x}}_k, \hat{\beta}_k$ are delivered by the MHE scheme based on the history of the measurements and fed to the MPC scheme to deliver the control policy. Assuming that this history constitutes a Markov state, one can define a policy performance index by the following expected value:

$$J(\pi_\theta) := \mathbb{E}_{\pi_\theta} \left[Q_{\pi_\theta}(\hat{\mathbf{x}}_k, \pi_\theta(\hat{\mathbf{x}}_k, \hat{\beta}_k)) \right] \quad (14)$$

where the expectation \mathbb{E}_{π_θ} is taken over the distribution of the Markov chain resulting from the real system in closed-loop with π_θ . The action-value function Q_{π_θ} is then defined as follows:

$$Q_{\pi_\theta}(\hat{\mathbf{x}}_k, \mathbf{a}_k) = L(\mathbf{y}_k, \mathbf{a}_k) + \gamma \mathbb{E}_{\pi_\theta} [V_{\pi_\theta}(\hat{\mathbf{x}}_{k+1}) \mid \hat{\mathbf{x}}_k, \mathbf{a}_k] \quad (15)$$

where, $L(\mathbf{y}_k, \mathbf{a}_k)$ reads as baseline cost (RL stage cost), which is a function of measurable states and actions at the current time

k . Based on the proposed DPG theorem by (Silver et al., 2014) and the fact that both the π_θ and Q_{π_θ} are functions of \hat{x}_k , the policy gradient equation is described as follows:

$$\nabla_\theta J(\pi_\theta) = \mathbb{E} [\nabla_\theta \pi_\theta(\hat{x}_k) \nabla_u Q_{\pi_\theta}(\hat{x}_k, \mathbf{a}_k) |_{\mathbf{a}_k = \pi_\theta}] \quad (16)$$

To represent the effect of the parameterized MHE upon the policy gradient, the sensitivity of the policy w.r.t θ can be updated such that the new policy gradient is described by the following expectation:

$$\nabla_\theta J(\pi_\theta) = \mathbb{E} [\Xi \nabla_u Q_{\pi_\theta}(\hat{x}_k, \mathbf{a}_k) |_{\mathbf{a}_k = \pi_\theta}] \quad (17)$$

where the Jacobian matrix Ξ is obtained by the following chain rule:

$$\Xi = (\nabla_\theta \pi_\theta + \nabla_\theta \hat{x}_k \nabla_{\hat{x}_k} \pi_\theta + \nabla_\theta \hat{\beta}_k \nabla_{\hat{\beta}_k} \pi_\theta) \quad (18)$$

In this paper, we adopt a *compatible deterministic actor-critic* algorithm (Silver et al., 2014) in which the action-value function $Q_{\pi_\theta}(\hat{x}_k, \mathbf{a}_k)$ can be replaced by a class of compatible function approximator $Q^w(\hat{x}_k, \mathbf{a}_k)$ such that the policy gradient is preserved. Therefore, the compatible function for a deterministic policy π_θ delivered by the parameterized MHE-MPC scheme can be expressed as follows:

$$Q^w = (\mathbf{a}_k - \pi_\theta)^\top \Xi^\top \mathbf{w} + V^\nu(\hat{x}_k) \quad (19)$$

The first term in the above compatible function as critic part is an estimation for the advantage function and the second term estimates a value function for the history of the measurements delivered as a summarized variable \hat{x}_k by the MHE scheme. Both functions can be computed by the linear function approximators as follows:

$$V^\nu(\hat{x}_k) = \Upsilon(\hat{x}_k)^\top \nu, \quad (20a)$$

$$A^w(\hat{x}_k, \mathbf{a}_k) = \Psi(\hat{x}_k, \mathbf{a}_k)^\top \mathbf{w} \quad (20b)$$

where $\Upsilon(\hat{x}_k)$ is the summarized measurement feature vector in order to constitute all monomials of the history of the measurements with degrees less than or equal to 2. The vector $\Psi(\hat{x}_k, \mathbf{a}_k) := \Xi(\mathbf{a}_k - \pi_\theta(\hat{x}_k, \hat{\beta}_k))$ includes the state-action features. Considering (19), the policy gradient (17) is then rewritten as follows:

$$\nabla_\theta J(\pi_\theta) = \mathbb{E} [\Xi \Xi^\top \mathbf{w}] \quad (21)$$

In this paper, the parameterized policy in the context of policy gradient RL is proposed to be captured by the MPC scheme (8). To evaluate the policy gradient (21), one need to calculate some sensitivities upon the MPC and MHE schemes in order to compute the Jacobian matrix Ξ . Hence, the Jacobian matrices $\nabla_\theta \pi_\theta$, $\nabla_{\hat{x}_k} \pi_\theta$ and $\nabla_{\hat{\beta}_k} \pi_\theta$ can be computed by the sensitivity analysis for the parameterized MPC scheme while the gradients $\nabla_\theta \hat{x}_k$ and $\nabla_\theta \hat{\beta}_k$ are obtained as a sensitivity term for the parameterized MHE scheme.

