Robust Adaptive Model Predictive Control

Mark Cannon

Department of Engineering Science, Oxford University

November 25, 2020

Motivation

Robust MPC paradigm:

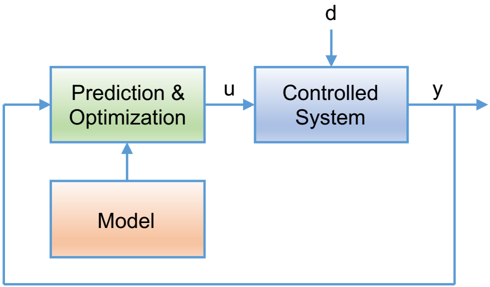

- Uncertain model & disturbances affect performance
- Large effort (time & money) spent on model identification offline

Motivation

Adaptive MPC paradigm:

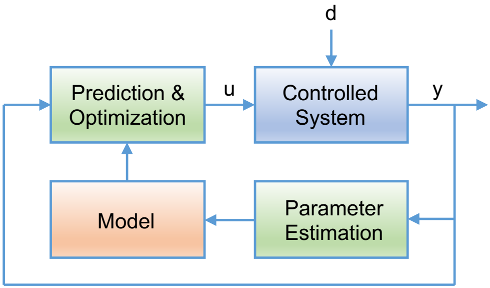

- Identify (or learn) model (or cost or constraints) online
- Require: robust constraint satisfaction
  closed loop stability & performance guarantees
  parameter convergence

Applications

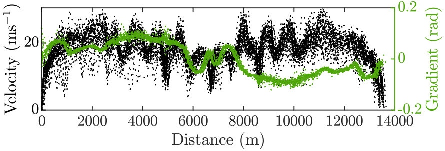

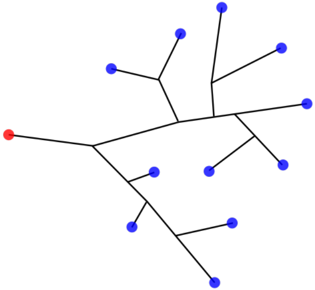

- Uncertain parameters, uncertain demand
- Networks of interacting locally controlled systems

Overview

An idea with a long history: e.g. self-tuning control, DMC, GPC . . .

[Clarke, Tuffs, Mohtadi, 1987]

Revisited with new tools:

- Set membership estimation

  [Bai, Cho, Tempo, 1998]

- Robust tube MPC

  [Langson, Chryssochoos, Rakovic, Mayne, 2004]

- Dual adaptive/predictive control

  [Lee & Lee, 2009]

Overview

Recent work on MPC with model adaptation

- Online learning & identification:
  - Persistence of excitation constraints

    [Marafioti, Bitmead, Hovd, 2014]
  - RLS parameter estimation with covariance matrix in cost

    [Heirung, Ydstie, Foss, 2017]
  - Gaussian process regression, particle filtering

    [Klenske, Zeilinger, Scholkopf, Hennig, 2016]
    [Bayard & Schumitzky, 2010]

- Robust constraint satisfaction and performance:
  - Constraints based on prior uncertainty set, online update of cost only

    [Aswani, Gonzalez, Sastry, Tomlin, 2013]
  - Set-based identification, stable FIR plant model

    [Tanaskovic, Fagiano, Smith, Morari, 2014]

Overview

This talk:

1. Set membership parameter estimation
2. Polytopic tube robust adaptive MPC
3. Persistent excitation
4. Differentiable MPC

Parameter set estimate

Plant model with unknown parameter vector $\theta^\star$ and disturbance w:

$$x_{k+1} = A(\theta^\star)x_k + B(\theta^\star)u_k + w_k$$

Assumption 1: model is affine in unknown parameters

$$x_{k+1} = D_k\theta^\star + d_k + w_k \qquad \begin{cases} D_k = D(x_k, u_k) \\ d_k = A_0 x_k + B_0 u_k \end{cases}$$

Assumption 2: stochastic disturbance $w_k \in \mathcal{W}$

$\mathcal{W} \ni 0$ is compact and convex

Unfalsified set: If x_k, x_{k-1}, u_{k-1} are known, then $\theta^\star \in \Delta_k$

$$\Delta_k = \{\theta : x_k = D_{k-1}\theta + d_{k-1} + w,\ w \in \mathcal{W}\}$$

Minimal parameter set estimate

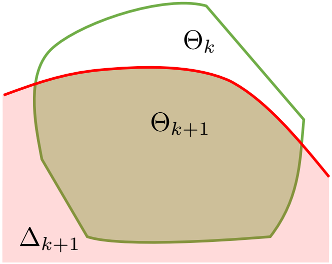

Minimal parameter set update:

$$\Theta_{k+1} = \Theta_k \cap \Delta_{k+1}$$

Assumption 3: $\mathcal{W}$ is a 'tight' bound: for all $w^0 \in \partial\mathcal{W}$ and $\epsilon > 0$

$$\Pr\{\|w_k - w^0\| < \epsilon\} \geq p_w(\epsilon)$$

where $p_w(\epsilon) > 0 \;\forall \epsilon > 0$

Assumption 4: persistent excitation: $\exists\ \alpha,\ \beta > 0,\ N$ such that

$$\|D_k\| \leq \alpha \quad \text{and} \quad \sum_{j=k}^{k+N-1} D_j^\top D_j \succeq \beta I \text{ for all } k$$

