Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnLine 54-20 (2021) 46–52

Learning-based Robust Model Predictive Control for Sector-bounded Lur'e Systems

Katrine Seel *,** **Mark Haring** ** **Esten I. Grøtli** **
Kristin Y. Pettersen * **Jan T. Gravdahl** *

* *Department of Engineering Cybernetics, Norwegian University of Science and Technology, Trondheim, Norway (e-mails: {katrine.seel, kristin.y.pettersen, jan.tommy.gravdahl}@ntnu.no)*
** *Department of Mathmatics and Cybernetics, SINTEF Digital, Trondheim, Norway (e-mails: {mark.haring, esteningar.grotli}@sintef.no)*

Abstract: For dynamical systems with uncertainty, robust controllers can be designed by assuming that the uncertainty is bounded. The less we know about the uncertainty in the system, the more conservative the bound must be, which in turn may lead to reduced control performance. If measurements of the uncertain term are available, this data may be used to reduce the uncertainty in order to make bounds as tight as possible. In this paper, we consider a linear system with a sector-bounded uncertainty. We develop a model predictive control algorithm to control the system, and use a weighted Bayesian linear regression model to learn the least conservative sector condition using measurements collected in closed-loop. The resulting robust model predictive control algorithm therefore reduces the conservativeness of the controller, and provides probabilistic guarantees of asymptotic stability and constraint satisfaction. The efficacy of the proposed method is shown in simulation.

Keywords: MPC, Bayesian linear regression, robust control

1. INTRODUCTION

Model predictive control (MPC) is a popular method for controlling systems with complex dynamics, subject to state and input constraints. The control method solves an optimization problem at every sampling instant, over some finite prediction horizon, in order to find a sequence of control inputs that optimizes the open-loop behaviour. For complex systems, MPC will often require solving nonlinear, possibly non-convex, optimization problems, resulting in a nonlinear MPC (NMPC) scheme.

Even though modern solvers are able to handle NMPC algorithms, solving the resulting optimization problem is in general challenging. The main reason is that for non-convex optimization problems, the solvers are often not guaranteed to find global minima, and may instead get stuck in a local minimum. A way to get around this, is to treat the nonlinearities present in the dynamics as uncertainties in a linear system. The linear system can then be controlled using algorithms that are robust with respect to some bounded uncertainty in the system, by solving convex optimization problems. In Kothare et al. (1996), robust MPC schemes are formulated for polytopic systems and for linear systems with structured uncertainty. A sector bound is particularly suited for modelling state-dependent uncertainty and is used with a similar MPC algorithm in Böhm et al. (2010), Böhm et al. (2009), Nguyen et al. (2018).

For uncertain systems, the smallest possible bound on the uncertainty may not be known apriori. Using more conservative bounds will in turn lead to more conservative controllers and correspondingly reduced control performance. If the uncertainty in the system can be measured or estimated, learning-based methods may be used to provide robust controllers with improved performance. In Ostafew et al. (2014), a Gaussian Process (GP) is used to model the disturbances in a vehicle model. The learned model is used to enhance a nominal prediction model in an MPC scheme, resulting in improved path-tracking performance.

A similar approach is taken in McKinnon and Schoellig (2019), using weighted Bayesian linear regression (wBLR) to learn unknown dynamics in the prediction model. In Hewing et al. (2019), GP regression is used to learn unmodelled dynamics to be used in stochastic MPC. The learned model complements the prediction model, and the model uncertainty is used to update the chance constraints. Compared to using only a nominal model, the addition of the learned model results in cautious control with improved performance. A GP model is also used in Soloperto et al. (2018) to model the uncertainty in a linear system, ensuring robust constraint satisfaction and resulting in less conservative control.

In this paper, we consider linear systems with sector-bounded uncertainties, as in Böhm et al. (2010), Böhm et al. (2009) and Nguyen et al. (2018). We develop a robust MPC algorithm similar to the one formulated in Böhm et al. (2009). However, instead of assuming that the smallest possible sector is known apriori, we use measurements of the uncertainty to learn the sector bound. This is done using a Bayesian linear regression (BLR)

10.1016/j.ifacol.2021.11.151

model (Murphy, 2012), with weighted data points as in McKinnon and Schoellig (2019), that is particularly suited for finding local approximations of nonlinear functions. To the best of the authors' knowledge, this has not been done before. The first contribution of this paper is therefore to use closed-loop measurements to tighten a sector bound of an uncertain nonlinear term. Because the wBLR model estimates distributions rather than deterministic model parameters, it provides a measure of model uncertainty, which is used to formulate a stochastic sector bound, that in turn provides probabilistic stability and constraint satisfaction guarantees for the closed-loop system. The third contribution of this paper, is a reformulation of the optimal infinite horizon problem, formulated in Böhm et al. (2009), for discrete-time systems.

The paper is structured as follows: Section 2 provides a description of the problem statement. In Section 3, we give a brief overview of how wBLR can be used to learn the sector bound, followed by a description of the resulting learning-based robust MPC algorithm in Section 4. Simulation results are provided in Section 5, and conclusions are given in Section 6.