4.2 Sensitivity Analysis and LSTD-based DPG

Sensitivity Computation We describe next how to compute the sensitivities (gradients) needed in the proposed policy gradient RL framework based on MHE-MPC. To that end, let us define the Lagrange functions $\hat{\mathcal{L}}_\theta, \mathcal{L}_\theta$ associated to the MHE and MPC schemes (7), (8) as follows:

$$\hat{\mathcal{L}}_\theta(\hat{z}) = \hat{\Lambda}_\theta + \hat{\lambda}^\top \hat{G}_\theta \quad (22)$$

$$\mathcal{L}_\theta(z) = \Lambda_\theta + \lambda^\top G_\theta + \mu^\top H_\theta \quad (23)$$

where Λ_θ and $\hat{\Lambda}_\theta$ are the total parameterized costs of the MPC and MHE schemes, respectively. The inequality constraints of

(8) are collected by H_θ while G_θ and \hat{G}_θ gather, respectively, the equality constraints in the MPC and MHE schemes. We then label $\lambda, \hat{\lambda}$ the Lagrange multipliers associated to the equality constraints G_θ, \hat{G}_θ of the MPC and MHE, respectively. Variables μ are the Lagrange multipliers associated to the inequality constraints of the MPC scheme. Let us label $\Gamma = \{x, u, \sigma\}$ and $\hat{\Gamma} = \{\hat{x}, \hat{\beta}\}$ the primal variables for the MPC and MHE, respectively. The associated primal-dual variables then read as $z = \{\Gamma, \lambda, \mu\}$ and $\hat{z} = \{\hat{\Gamma}, \hat{\lambda}\}$.

The sensitivity of the policy delivered by the MPC scheme (8) w.r.t policy parameters and the sensitivity of the estimated state associated to the MHE scheme (7) can be obtained via using the Implicit Function Theorem (IFT) on the Karush Kuhn Tucker (KKT) conditions underlying the parametric NLP. Assuming that Linear Independence Constraint Qualification (LICQ) and Second Order Sufficient Condition (SOSC) hold (Nocedal and Wright, 2006) at z^* and \hat{z}^* , then, the following holds:

$$\frac{\partial z^*}{\partial \theta} = -\frac{\partial \kappa_\theta}{\partial z}^{-1} \frac{\partial \kappa_\theta}{\partial \theta}, \quad (24a)$$

$$\frac{\partial \hat{z}^*}{\partial \theta} = -\frac{\partial \hat{\kappa}_\theta}{\partial \hat{z}}^{-1} \frac{\partial \hat{\kappa}_\theta}{\partial \theta} \quad (24b)$$

where

$$\kappa_\theta = \begin{bmatrix} \nabla_\Gamma \mathcal{L}_\theta \\ G_\theta \\ \text{diag}(\mu) H_\theta \end{bmatrix}, \quad \hat{\kappa}_\theta = \begin{bmatrix} \nabla_{\hat{\Gamma}} \hat{\mathcal{L}}_\theta \\ \hat{G}_\theta \end{bmatrix} \quad (25)$$

are the KKT conditions associated to the MPC and MHE schemes, respectively. As π_θ and $(\hat{x}, \hat{\beta})$ are, respectively, part of z^* and \hat{z}^* . Then, the sensitivity of the MPC policy $\nabla_\theta \pi_\theta$ and the sensitivity of the MHE solution $(\nabla_\theta \hat{x}, \nabla_\theta \hat{\beta})$ required in (21) can be extracted from gradients $\frac{\partial z^*}{\partial \theta}$ and $\frac{\partial \hat{z}^*}{\partial \theta}$, respectively.

LSTD-based Policy Gradient In the context of compatible DPG, one can evaluate the optimal parameters \mathbf{w} and ν of the action-value function approximation (19) as solutions of the following Least Squares (LS) problem:

$$\min_{\mathbf{w}, \nu} \mathbb{E} [(Q_{\pi_\theta}(\hat{x}_k, \mathbf{a}_k) - Q^w(\hat{x}_k, \mathbf{a}_k))^2], \quad (26)$$