Minimal parameter set estimate

Unfalsified set: $$\begin{aligned}\Delta_{k+1} &= \{\theta : x_{k+1} - D_k\theta - d_k \in \mathcal{W}\} \\ &= \{\theta : D_k(\theta^* - \theta) + w_k \in \mathcal{W}\}\end{aligned}$$

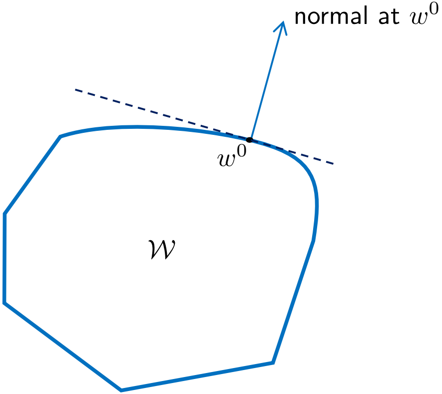

Minimal parameter set estimate

Unfalsified set: $\Delta_{k+1} = \{\theta : x_{k+1} - D_k\theta - d_k \in \mathcal{W}\}$
$= \{\theta : D_k(\theta^* - \theta) + w_k \in \mathcal{W}\}$

For any given $\theta^0 \in \Theta_k$:

- pick $w^0 \in \partial\mathcal{W}$ so that $D_k(\theta^* - \theta^0)$ is normal to $\partial\mathcal{W}$ at w^0
- let $\epsilon = \|w_k - w^0\|$

then $\theta^0 \notin \Delta_{k+1}$ if $\epsilon < \|D_k(\theta^* - \theta^0)\|$

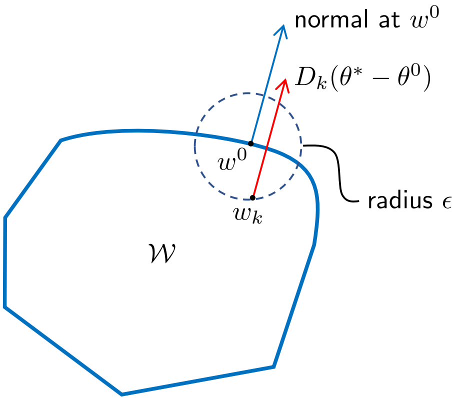

Minimal parameter set estimate

If Assumptions 1-4 hold, then $\Theta_k \to \{\theta^*\}$ as $k \to \infty$ w.p. 1

This follows from:

- **A** For any $\theta^0 \in \Theta_k$, if $\|\theta^* - \theta^0\| \geq \epsilon$, then

$$\Pr\{\theta^0 \notin \Delta_j\} \geq p_w\big(\epsilon\sqrt{\beta/N}\big)$$

for all k, all $\epsilon > 0$, and some $j \in \{k+1, \ldots, k+N\}$

- **B** For any $\theta^0 \in \Theta_0$ such that $\|\theta^0 - \theta^*\| \geq \epsilon$,

$$\Pr\{\theta^0 \in \Theta_k\} \leq \Big[1 - p_w\big(\epsilon\sqrt{\beta/N}\big)\Big]^{\lfloor k/N \rfloor}$$

for all k and all $\epsilon > 0$, so

$$\sum_{k=0}^{\infty} \Pr\{\theta^0 \in \Theta_k\} = 0 \overset{\text{Borel-Cantelli}}{\underset{\text{Lemma}}{\Longrightarrow}} \Pr\{\theta^0 \in \bigcap_{k=0}^{\infty} \Theta_k\} = 0$$

Minimal parameter set estimate

The complexity of Θ_k is unbounded in general

e.g. Minimal parameter set Θ_k for $k = 1, \ldots, 6$ with polytopic $\mathcal{W}$ and Θ_0

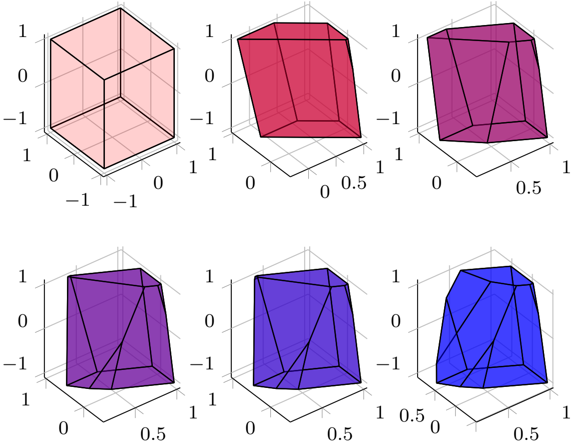

Fixed complexity polytopic parameter set estimate

- Define $\Theta_k = \{\theta : H_\Theta \theta \le h_k\}$ for a fixed matrix H_Θ
- Update Θ_{k+1} by solving, for each row i:

$$[h_{k+1}]_i = \max_{\substack{w_0 \in \mathcal{W}, \dots, w_{N-1} \in \mathcal{W} \\ \theta \in \Theta_k}} [H_\Theta]_i \theta$$

subject to

$$x_{k-N+2} = D_{k-N+1}\theta + d_{k-N+1} + w_0$$
$$\vdots$$
$$x_{k+1} = D_k \theta + d_k + w_{N-1}$$