2. PROBLEM STATEMENT

We consider a subclass of discrete-time nonlinear systems, namely sector-bounded Lur'e systems, which can be written in the form

$$\begin{aligned} x_{k+1} &= Ax_k + G\gamma(z_k) + Bu_k, \\ z_k &= Hx_k, \end{aligned} \tag{1}$$

where x_k is the state vector, u_k is the control input and z_k is the input of the nonlinearity $\gamma(z) : \mathbb{R} \to \mathbb{R}$ for $z \in \mathbb{R}$. The system matrices have dimensions $A \in \mathbb{R}^{n\times n}$, $B \in \mathbb{R}^{n\times m}$, $G \in \mathbb{R}^{n\times 1}$ and $H \in \mathbb{R}^{1\times n}$. We assume that the nonlinearity satisfies

$$(uz - \gamma(z))(\gamma(z) - lz) \geq 0 \quad \forall z, \tag{2}$$

where $u, l \in \mathbb{R}^+$, i.e. is bounded by the sector condition as visualized in Figure 1.

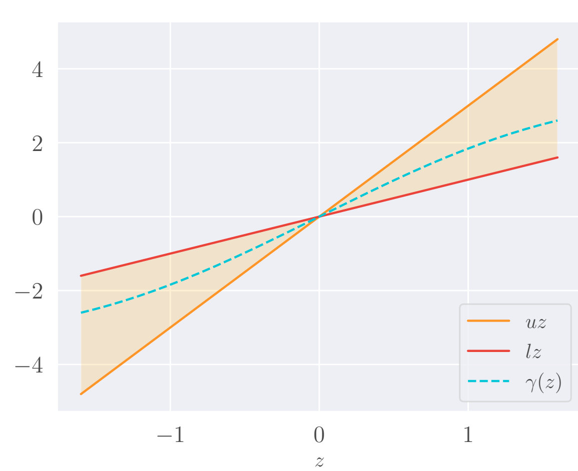

Fig. 1. Sector-bounded nonlinear function, $\gamma(z)$.

The system is subject to r polytopic state and input constraints, of the form

$$\mathscr{C}_k = \left\{ \begin{bmatrix} x_k \\ u_k \end{bmatrix} \in \mathbb{R}^{n+m} : c_j x_k + d_j u_k \leq 1, j = 1, ..., r \right\}, \tag{3}$$

that must be satisfied at every time instant k. The control objective is to steer the system (1) to the origin, for all nonlinearities that satisfy the sector condition (2), and for the input- and state-constraints in (3).

To this end, a sector condition that is as small as possible can improve the control performance by making it less conservative. For this purpose, both open-loop and closed-loop measurements can be used to reduce the bounds on the uncertainty. In this paper, we focus on the latter and propose using closed-loop measurements to tighten an initially conservative sector bound. For systems of the form (1), where the full state is sampled for $k \geq 0$, we can estimate $\gamma(z_k)$ using the available closed-loop measurements. For each time step, we use x_{k+1} and x_k, in combination with (1) and the known system matrices A, B, G and H, to obtain an estimate of $\gamma(z_k)$.

3. LEARNING THE SECTOR BOUND

The goal of this section is to describe how closed-loop measurements can be used to learn a sector bound, in order to make it as tight as possible. For this purpose we use wBLR, which is an extension of BLR as presented in Murphy (2012), and a modification of McKinnon and Schoellig (2019). We assume that we have a dataset of n training samples, $\mathcal{D} = \{z_i, y_i\}_{i=1}^n$, with $y_i = \gamma(z_i)$. Consider a local model

$$y_i = wz_i + \epsilon, \tag{4}$$

where z_i and y_i is the scalar input and output, respectively, and with zero mean Gaussian noise, $\epsilon \sim \mathcal{N}(0, \sigma^2)$, where σ^2 is the variance. For the sampled region of input space, we want to approximate a locally linear model of the form

$$\hat{y} = wz, \tag{5}$$

where w is a stochastic variable. Because the available data points are sampled from a closed-loop system with a fixed sampling frequency, we use scalar weights $l_i \in [0, 1]$ to determine the importance of each data point as done in McKinnon and Schoellig (2019). As in McKinnon and Schoellig (2019), we assume that the data points are weighted ahead of learning. However, we weigh the datapoints differently, namely by considering the density of data points in input space. To this end, the closed-loop measurements are sorted, and then weighted according to the input points' similarity with both the previous and the next data point. The weight is then scaled using the largest difference between subsequent data points, according to

$$l_i = \frac{0.5\|z_i - z_{i-1}\| + 0.5\|z_i - z_{i+1}\|}{\max(\{\|z_2 - z_1\|, \ldots, \|z_n - z_{n-1}\|\})}, \tag{6}$$

for $2 \leq i \leq n-1$. For the first and last data point we let $l_i = 0$. For a sample z_i that is very similar to the previous sample z_{i-1} and the next sample z_{i+1}, the weight is small, $l_i \approx 0$, and the data point will have little influence on the regression. For the opposite case, the sample will be weighted with $l_i \approx 1$, and the data point is fully included in the regression. If all weights are 1, we obtain standard BLR.