In the context of RL, the LSTD-based algorithms offer an efficient use of data and tend to converge faster than other methods, e.g., see (Esfahani et al., 2021a). The LSTD update rules for a policy gradient RL are then obtained as follows:

$$\nu \leftarrow A_\nu^{-1} b_\nu, \quad (27a)$$

$$\mathbf{w} \leftarrow A_w^{-1} b_w, \quad (27b)$$

$$\theta \leftarrow \theta - \alpha b_\theta \quad (27c)$$

where the matrices $\Omega_{(\cdot)}$ and the vectors $b_{(\cdot)}$ are calculated by taking expectation (\mathbb{E}_m) over m episodes as follows:

$$A_{\nu} = \mathbb{E}_m \left[\sum_{k=1}^{T_f} \left[\Upsilon(\hat{\mathbf{x}}_k) (\Upsilon(\hat{\mathbf{x}}_k) - \gamma \Upsilon(\hat{\mathbf{x}}_{k+1}))^\top \right] \right] \quad (28a)$$

$$A_w = \mathbb{E}_m \left[\sum_{k=1}^{T_f} \left[\Psi(\hat{\mathbf{x}}_k, \mathbf{a}_k) \Psi(\hat{\mathbf{x}}_k, \mathbf{a}_k)^\top \right] \right], \quad (28b)$$

$$b_{\nu} = \mathbb{E}_m \left[\sum_{k=1}^{T_f} \Upsilon(\hat{\mathbf{x}}_k) L(\mathbf{y}_k, \mathbf{a}_k) \right], \quad (28c)$$

$$b_w = \mathbb{E}_m \left[\sum_{k=1}^{T_f} \left[(L(\mathbf{y}_k, \mathbf{a}_k) + \gamma V^{\nu}(\hat{\mathbf{x}}_{k+1}) - V^{\nu}(\hat{\mathbf{x}}_k)) \Psi(\hat{\mathbf{x}}_k, \mathbf{a}_k) \right] \right], \quad (28d)$$

$$b_{\theta} = E_m \left[\sum_{k=1}^{T_f} \Xi \Xi^\top w \right] \quad (28e)$$

where T_f is the final time instant at the end of each episode.

5. ILLUSTRATIVE EXAMPLE

To demonstrate the effectiveness of the proposed learning-based control approach, we choose a stabilization problem for a mass-spring-damper system where the spring and damper coefficients are considered as the exogenous scheduling parameters.

$$\frac{d}{dt} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -k/m & -d/m \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} 0 \\ 1/m \end{bmatrix} u \quad (29)$$

Then, the LPV model can be represented as a linear system with polytopic uncertainty including four vertices (LTI models), which are constructed based on the scheduling bounds $1 \leq k \leq 2, 0 \leq b \leq 0.5$ shown in Figure 1.

$$\begin{aligned} A_1 &= \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}, & A_2 &= \begin{bmatrix} 0 & 1 \\ -1 & -0.5 \end{bmatrix} \\ A_3 &= \begin{bmatrix} 0 & 1 \\ -2 & 0 \end{bmatrix}, & A_4 &= \begin{bmatrix} 0 & 1 \\ -2 & -0.5 \end{bmatrix} \end{aligned} \quad (30)$$

and $B_{1,\dots,4} = [0, 1]^\top$. On the other hand, we select the scheduling bounds used in the real system different from those adopted in four LTI models $0.5 \leq k \leq 2, 0 \leq b \leq 0.2$. The position of the mass (x_1) is selected as measurement while the velocity ($x_2 = \dot{x}_1$) is estimated. We consider a constraint on the control input $-1 \leq u \leq 1$. As it can be observed in Figure




Fig. 1. Mass-Spring-Damper with variable spring constant and damping factor as scheduling parameters

2, the stabilization scenario is affected since the scheduling

parameters violate their correct values used in four LTI models. We then adopt a policy gradient RL to adjust the MHE-MPC parameters in order to cope with this problem and improve the closed-loop performance $J(\pi_\theta)$. The evolution of some RL parameters and the closed-loop performance are shown in Figure 3. The evolution during learning shows that the optimal policy (Figure 2: π_θ in black color) can be delivered after 35 RL step and the stabilization goal is perfectly achieved. Note that the system's behavior in terms of β -parameters and control policy without learning is highlighted in red color shown in Figure 2. It is obvious that the estimated combination vector β is improved by adjusting the MHE scheme under wrong LTI models. As a comparative study shown in Figure 4, the performance of the stabilization has been dramatically improved after learning the MHE-MPC scheme where the LTI models cannot capture the correct polytope based on the scheduling bounds used in the real system.

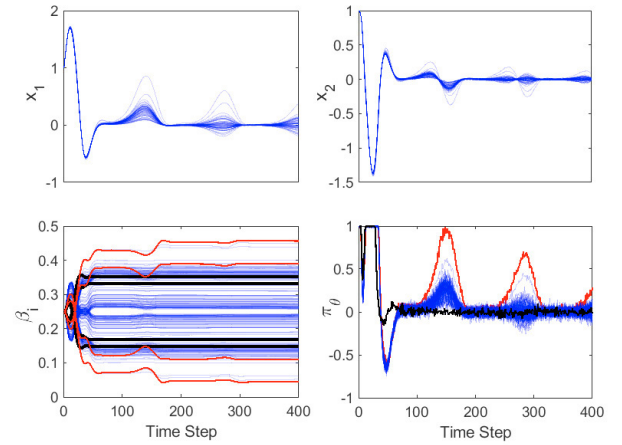


Fig. 2. Evolution during reinforcement learning: states, policy and β -parameters