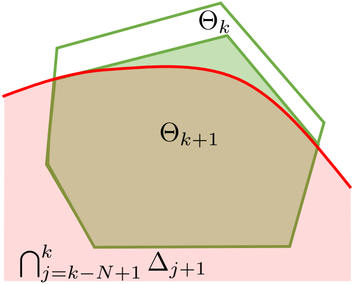

- Then $\Theta_{k+1} \subseteq \Theta_k \subseteq \cdots \subseteq \Theta_0$,
  and Θ_{k+1} is the minimum volume set such that

$$\Theta_{k+1} \supseteq \Theta_k \cap \bigcap_{j=k-N+1}^{k} \Delta_{j+1}$$

Fixed complexity polytopic parameter set estimate

If Assumptions 1-4 hold, then $\Theta_k \to \{\theta^*\}$ as $k \to \infty$ w.p. 1

This follows from:

- **A** If $[h_k]_i - [H_\Theta]_i \theta^* \geq \epsilon$, then

$$\Pr\left\{\{\theta : [H_\Theta]_i \theta = [h_k]_i\} \cap \bigcap_{j=k-N+1}^{k} \Delta_{j+1} = \emptyset\right\} \geq \left[p_w\left(\frac{\epsilon\beta}{\alpha N}\right)\right]^N$$

  for all i, k, and all $\epsilon > 0$

- **B** For any θ^0 such that $[H_\Theta]_i(\theta^0 - \theta^*) \geq \epsilon$ for some row i,

$$\Pr\{\theta^0 \in \Theta_k\} \leq \left\{1 - \left[p_w\left(\frac{\epsilon\beta}{N\alpha}\right)\right]^N\right\}^{\lfloor k/N \rfloor}$$

  for all k and all $\epsilon > 0$, so

$$\sum_{k=0}^{\infty} \Pr\{\theta^0 \in \Theta_k\} = 0 \;\overset{\text{Borel-Cantelli}}{\underset{\text{Lemma}}{\Longrightarrow}}\; \Pr\{\theta^0 \in \bigcap_{k=0}^{\infty} \Theta_k\} = 0$$

Example: fixed complexity parameter set estimate

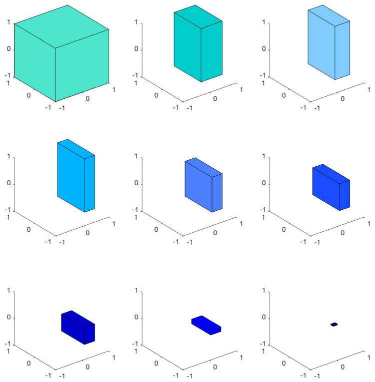

Parameter set Θ_k at time
$k \in \{0, 1, 2; 10, 25, 50; 100, 500, 5000\}$

Θ set	Volume (%)	Cost*
Θ_0	100	62.22
Θ_1	26.1	61.13
Θ_2	18.3	61.03
Θ_{10}	12.7	60.96
Θ_{25}	8.3	60.93
Θ_{50}	6.3	60.77
Θ_{100}	3.4	59.45
Θ_{500}	0.7	57.94
Θ_{5000}	0.0089	53.95
$\theta^\star$	-	52.70

Volume of Θ_k and Cost* for same x_0

Inexact disturbance bounds

What if $\mathcal{W}$ is not exactly known?

Suppose $w_k \in \widehat{\mathcal{W}}$ for all k, for known $\widehat{\mathcal{W}}$

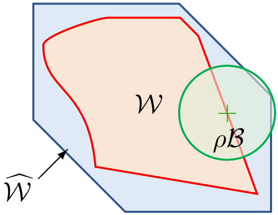

Assumption 5: $\widehat{\mathcal{W}}$ is compact and convex, and $\mathcal{W} \subseteq \widehat{\mathcal{W}} \subseteq \mathcal{W} \oplus \rho\mathcal{B}$ for some $\rho \geq 0$, and $\mathcal{B} = \{x : \|x\| \leq 1\}$

Replace $\mathcal{W}$ with $\widehat{\mathcal{W}}$ in the fixed complexity polytopic parameter set update

then $\theta^\star \in \widehat{\Delta}_{k+1} = \{\theta : x_{k+1} = D_k\theta + d_k + w,\ w \in \widehat{\mathcal{W}}\}$, and

if Assumptions 1-5 hold, then $\Theta_k \to \{\theta^*\} \oplus \rho\sqrt{N/\beta}\,\mathcal{B}$ as $k \to \infty$ w.p. 1

Inexact disturbance bounds

What if $\mathcal{W}$ is not exactly known?

Suppose $w_k \in \widehat{\mathcal{W}}$ for all k, for known $\widehat{\mathcal{W}}$

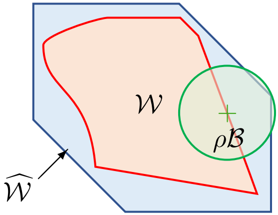

Assumption 5: $\widehat{\mathcal{W}}$ is compact and convex, and $\mathcal{W} \subseteq \widehat{\mathcal{W}} \subseteq \mathcal{W} \oplus \rho\mathcal{B}$ for some $\rho \geq 0$, and $\mathcal{B} = \{x : \|x\| \leq 1\}$

Replace $\mathcal{W}$ with $\widehat{\mathcal{W}}$ in the fixed complexity polytopic parameter set update

then $\theta^\star \in \widehat{\Delta}_{k+1} = \{\theta : x_{k+1} = D_k\theta + d_k + w,\ w \in \widehat{\mathcal{W}}\}$, and

if Assumptions 1-5 hold, then $\Theta_k \to \{\theta^*\} \oplus \rho\sqrt{N/\beta}\,\mathcal{B}$ as $k \to \infty$ w.p. 1