For the weighted data set $\mathcal{D}^l = \{z_i, y_i, l_i\}_{i=1}^n$, we assume that each data point is independent, and distributed according to

$$p(y|Z, w, \sigma^2) = \prod_{i=1}^{n} \mathcal{N}(y_i|wz_i, \sigma^2)^{l_i}, \quad (7)$$

known as the likelihood function, where Z is a matrix with rows z_i and y is a vector with elements y_i. For this likelihood function, the conjugate prior is a Normal Inverse Gamma (NIG) distribution, resulting in the following priors for w and σ^2 (Murphy, 2012)

$$p(w|\sigma^2) \sim \mathcal{N}(w|w_0, \sigma^2 V_0), \quad (8a)$$

$$p(\sigma^2) \sim \mathrm{IG}(\sigma^2|a_0, b_0), \quad (8b)$$

where w_0 is the prior mean and $\sigma^2 V_0$ is the prior variance of the regression weight, and a_0 and b_0 are the initial parameters determining the Inverse Gamma (IG) distribution. For the specified likelihood (7) and prior (8), one can show that the posterior distribution over w and σ^2 has the form (Murphy, 2012)

$$p(w, \sigma^2|\mathcal{D}^l) = \mathrm{NIG}(w, \sigma^2|w_N, V_N, a_N, b_N), \quad (9)$$

with

$$w_N = V_N(V_0^{-1}w_0 + Z^T L y), \quad (10a)$$

$$V_N = (V_0^{-1} + Z^T L Z)^{-1}, \quad (10b)$$

$$a_N = a_0 + \frac{\mathrm{tr}(L)}{2}, \quad (10c)$$

$$b_N = b_0 + \frac{1}{2}(w_0 V_0^{-1} w_0 + y^T L y - w_N V_N^{-1} w_N), \quad (10d)$$

where $\mathrm{tr}(\cdot)$ is the trace operator, and L is a diagonal matrix with the data weights l_i. From the joint posterior distribution (9), we obtain the marginal distributions for w and σ^2 (Murphy, 2012)

$$p(w|\mathcal{D}^l) = \mathcal{T}(w|w_N, \frac{b_N}{a_N}V_N, 2a_N), \quad (11a)$$

$$p(\sigma^2|\mathcal{D}^l) = \mathrm{IG}(\sigma^2|a_N, b_N). \quad (11b)$$

In order to define a stochastic model of the sector bound, we use the uncertainty of w, given by the Student t-distribution as specified in (11a). To make use of the stochastic sector bound in the control design, we use the confidence interval for w to define the upper $\hat{l}$ and lower $\hat{u}$ bound on the local linear approximation of the nonlinear function $\gamma(z)$.

Assumption 1. *For the local linear approximation $\hat{y}$*

$$Pr\big((\hat{u}z - \hat{y})(\hat{y} - \hat{l}z) \geq 0\big) \geq p_s \quad \forall \hat{y}, \quad (12)$$

holds with probability, $Pr(\cdot)$, at least p_s.

4. LEARNING-BASED ROBUST MPC

For the system (1) and the control objective as described in Section 2, we propose to use NMPC. The following Section is dedicated to describing the control design, as well as analyzing the stability and recursive feasibility of the closed-loop system.

The basic idea of MPC is to solve an optimal control problem online, at each time instant k, based on the measured system state $x_{k|k}$. The goal of the optimal control problem, is to find a control sequence that minimizes the infinite horizon cost function

$$J_{\infty|k} = \sum_{i=0}^{\infty} F(x_{k+i|k}, u_{k+i|k}), \quad (13)$$

where $i \geq 0$ is the prediction time variable and

$$F(x_{k+i|k}, u_{k+i|k}) = x_{k+i|k}^T Q x_{k+i|k} + u_{k+i|k}^T R u_{k+i|k}, \quad (14)$$

with $Q > 0 \in \mathbb{R}^{n\times n}$ and $R > 0 \in \mathbb{R}^{m\times m}$. Because of the infinite horizon cost function, the resulting constrained optimization problem is in general not solvable. This is overcome by assuming a controller of the form

$$u_{k+i|k} = K_k x_{k+i|k}, \quad (15)$$

and at each time instant k calculating a constant feedback matrix $K_k \in \mathbb{R}^{m\times n}$, instead of the complete input trajectory. The resulting optimal control problem is a min-max problem, that determines the feedback matrix K_k, that minimizes the upper bound on the infinite horizon cost. The uncertain nonlinear term is treated as a sector condition, so that the resulting optimization problem is convex, which we can solve efficiently and for which we can find a global minimum. By recalculating K_k, the controller can be more aggressive as the state evolves closer to the origin, compared to the corresponding static feedback law (Kothare et al., 1996).