Fig. 3. Closed-loop performance and parameters of the policy gradient reinforcement learning




Fig. 4. Comparative analysis

6. CONCLUSION

In this paper, an MHE-MPC scheme was proposed to deal with the linear systems with polytopic uncertainty in the context of the LPV systems with exogenous scheduling parameters. However, the performance of the adopted multi-model approach can be affected where the bounds of the scheduling parameters are different from their real values applied to the real system. To address a data-driven MHE-MPC approach for the polytopic LPV systems with inexact LTI models, we proposed to adopt a policy gradient reinforcement learning to capture the optimal policy and achieve the best closed-loop performance by adjusting the MHE and MPC schemes.

ACKNOWLEDGEMENTS

This work was supported by the Research Council of Norway (RCN), project “Safe Reinforcement Learning using MPC (SARLEM)”.

REFERENCES

- Cai, W., Kordabad, A.B., Esfahani, H.N., Lekkas, A.M., and Gros, S. (2021). Mpc-based reinforcement learning for a simplified freight mission of autonomous surface vehicles. In *2021 60th IEEE Conference on Decision and Control (CDC)*, 2990–2995.
- Cairano, S.D. (2012). An industry perspective on mpc in large volumes applications: Potential benefits and open challenges. *IFAC Proceedings Volumes*, 45(17), 52–59. 4th IFAC Conference on Nonlinear Model Predictive Control.
- Di Cairano, S. (2016). Indirect adaptive model predictive control for linear systems with polytopic uncertainty. In *2016 American Control Conference (ACC)*, 3570–3575.
- Di Cairano, S. and Danielson, C. (2021). Indirect adaptive model predictive control and its application to uncertain linear systems. *International Journal of Robust and Nonlinear Control*, 31(18), 8678–8702.
- Ding, B. and Pan, H. (2016). Output feedback robust mpc for lpv system with polytopic model parametric uncertainty and bounded disturbance. *International Journal of Control*, 89(8), 1554–1571.
- Esfahani, H.N., Kordabad, A.B., and Gros, S. (2021a). Approximate robust nmpc using reinforcement learning. In *2021 European Control Conference (ECC)*, 132–137.
- Esfahani, H.N., Kordabad, A.B., and Gros, S. (2021b). Reinforcement learning based on mpc/mhe for unmodeled and partially observable dynamics. In *2021 American Control Conference (ACC)*, 2121–2126.
- Kordabad, A.B., Esfahani, H.N., Lekkas, A.M., and Gros, S. (2021). Reinforcement learning based on scenario-tree mpc for asvs. In *2021 American Control Conference (ACC)*, 1985–1990.
- Langson, W., Chrysochoos, I., Raković, S., and Mayne, D. (2004). Robust model predictive control using tubes. *Automatica*, 40(1), 125–133.
- Leith, D.J. and Leithead, W.E. (2000). Survey of gain-scheduling analysis and design. *International Journal of Control*, 73(11), 1001–1025.
- Nocedal, J. and Wright, S. (2006). *Numerical Optimization*. Springer, 2 edition.
- Pipino, H.A. and Adam, E.J. (2019). Mpc for linear systems with parametric uncertainty. In *2019 XVIII Workshop on Information Processing and Control (RPIC)*, 42–47. doi: 10.1109/RPIC.2019.8882151.
- Pipino, H.A., Morato, M.M., Bernardi, E., Adam, E.J., and Normey-Rico, J.E. (2020). Nonlinear temperature regulation of solar collectors with a fast adaptive polytopic lpv mpc formulation. *Solar Energy*, 209, 214–225.
- Pour, F.K., Puig, V., and Ocampo-Martinez, C. (2017). Comparative assessment of lpv-based predictive control strategies for a pasteurization plant. In *2017 4th International Conference on Control, Decision and Information Technologies (CoDIT)*, 0821–0826.
- Rawlings, J.B., Mayne, D.Q., and Diehl, M. (2017). *Model predictive control: theory, computation, and design*, volume 2. Nob Hill Publishing Madison, WI.
- Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., and Riedmiller, M. (2014). Deterministic policy gradient algorithms. In *Proceedings of the 31st International Conference on International Conference on Machine Learning*, I–387–I–395. JMLR.org.
- Sutton, R.S. and Barto, A.G. (2018). *Reinforcement learning: An introduction*. MIT press.
- Zhou, J., Di Cairano, S., and Danielson, C. (2017). Indirect adaptive mpc for output tracking of uncertain linear polytopic systems. In *2017 American Control Conference (ACC)*, 3054–3059.