Noisy measurements

Let $y_k = x_k + s_k$ be an estimate of x_k

Assumption 6: i.i.d. noise $s_k \in \mathcal{S}$ for all k
where $\mathcal{S} \ni 0$ is a compact, convex polytope

Assumption 7: the noise bound is tight, i.e. for all $s^0 \in \partial\mathcal{S}$ and $\epsilon > 0$

$$\Pr\big\{\|s_k - s^0\| < \epsilon\big\} \geq p_s(\epsilon)$$

where $p_s(\epsilon) > 0$ for all $\epsilon > 0$

Then $\mathcal{S} = \mathrm{co}\{s^{(1)}, \ldots, s^{(h)}\}$ implies $\theta^\star \in \mathrm{co}\{\widehat{\Delta}^{(1)}_{k+1}, \ldots, \widehat{\Delta}^{(h)}_{k+1}\}$, where

$$\widehat{\Delta}^{(j)}_{k+1} = \Big\{\theta : y_{k+1} - D(y_k - s^{(j)}, u_k)\theta - d(y_k - s^{(j)}_k, u_k) \in \widehat{\mathcal{W}} \oplus \mathcal{S}\Big\}$$

If Assumptions 1-7 hold, then $\Theta_k \to \{\theta^*\} \oplus \rho\sqrt{N/\beta}\,\mathcal{B}$ as $k \to \infty$ w.p. 1

Parameter point estimate

Define a point estimate $\hat{\theta}_k$ of $\theta^\star$

$\hat{\theta}_k$:	defines a nominal predicted performance index
Θ_k:	enforces constraints robustly

Given a parameter estimate $\hat{\theta}_k$:

- Least mean squares (LMS) filter estimate update is

$$\tilde{\theta}_{k+1} = \hat{\theta}_k + \mu D^\top(x_k, u_k)(x_{k+1} - \hat{x}_{1|k})$$
$$\hat{\theta}_{k+1} = \Pi_{\Theta_{k+1}}(\tilde{\theta}_{k+1})$$

  where

  - $\hat{x}_{1|k} = D(x_k, u_k)(\hat{\theta}_k)$
  - $\mu > 0$ satisfies $1/\mu > \sup_{(x,u)\in\mathcal{Z}} \|D(x,u)\|^2$
  - $\Pi_\Theta(\hat{\theta}) = \arg\min_{\theta\in\Theta} \|\theta - \hat{\theta}\|$ projects onto Θ

- For $\mu = 0$ the update is $\hat{\theta}_{k+1} = \Pi_{\Theta_{k+1}}(\hat{\theta}_k)$

Parameter point estimate

The LMS filter ($\mu > 0$) ensures the l^2 gain bound:

If $\sup_{k \in \mathbb{N}} \|x_k\| < \infty$ and $\sup_{k \in \mathbb{N}} \|u_k\| < \infty$, then $\hat{\theta}_k \in \Theta_k$ for all k and

$$\sup_{T \in \mathbb{N}, w_k \in \mathcal{W}, \hat{\theta}_0 \in \Theta_0} \frac{\sum_{k=0}^{T} \|\tilde{x}_{1|k}\|^2}{\frac{1}{\mu}\|\hat{\theta}_0 - \theta^\star\|^2 + \sum_{k=0}^{T} \|w_k\|^2} \leq 1$$

where $\tilde{x}_{1|k} = A(\theta^\star)x_k + B(\theta^\star)u_k - \hat{x}_{1|k}$ is the 1-step prediction error

Control Problem

Consider robust regulation of the system

$$x_{k+1} = A(\theta)x_k + B(\theta)u_k + w_k$$

with $\theta \in \Theta_k$, $w_k \in \mathcal{W}$, subject to the state and control constraints

$$Fx_k + Gu_k \leq \mathbf{1} = [1 \quad \cdots \quad 1]^\top$$

Assumption (Robust stabilizability):
There exists a set $\mathcal{X} = \{x : Vx \leq \mathbf{1}\}$ and feedback gain K such that $\mathcal{X}$ is λ-contractive for some $\lambda \in [0, 1)$, i.e.

$$V\Phi(\theta)x \leq \lambda\mathbf{1}, \quad \text{for all } x \in \mathcal{X}, \theta \in \Theta_0.$$

where $\Phi(\theta) = A(\theta) + B(\theta)K$.

Control Problem

- State sequence predicted at time k: $x_{1|k}, x_{2|k}, \ldots$
- Control sequence predicted at time k: $u_{0|k}, u_{1|k}, \ldots$:

$$u_{i|k} = \begin{cases} Kx_{i|k} + v_{i|k} & i = 0, 1, \ldots, N-1 \\ Kx_{i|k} & i = N, N+1, \ldots \end{cases}$$

where $\mathbf{v} = (v_{0|k}, \ldots, v_{N|k})$ is a decision variable

Nominal predicted performance index

$$J_N(x_k, \hat{\theta}_k, \mathbf{v}_k) = \sum_{i=0}^{N-1} \Big(\|\hat{x}_{i|k}\|_Q^2 + \|\hat{u}_{i|k}\|_R^2 \Big) + \|\hat{x}_{N|k}\|_P^2$$

where $\hat{x}_{0|k} = x_k$

$$\hat{u}_{i|k} = K\hat{x}_{i|k} + v_{i|k}$$

$$\hat{x}_{i+1|k} = A(\hat{\theta}_k)\hat{x}_{i|k} + B(\hat{\theta}_k)\hat{u}_{i|k}, \qquad \hat{\theta}_k = \text{nominal estimate}$$

and $P \succeq \Phi^\top(\theta) P \Phi(\theta) + Q + K^\top R K$ for all $\theta \in \Theta_k$