With the purpose of incorporating the learned sector from Section 3 in the control design, we define the following parameters

$$\delta = \frac{\hat{u} + \hat{l}}{2}, \quad (16)$$

$$\nu = \frac{\hat{u} - \hat{l}}{2}, \quad (17)$$

where $\hat{l}$ and $\hat{u}$ are the lower and upper bound of the learned sector, as given by the confidence interval for (11a). The parameters are used to shift the nonlinearity according to

$$\varphi(z) = \gamma(z) - \delta z. \quad (18)$$

This is done in order to work with a more convenient expression for the sector bound in the control design. To this end, we use (18) to express $\gamma(z)$ in the original sector condition (2), and obtain the following inequality

$$(\nu z - \varphi(z))(\varphi(z) + \nu z) \geq 0, \quad (19)$$

which implies that

$$|\varphi(z)| \leq |\nu H x|. \quad (20)$$

Let $E := \nu H \in \mathbb{R}^{1\times n}$ and rewrite so that

$$\varphi^2(z) \leq x^T E^T E x, \quad (21)$$

which in matrix form becomes

$$\begin{bmatrix} x \\ \varphi \end{bmatrix}^T \begin{bmatrix} E^T E & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} x \\ \varphi \end{bmatrix} \geq 0. \quad (22)$$

We now define

$$\bar{A} = A + \delta G H, \quad (23)$$

so that the system (1) can be written as

$$\begin{aligned} x_{k+1} &= \bar{A}x_k + G\varphi(z_k) + Bu_k, \\ z_k &= Hx_k. \end{aligned} \quad (24)$$

The following Lemma provides conditions for obtaining a stabilizing feedback law (15) for a system (1) with sector condition (2) and an upper bound on the infinite horizon cost (13). This is similar to Lemma 2 in Böhm et al. (2009),but is modified to apply for discrete-time systems and formulates a convex optimization problem without having to fix any of the optimization variables. The Lemma is also similar to Theorem 1 in Kothare et al. (1996), but is derived using the sector condition as formulated in (2).

Lemma 1. *Let Assumption 1 hold, so that for system (1), with a sector bound as given by (2), and for the matrices* $\Xi = \Xi^T > 0 \in \mathbb{R}^{n\times n}$, $Y \in \mathbb{R}^{m\times n}$, *and scalars* $\lambda > 0$, $\alpha > 0$, *the following inequality*

$$\begin{bmatrix} \Xi & Y^T R^{\frac{1}{2}} & \Xi Q^{\frac{1}{2}} & \Xi E^T & \Xi \bar{A}^T + Y^T B^T \\ R^{\frac{1}{2}} Y & \alpha I & 0 & 0 & 0 \\ Q^{\frac{1}{2}} \Xi & 0 & \alpha I & 0 & 0 \\ E\Xi & 0 & 0 & \lambda & 0 \\ \bar{A}\Xi + BY & 0 & 0 & 0 & \Xi - G\lambda G^T \end{bmatrix} \geq 0. \quad (25)$$

is satisfied. For $K = Y\Xi^{-1}$ *and* $P = \alpha \Xi^{-1}$, *it holds that*

- a. *The feedback law* $u_{k+i|k} = Kx_{k+i|k}$ *asymptotically stabilizes the system (1), with the sector-bounded nonlinearity described by (2).*
- b. $V(x_{k|k}) = x_{k|k}^T P x_{k|k}$ *is an upper bound on the infinite horizon cost (13).*

Proof. We define $\lambda = \frac{\alpha}{\tau}$ and use that $P = \alpha\Xi^{-1}$, $K = Y\Xi^{-1}$. After some matrix manipulations we take the Schur complement to the matrix (25), and obtain that

$$\begin{bmatrix} \begin{bmatrix} (\bar{A}+BK)^T P (\bar{A}+BK) \\ -P + Q + K^T R K + \tau E^T E \end{bmatrix} & (\bar{A}+BK)^T P G \\ G^T P (\bar{A}+BK) & G^T P G - \tau \end{bmatrix} \leq 0. \quad (26)$$

Applying the lossless S-procedure, see e.g. Boyd et al. (1994), to (26), it follows that $v^T Q v \leq 0$, with $v^T = \begin{bmatrix} x^T & \varphi \end{bmatrix}$ and

$$Q = \begin{bmatrix} \begin{bmatrix} (\bar{A}+BK)^T P (\bar{A}+BK) \\ -P + Q + K^T R K \end{bmatrix} & (\bar{A}+BK)^T P G \\ G^T P (\bar{A}+BK) & G^T P G \end{bmatrix} \leq 0, \quad (27)$$

holds for all $x = x_{k+i|k}$ and $\varphi = \varphi(z_{k+i|k})$ that satisfies (22). This holds probabilistically under Assumption 1. The inequality (27), is equivalent to

$$\begin{aligned} x^T((\bar{A}+BK)^T P(\bar{A}+BK) - P + Q + K^T RK)x \\ +\varphi G^T P(\bar{A}+BK)x_k + x^T(\bar{A}+BK)^T PG\varphi \\ +\varphi G^T PG\varphi \leq 0. \end{aligned} \quad (28)$$

For a feedback law $u_{k+i|k} = Kx_{k+i|k}$ and a quadratic function of the form $V(x_{k|k}) = x_{k|k}^T P x_{k|k}$ where $V(0) = 0$ and $P > 0$, we then know that at sampling time k and $i \geq 0$

$$\begin{aligned} V(x_{k+i+1|k}) - V(x_{k+i|k})) \leq \\ -(x_{k+i|k}^T Q x_{k+i|k} + u_{k+i|k}^T R u_{k+i|k}). \end{aligned} \quad (29)$$

Thus, for $Q > 0$ and $R > 0$, $V(x_{k|k}) = x_{k|k}^T P x_{k|k}$ is a Lyapunov function and the control law asymptotically stabilizes the system (1), proving part (a) of Lemma 1.