Tube MPC

A sequence of sets (a “tube”) is constructed to bound the predicted state $x_{i|k}$, with ith cross section, $\mathcal{X}_{i|k}$:

$$\mathcal{X}_{i|k} = \{x : Vx \le \alpha_{i|k}\}$$

where V is determined offline and $\alpha_{i|k}$ are online decision variables

(A) For robust satisfaction of $x_{i|k} \in \mathcal{X}_{i|k}$, we require

$$V\Phi(\theta)x + VB(\theta)v_{i|k} + \bar{w} \le \alpha_{i+1|k} \quad \text{for all } x \in \mathcal{X}_{i|k},\ \theta \in \Theta_k$$

where $[\bar{w}]_i = \max_{w\in\mathcal{W}}[V]_i w$

(B) For robust satisfaction of $Fx_{i|k} + Gu_{i|k} \le \mathbf{1}$, we require

$$(F + GK)x + Gv_{i|k} \le \mathbf{1} \quad \text{for all } x \in \mathcal{X}_{i|k}$$

Condition (A) is bilinear in x and θ, but can be expressed in terms of linear inequalities using a vertex representation of either $\mathcal{X}_{i|k}$ or Θ_k

Tube MPC

We generate the vertex representation:

$$\mathcal{X}_{i|k} = \mathrm{co}\{x^1_{i|k}, \ldots x^m_{i|k}\}$$

using the fact that $\{x : [V]_r x \le [\alpha_{i|k}]_r\}$ is a supporting hyperplane of $\mathcal{X}_{i|k}$

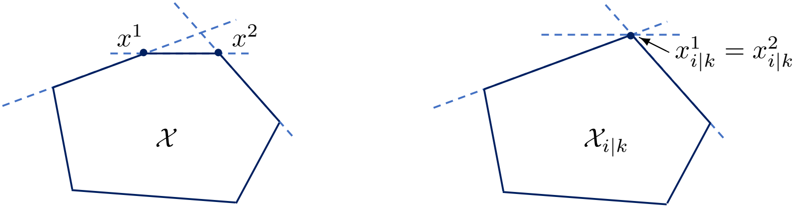

Hence each vertex $x^j_{i|k}$ is defined by a fixed set of rows of V, so

$$x^j_{i|k} = U^j \alpha_{i|k}$$

where U^j is determined offline from the vertices of $\mathcal{X} = \{x : Vx \le \mathbf{1}\}$

Tube MPC

Using the hyperplane and vertex descriptions of $\mathcal{X}_{i|k}$, the robust tube constraints become

- (A) $V\Phi(\theta)U^j\alpha_{i|k} + VB(\theta)v_{i|k} + \bar{w} \le \alpha_{i+1|k}$ for all $\theta \in \Theta_k$, $j = 1, \dots, m$
- (B) $(F + GK)U^j\alpha_{i|k} + Gv_{i|k} \le \mathbf{1}$, $j = 1, \dots, m$

Now condition (B) is linear and (A) can be equivalently written as linear constraints using

Polyhedral set inclusion lemma

Let $\mathcal{P}_i = \{x : F_i x \le f_i\} \subset \mathbb{R}^n$ for $i = 1, 2$. Then $\mathcal{P}_1 \subseteq \mathcal{P}_2$ iff

$$\exists \Lambda \ge 0 \text{ such that } \Lambda F_1 = F_2 \text{ and } \Lambda f_1 \le f_2$$

Robust MPC online optimization problem

Summary of constraints in the online MPC optimization at time k:

$$V x_k \leq \alpha_{0|k}$$

$$\Lambda^j_{i|k} H_\Theta = V D(U^j \alpha_{i|k}, K U^j \alpha_{i|k} + v_{i|k})$$
$$\Lambda^j_{i|k} h_k \leq \alpha_{i+1|k} - V d(u^j \alpha_{i|k}, K U^j \alpha_{i|k} + v_{i|k}) - \bar{w}$$
$$\Lambda^j_{i|k} \geq 0$$

$$(F + GK) U^j \alpha_{i|k} + G v_{i|k} \leq \mathbf{1}$$

$$\Lambda^j_{N|k} H_\Theta = V D(U^j \alpha_{N|k}, K U^j \alpha_{N|k})$$
$$\Lambda^j_{N|k} h_k \leq \alpha_{N|k} - V d(u^j \alpha_{N|k}, K U^j \alpha_{N|k}) - \bar{w}$$
$$\Lambda^j_{N|k} \geq 0$$

$$(F + GK) U^j \alpha_{N|k} \leq \mathbf{1}$$

for $i = 0, \ldots, N-1$, $j = 1, \ldots, m$

Let $\mathcal{F}(x_k, \Theta_k)$ be the feasible set for the decision variables $\mathbf{v}_k, \boldsymbol{\alpha}_k, \boldsymbol{\Lambda}_k$

Robust adaptive MPC algorithm

Offline: Choose Θ_0, $\mathcal{X}$, feedback gain K, and compute P

Online, at each time $k = 1, 2, \ldots$:

1. Given x_k, update Θ_k and $\hat{\theta}_k$
2. Compute the solution $(\mathbf{v}_k^*, \boldsymbol{\alpha}_k^*, \boldsymbol{\Lambda}_k^*)$ of the QP:

$$\min_{\mathbf{v}_k, \boldsymbol{\alpha}_k, \boldsymbol{\Lambda}_k} J(x_k, \hat{\theta}_k, \mathbf{v}_k)$$
$$\text{subject to } (\mathbf{v}_k, \boldsymbol{\alpha}_k, \boldsymbol{\Lambda}_k) \in \mathcal{F}(x_k, \Theta_k)$$