Using that $x_{\infty|k} = 0$, so that $V(x_{\infty|k}) = 0$, and summing (29) from $i = 0$ to $i = \infty$, we get

$$J_{\infty|k} \leq V(x_{k|k}), \quad (30)$$

i.e. $V(x_{k|k})$ is an upper bound on the infinite horizon cost $J_{\infty|k}$, proving part (b) of Lemma 1. □

The following Lemma is used for showing constraint satisfaction:

Lemma 2. *The ellipsoid* $\mathscr{E}_k = \{x_k \in \mathbb{R}^n : x_k^T P_k x_k \leq \alpha_k\}$ *is contained in the constraint set* $\mathscr{C}_k$, *as described in (3), at time instant k if and only if*

$$(c_j + d_j K_k)(\alpha_k P_k^{-1})(c_j + d_j K_k)^T \leq 1, j = 1, \ldots, r. \quad (31)$$

Proof. See e.g. Boyd et al. (1994). □

Using the results from Lemma 1, we state the following theorem for an NMPC controller with probabilistic stability and constraint satisfation guarantees, that minimizes the upper bound on the infinite horizon cost function (13). This is a discrete-time version of Theorem 1 in Böhm et al. (2009). Contrary to this theorem, all matrix inequalities are linear, making the resulting optimization problem easier to solve.

Theorem 1. *Let Assumption 1 hold, so that for the system (1), with a sector-bounded nonlinearity (2), an NMPC scheme is given by solving the following optimization problem:*

$$\min_{\alpha_k, \Xi_k, Y_k, \lambda_k} \alpha_k \quad (32)$$

subject to

$$\begin{bmatrix} 1 & x_{k|k}^T \\ x_{k|k} & \Xi_k \end{bmatrix} \geq 0 \quad (33a)$$

$$\begin{bmatrix} \Xi & Y_k^T R^{\frac{1}{2}} & \Xi_k Q^{\frac{1}{2}} & \Xi_k E^T & \Xi_k \bar{A}^T + Y_k^T B^T \\ R^{\frac{1}{2}} Y_k & \alpha_k I & 0 & 0 & 0 \\ Q^{\frac{1}{2}} \Xi_k & 0 & \alpha_k I & 0 & 0 \\ E\Xi_k & 0 & 0 & \lambda_k & 0 \\ \bar{A}\Xi_k + BY_k & 0 & 0 & 0 & \Xi_k - G\lambda_k G^T \end{bmatrix} \geq 0 \quad (33b)$$

$$\begin{bmatrix} 1 & c_j \Xi_k + d_j Y_k \\ (c_j \Xi_k + d_j Y_k)^T & \Xi_k \end{bmatrix} \geq 0 \quad (33c)$$

$$j = 1, ..., r$$

with $P_k = \alpha_k \Xi_k^{-1}$ *and* $K_k = Y_k \Xi_k^{-1}$. *The NMPC scheme has the following properties:*

- a. *The optimization problem is feasible for all future time instants k if it is feasible at* $k = 0$.
- b. *The solution to the optimization problem (32)-(33) minimizes the upper bound* $V(x_{k|k}) = x_{k|k}^T P_k x_{k|k}$ *on the infinite-horizon cost (13) at each time instant k.*
- c. *If the optimization problem (32)-(33) is feasible at* $k = 0$, *then the control law*

$$u_{k+i|k} = K_k x_{k+i|k}, \quad i \geq 0, \quad (34)$$

asymptotically stabilizes the origin of the system (1), with sector condition (2) and state and input constraints (3) for all times $k \geq 0$.

Proof. The proof is divided into three parts, to show that the properties (a)-(c) hold.

Part (a): As only (33a) depends on $x_{k|k}$, we know that the solution to the optimization problem (32)-(33) satisfies constraints (33b) and (33c). As inequality (33b) is identical to (25) from Lemma 1, (29) in combination with (33a) means that $x_{k+1|k}^T P_k x_{k+1|k} \leq x_{k|k}^T P_k x_{k|k} \leq \alpha_k$. Hence, the solution to the optimization problem at time k is also a solution at time $k+1$. By induction, feasibility at time $k+1$ leads to feasibility at $k+2, k+3, \dots$

Part (b): From Lemma 1 we have that $V(x_{k|k})$ is an upper bound on the cost function (13) at time k. Because (33a) is equivalent to $x_{k|k}^T P_k x_{k|k} \leq \alpha_k$, minimizing α_k implies minimizing the upper bound on the cost function.