3. Apply the control law $u_k^* = K x_k + v_{0|k}^*$

Robust adaptive MPC algorithm

The MPC algorithm has the following closed loop properties:

If $\theta^\star \in \Theta_0$ and $\mathcal{F}(x_0, \Theta_0) \neq \emptyset$, then for all $k > 0$:

1. $\theta^\star \in \Theta_k$
2. $\mathcal{F}(x_k, \Theta_k) \neq \emptyset$
3. $Fx_k + Gu_k \leq \mathbf{1}$

If $\mu > 0$, then

4. the closed loop system is finite-gain l^2-stable, i.e.

$$\sum_{k=0}^{T} \|x_k\|^2 \leq c_0\|x_0\|^2 + c_1\|\hat{\theta}_0 - \theta^\star\|^2 + c_2 \sum_{k=0}^{T} \|w_k\|^2$$

for some constants $c_0, c_1, c_2 > 0$, for all $T > 0$

Robust adaptive MPC algorithm

The MPC algorithm has the following closed loop properties:

If $\theta^\star \in \Theta_0$ and $\mathcal{F}(x_0, \Theta_0) \neq \emptyset$, then for all $k > 0$:

1. $\theta^\star \in \Theta_k$
2. $\mathcal{F}(x_k, \Theta_k) \neq \emptyset$
3. $Fx_k + Gu_k \leq \mathbf{1}$

If $\mu = 0$, then

4. the closed loop system is input-to-state stable (ISS)

$$\|x_T\| \leq \eta(\|x_0\|, T) + \zeta(\|\hat{\theta}_0 - \theta^*\|) + \psi(\max_{k \in \{0,\ldots,T-1\}} \|w_k\|)$$

for some $\mathcal{KL}$-function η, some $\mathcal{K}$-functions ψ, ζ and all k, T.

Regulation example

Linear system with

$$(A(\theta), B(\theta)) = (A_0, B_0) + \sum_{i=1}^{3} (A_i, B_i)\theta_i$$

$$A_0 = \begin{bmatrix} 0.5 & 0.2 \\ -0.1 & 0.6 \end{bmatrix} \quad A_1 = \begin{bmatrix} 0.042 & 0 \\ 0.072 & 0.03 \end{bmatrix} \quad A_2 = \begin{bmatrix} 0.015 & 0.019 \\ 0.009 & 0.035 \end{bmatrix} \quad A_3 = 0_{2\times 2}$$

$$B_0 = \begin{bmatrix} 0 \\ 0.5 \end{bmatrix} \quad B_1 = 0_{2\times 1} \quad B_2 = 0_{2\times 1} \quad B_3 = \begin{bmatrix} 0.0397 \\ 0.059 \end{bmatrix}$$

- true parameter $\theta^\star = [0.8 \;\; 0.2 \;\; -0.5]^\top$, initial set $\Theta_0 = \{\theta : \|\theta\|_\infty \leq 1\}$.
- disturbance uniformly distributed on $\mathcal{W} = \{w \in \mathbb{R}^2 : \; \|w\|_\infty \leq 0.1\}$
- state and input constraints: $[x]_2 \geq -0.3$, $u \leq 1$.

Regulation example: constraint satisfaction

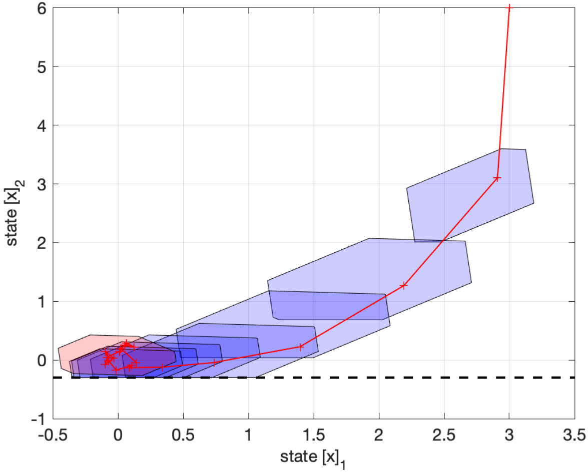

red: Closed loop trajectory from initial condition $x_0 = (3, 6)$

blue: Predicted state tube at time $k = 0$

pink: Terminal set

Tracking example

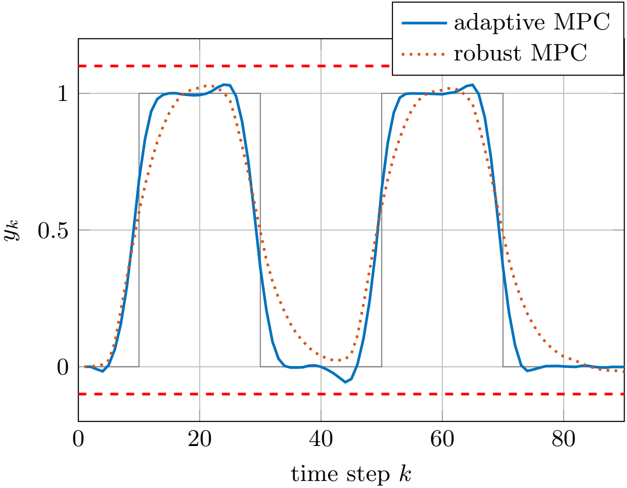

Closed loop setpoint tracking with and without model updates

Time-varying parameters

Assumption (time-varying parameters)

There exists a constant r_θ *such that the parameter vector* $\theta_k^\star$ *satisfies* $\theta_k^\star \in \Theta_0$ *for all* k *and* $\|\theta_{k+1}^\star - \theta_k^\star\| \leq r_\theta$

Define the dilation operator:

$$R_k(\Theta) = \{\theta : H_\Theta \theta \leq h + k r_\theta \mathbf{1}\}$$

Then the minimal parameter set at $k+1$ is

$$\Theta_{k+1} = R_1(\Theta_k \cap \Delta_{k+1}) \cap \Theta_0$$

and Θ_k is replaced in the tube MPC constraints by

$$\Theta_{i|k} = R_i(\Theta_k) \cap \Theta_0$$

Robust adaptive MPC algorithm with time-varying parameters

Parameter estimate bounds and recursive feasibility properties are unchanged:

Theorem (Closed loop properties)

If $\theta^\star \in \Theta_0$ and $\mathcal{F}(x_0, \Theta_0) \neq \emptyset$, then for all $k > 0$:

1. $\theta^\star \in \Theta_k$
2. $\mathcal{F}(x_k, \Theta_k) \neq \emptyset$
3. $Fx_k + Gu_k \leq \mathbf{1}$

But the LMS filter has an additional tracking error, which invalidates the l^2-stability properties, i.e. "certainty equivalence" no longer applies

Time-varying parameters example

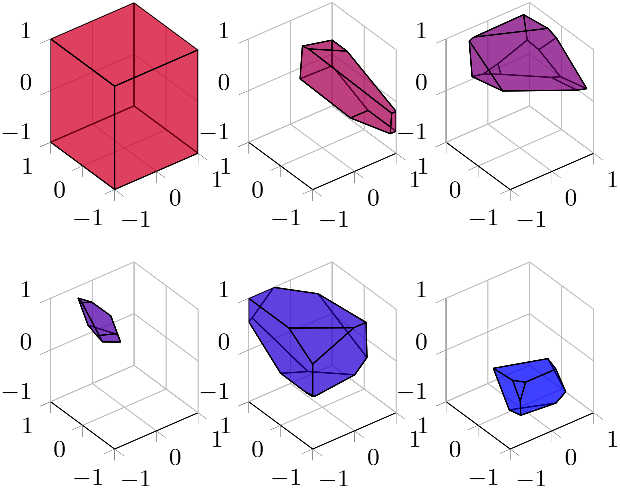

Parameter set Θ_k at time $k \in \{0, 100, 200, 300, 400, 500\}$ for time-varying system with $r_\theta = 0.01$

Time-varying parameters example

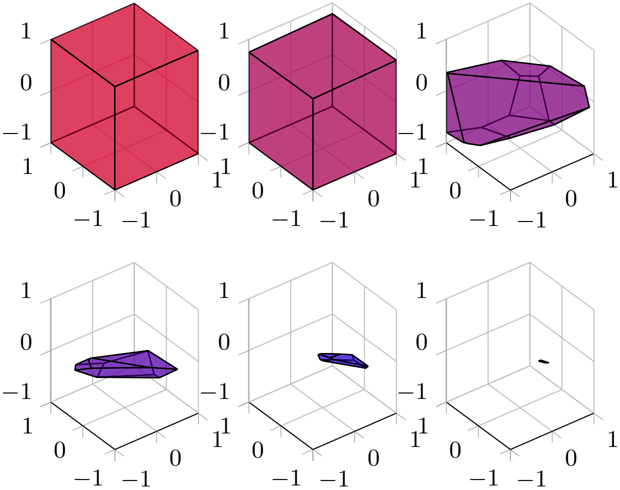

Parameter set Θ_k at time $k \in \{0, 5, 25, 70, 120, 500\}$ for time-invariant system for comparison

Persistent excitation

PE condition evaluated over a future horizon is nonconvex in $u_{i|k}, x_{i|k}$:

$$\text{(PE):} \qquad \sum_{i=0}^{N_p-1} D^\top(x_{i|k}, u_{i|k})\, D(x_{i|k}, u_{i|k}) \succeq \beta I$$

Linearise:

- let $(x, u) = (\bar{x}, \bar{u}) + (\check{x}, \check{u})$ where $\bar{x}_{0|k} = x_k$ and

$$\bar{u}_{i|k} = K\bar{x}_{i|k} + v^*_{i+1|k-1}$$
$$\bar{x}_{i+1|k} = A(\hat{\theta}_k)\bar{x}_{i|k} + B(\hat{\theta}_k)\bar{u}_{i|k}$$

- then $D = \bar{D} + \check{D}$, where $\bar{D} = D(\bar{x}, \bar{u})$, $\check{D} = D(\check{x}, \check{u})$

$$\begin{aligned} D^\top D &= \check{D}^\top \bar{D} + \bar{D}^\top \check{D} + \bar{D}^\top \bar{D} + \check{D}^\top \check{D} \\ &\succeq \check{D}^\top \bar{D} + \bar{D}^\top \check{D} + \bar{D}^\top \bar{D} \end{aligned}$$

- so $\quad \check{D}^\top \bar{D} + \bar{D}^\top \check{D} + \bar{D}^\top \bar{D} \succeq \beta I \quad \Longrightarrow \quad D^\top D \succeq \beta I$