Part (c): We now consider stability for when P is recalculated at every sampling instant, k. From Lemma 1, we have that applying the control law (34), leads to $x_{k+1|k}^T P_k x_{k+1|k} \leq x_{k|k}^T P_k x_{k|k}$. At the next sampling instant, the previous solution to the optimization problem is feasible, but not necessarily optimal, i.e. $x_{k+1|k+1}^T P_{k+1} x_{k+1|k+1} \leq x_{k+1|k+1}^T P_k x_{k+1|k+1}$. Combining these two inequalities, yields $x_{k+1|k+1}^T P_{k+1} x_{k+1|k+1} \leq x_{k|k}^T P_k x_{k|k}$. Thus, $x_{k|k}^T P_k x_{k|k}$ is a strictly decreasing Lyapunov function for the closed-loop system, implying that $x_{k|k} \to 0$ as $k \to \infty$.

It remains to show constraint satisfaction. Having satisfied the conditions of Lemma 1, we know that the state lies in the ellipsoid $\mathscr{E}_k = \{x_k \in \mathbb{R}^n : x_k^T P_k x_k \leq \alpha_k\}$. From Lemma 2, $\mathscr{E}_k$ lies in the constraint set $\mathscr{C}_k$ if (31) holds. It can be shown that (33a) is equivalent to (31). We let $u_{k+i|k} = K_k x_{k+i|k}$ replace u in (3), so that $\mathscr{C}_k = \{x_k \in \mathbb{R}^n : (c_j + d_j K_k)x_k \leq 1, j = 1, ..., r\}$, followed by some matrix manipulation. Because $\mathscr{E}_k$ is invariant, and contained in the constraint set $\mathscr{C}_k$, all states are guaranteed to satisfy input and state constraints. □

Remark 1. *Closed-loop stability, input and state constraint satisfaction are guaranteed probabilistically by feasibility of the linear matrix inequalities at initial time. Because we consider a confidence interval at every time step, the resulting probability may be smaller than that given by Assumption 1.*

Remark 2. *For some initial conditions the sector condition may be too restrictive, so that the set of linear matrix inequalities does not have a feasible solution. Reduction of the conservative sector using learning will increase the number of feasible initial conditions.*

5. SIMULATION RESULTS

To test the proposed control design, we consider the dynamics of a flexible link robotic arm, as given in e.g Böhm et al. (2009). The system is discretized using the forward Euler method, to obtain the same form as (1), resulting in

$$A = \begin{bmatrix} 1 & 1 & 0 & 0 \\ -48.6\Delta t & -0.25 & -48.6\Delta t & 0 \\ 0 & 0 & 1 & \Delta t \\ 19.5\Delta t & 0 & -16.7\Delta t & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ 21.6\Delta t \\ 0 \\ 0 \end{bmatrix}$$
$$G^T = [0\ 0\ 0\ -3.33\Delta t], \quad H = [0\ 0\ 1\ 0], \tag{35}$$

with time step $\Delta t = 0.01$s. For the robotic arm we have that $[x_1, x_2, x_3, x_4]^T = [\theta_1, \dot{\theta}_1, \theta_2, \dot{\theta}_2]^T$, where θ_1, θ_2 are angles, and $\dot{\theta}_1, \dot{\theta}_2$ are angle rates. The nonlinearity, given by

$$\gamma(z_k) = \sin(z_k) + z_k, \tag{36}$$

is assumed unknown, but estimated as described in Section 2. The following state and input constraints apply

$$u_k \in [-1.5, 1.5], \quad x_{1,k}, x_{3,k} \in [-\frac{\pi}{2}, \frac{\pi}{2}], \quad k \geq 0. \tag{37}$$

The control objective is to steer the system (35) to the origin. In the control design, we used matrices

$$Q = \text{diag}([1, 0.1, 1, 0.1]), \quad R = 0.1. \tag{38}$$

A conservative sector was defined using the prior distribution for w and σ^2 (8). The parameters of the prior distribution were specified such that the resulting conservative sector was contained within the first and third quadrant of the input coordinate plane. We used $w_0 = 3.0$ and $\sigma^2 V_0 = 2.25$, with $a_0 = 90$, $b_0 = 50$. The conservative sector was then defined using the 95% confidence interval for the prior distribution, and is visualized in Figure 2. The first simulation was run for initial condition $x_0 = [1.2, 0.1, 0.1, 0.1]$, with the conservative sector formulation. The closed-loop measurements were then sorted, and weighted according to (6). The weighted data set was then used to train the wBLR model, using the equations (10). A 95% confidence interval for the posterior distribution of w (11a), resulted in the following upper bound, $\hat{u} = 2.63$, and lower bound, $\hat{l} = 1.39$, for the learned sector. The learned sector is visualized in Figure 2.

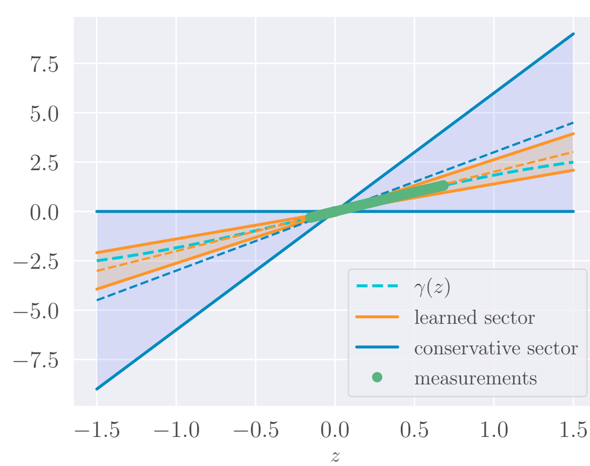

Fig. 2. A conservative (blue) and a learned (orange) sector bound on the nonlinear function, $\gamma(z)$.

For the next 20 simulations, the MPC algorithm was tested with two versions of the sector, namely

(1) the conservative sector, as defined above
(2) the sector learned offline with data from the first simulation

The optimization problem (32)-(33) was solved using *CVXOPT* in Python (Andersen et al., 2013). We tested 20 different initial conditions in a region around $x_0 = [1.2, 0, 0, 0]$, according to

$$\tilde{x}_{i,0} = x_{i,0} + [-\delta, \delta], \quad i = 1, \dots, 4 \tag{39}$$

with $\delta = 0.2$. The simulation results for when both sectors were feasible in closed-loop, are plotted in Figure 3 and 4. In order to compare the closed-loop performance, we calculated the closed-loop cost according to

$$\phi = \sum_{k=1}^{k_{\text{sim}}} x_k^T Q x_k + u_k^T R u_k, \tag{40}$$

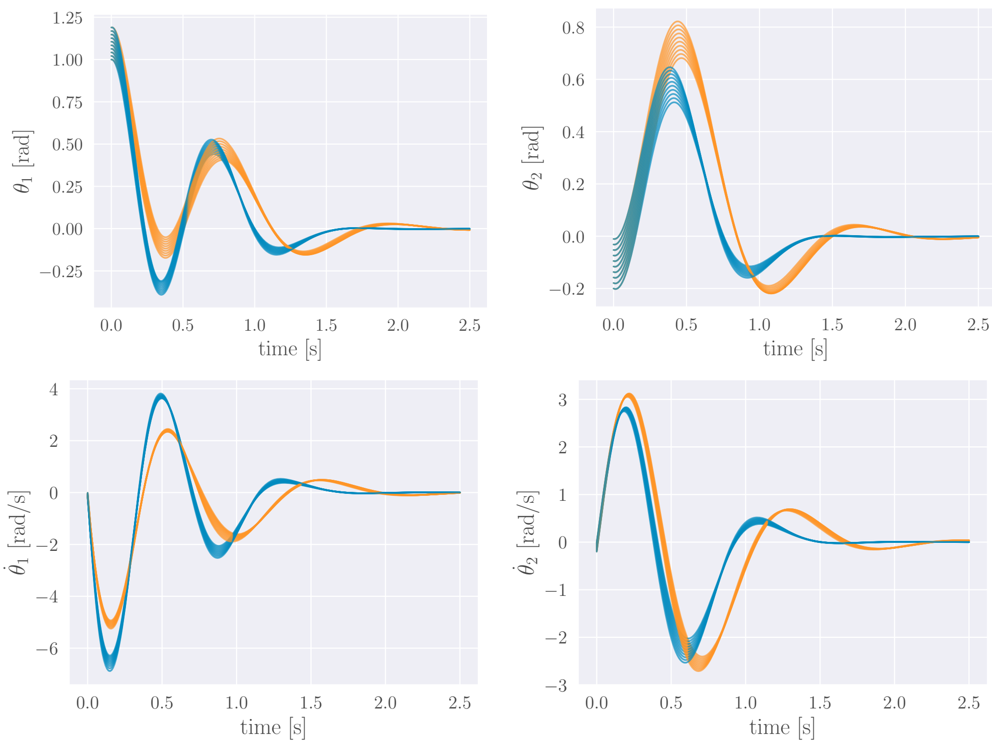

Fig. 3. Simulation results generated using (1) a conservative sector (blue), (2) a sector learned offline (orange).

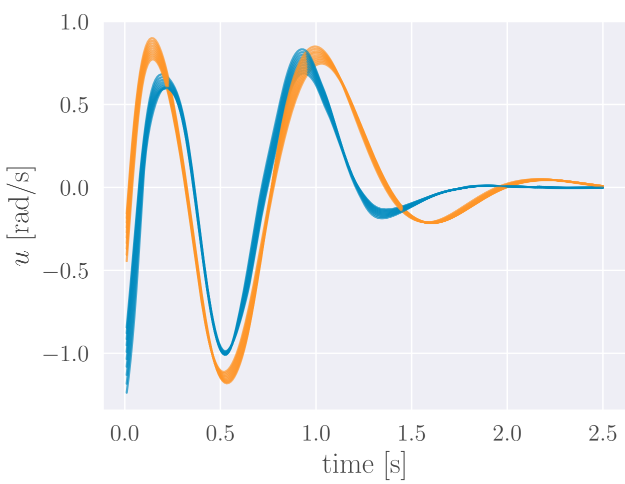

Fig. 4. Control inputs using a conservative (blue) and a learned (orange) sector bound on the nonlinear function, $\gamma(z)$.

resembling (14), but summed over the simulation length denoted by the time index k_{sim}. Using the conservative sector, fewer of the initial conditions rendered feasible optimization problems due to a more conservative sector bound in Theorem 1. Table 1 shows the mean of (40) for all simulations where both versions of the sectors were feasible, and in addition the total number of feasible simulations.

For the proposed control design optimization is performed at every time step in order to recalculate K. Because the resulting optimization problem is convex, this can be solved very efficiently. For high frequency systems, optimization may also be done at a lower frequency than control, maintaining stability but with a small performance loss in between updates of K.

Table 1. Comparison of closed-loop performance

MPC scheme w/	ϕ_{mean}	Num. of feasible simulations
(1) Conservative sector	171.97	10
(2) Learned sector	156.55	20

Simulations verify that the proposed design can be used for robust control of linear systems with sector-bounded uncertainties, where the sector bound is not initially known. The stochastic sector allows for a conservative initial formulation, based on a best guess on the uncertainty. By exploiting available closed-loop measurements, the uncertainty of the initial, stochastic sector can be reduced, resulting in a smaller sector bound. For the smaller sector the quadratic cost of the input and state is reduced, and the feasible region of the optimization problem is enlarged.

6. CONCLUSION

In this paper, we have shown how measurements in closed-loop can be used to define a stochastic sector condition, for robust control of linear systems with sector-bounded nonlinearities. A wBLR model is trained with closed-loop measurements of the nonlinearity, and used to reduce conservativeness of the robust controller. The proposed MPC design is tested in simulations of a flexible link robotic arm. Comparing the closed-loop performance for a conservative sector with a sector that incorporates learning, shows that the latter reduces the quadratic cost of the state and input over the simulation length, and results in a enlarged feasible region for the control optimization problem. The approach is currently limited to systems without disturbances and with a certain constraint formulation. Further work will aim to investigate similar control design for a broader class of Lur'e systems and constraint formulations.

ACKNOWLEDGEMENTS

This work was supported by the industry partners Borregaard, Elkem, Yara, Hydro and the Research Council of Norway through the project TAPI: Towards Autonomy in Process Industries, project number: 294544

REFERENCES

Andersen, M.S., Dahl, J., Vandenberghe, L., et al. (2013). Cvxopt: A python package for convex optimization. *Available at cvxopt. org*, 54.

Böhm, C., Findeisen, R., and Allgöwer, F. (2010). Robust control of constrained sector bounded Lur'e systems with applications to nonlinear model predictive control. *Dynamics of Continuous, Discrete and Impulsive Systems Series B: Applications and Algorithms*, 17(6), 935–958.

Böhm, C., Yu, S., Findeisen, R., and Allgöwer, F. (2009). Predictive control for Lure systems subject to constraints using LMIs. *European Control Conference*, 3389–3394.

Boyd, S., El Ghaoui, L., Feron, E., and Balakrishnan, V. (1994). *Linear matrix inequalities in system and control theory*. SIAM.

Hewing, L., Kabzan, J., and Zeilinger, M.N. (2019). Cautious model predictive control using gaussian process regression. *IEEE Transactions on Control Systems Technology*, 28(6), 2736–2743.

Kothare, M.V., Balakrishnan, V., and Morari, M. (1996). Robust constrained model predictive control using linear matrix inequalities. *Automatica*, 32(10), 1361–1379.

McKinnon, C.D. and Schoellig, A.P. (2019). Learn fast, forget slow: Safe predictive learning control for systems with unknown and changing dynamics performing repetitive tasks. *IEEE Robotics and Automation Letters*, 4(2), 2180–2187.

Murphy, K.P. (2012). *Machine learning: a probabilistic perspective*. MIT press.

Nguyen, H.H., Savchenko, A., Yu, S., and Findeisen, R. (2018). Improved Robust Predictive Control for Lur'e Systems Using Set-based Learning. *IFAC-PapersOnLine (Proc. 6th IFAC Conference on Nonlinear Model Predictive Control NMPC)*, 51(20), 487–492.

Ostafew, C.J., Schoellig, A.P., and Barfoot, T.D. (2014). Learning-based nonlinear model predictive control to improve vision-based mobile robot path-tracking in challenging outdoor environments. *Proc. IEEE International Conference on Robotics and Automation*, 4029–4036.

Soloperto, R., Müller, M.A., Trimpe, S., and Allgöwer, F. (2018). Learning-Based Robust Model Predictive Control with State-Dependent Uncertainty. *IFAC-PapersOnLine (Proc. 6th IFAC Conference on Nonlinear Model Predictive Control NMPC)*, 51(20), 442–447.