Persistent excitation

PE condition evaluated over a future horizon is nonconvex in $u_{i|k}, x_{i|k}$:

$$\text{(PE):} \qquad \sum_{i=0}^{N_p-1} D^\top(x_{i|k}, u_{i|k})\, D(x_{i|k}, u_{i|k}) \succeq \beta I$$

Linearise:

- ⋆ let $(x, u) = (\bar{x}, \bar{u}) + (\check{x}, \check{u})$ where $\bar{x}_{0|k} = x_k$ and

$$\bar{u}_{i|k} = K\bar{x}_{i|k} + v^*_{i+1|k-1}$$
$$\bar{x}_{i+1|k} = A(\hat{\theta}_k)\bar{x}_{i|k} + B(\hat{\theta}_k)\bar{u}_{i|k}$$

- ⋆ then $D = \bar{D} + \check{D}$, where $\bar{D} = D(\bar{x}, \bar{u})$, $\check{D} = D(\check{x}, \check{u})$

$$D^\top D = \check{D}^\top \bar{D} + \bar{D}^\top \check{D} + \bar{D}^\top \bar{D} + \check{D}^\top \check{D}$$
$$\succeq \check{D}^\top \bar{D} + \bar{D}^\top \check{D} + \bar{D}^\top \bar{D}$$

- ⋆ so $\quad \check{D}^\top \bar{D} + \bar{D}^\top \check{D} + \bar{D}^\top \bar{D} \succeq \beta I \quad \Longrightarrow \quad D^\top D \succeq \beta I$

Persistent excitation

- A sufficient condition for $\sum_{i=0}^{N_p-1} D_{i|k}^\top D_{i|k} \succeq \beta I$ is

$$\text{(PE-LMI):} \qquad \sum_{i=0}^{N_p-1} \left(\check{D}_{i|k}^\top \bar{D}_{i|k} + \bar{D}_{i|k}^\top \check{D}_{i|k} + \bar{D}_{i|k}^\top \bar{D}_{i|k}\right) \succeq \beta I.$$

- This can be expressed in terms of

$$\check{x}_{i|k} \in \mathcal{X}_{i|k} - \{\bar{x}_{i|k}\}$$
$$\check{u}_{i|k} \in K(\mathcal{X}_{i|k} - \{\bar{x}_{i|k}\}) + \{v_{i|k}\} - \{v^*_{i+1|k-1}\}$$

using

$$\check{D}_{i|k} \in \mathrm{co}\Big\{D\big(U^j \alpha_{i|k} - \bar{x}_{i|k}, K(U^j \alpha_{i|k} - \bar{x}_{i|k}) + v_{i|k} - v^*_{i+1|k-1}\big)\Big\}$$

Hence (PE-LMI) is equivalent to an LMI in variables $\mathbf{v}_k, \boldsymbol{\alpha}_k, \beta$

Robust adaptive multiobjective MPC algorithm

Offline: Choose Θ_0, $\mathcal{X}$, γ, N_p, K, and compute P

Online, at each time $k = 1, 2, \ldots$:

1. Given x_k, update set (Θ_k) and point $(\hat{\theta}_k)$ parameter estimates, and compute $\bar{x}_{i|k}, \bar{u}_{i|k}$, $i = 0, \ldots, N-1$
2. Compute the solution $(\mathbf{v}_k^*, \boldsymbol{\alpha}_k^*, \boldsymbol{\Lambda}_k^*)$ of the semidefinite program

$$\min_{\mathbf{v}_k, \boldsymbol{\alpha}_k, \boldsymbol{\Lambda}_k, \beta} J(x_k, \hat{\theta}_k, \mathbf{v}_k) - \gamma\beta$$

subject to $(\mathbf{v}_k, \boldsymbol{\alpha}_k, \boldsymbol{\Lambda}_k) \in \mathcal{F}(x_k, \Theta_k)$ and (PE-LMI)

3. Apply the control law $u_k^* = Kx_k + v_{0|k}^*$

How to choose γ? Stability? Closed loop PE?

Robust adaptive MPC algorithm with PE

Let $D(x, Kx) = \sum_{j=1}^{p} \Phi_j [\theta]_j x$, where $\Phi_j = A_j + B_j K \quad j = 1, \ldots, p$

The terminal feedback law $u = Kx$ is *on average* PE if

(a). $\underline{\sigma}\big([\mathrm{vec}(\Phi_1) \ \cdots \ \mathrm{vec}(\Phi_p)]\big) = \sigma_K > 0$

(b). $\mathbb{E}\{ww^\top\} \succeq \epsilon_w I$

Here (a) $\Rightarrow$ $\big\|[\mathrm{vec}(\Phi_1) \ \cdots \ \mathrm{vec}(\Phi_p)]\theta\big\| \geq \sigma_K \|\theta\|$

(b) $\Rightarrow$ $\mathbb{E}\{xx^\top\} \succeq \epsilon_w I$

so that

$$\theta^\top \sum_{i=0}^{N_p - 1} \mathbb{E}\{D(x_i, Kx_i)^\top D(x_i, Kx_i)\}\theta \geq \epsilon_w \sigma_K^2 \|\theta\|^2$$

$$\Longrightarrow \quad \sum_{i=\kappa}^{\kappa + N_p - 1} \mathbb{E}\{D_{i|k}^\top D_{i|k}\} \succeq \epsilon_w \sigma_K^2 \qquad \forall \kappa \geq N$$

Robust adaptive MPC algorithm with PE

Impose PE conditions on predictions in a chain of windows:

$$\text{(PE-LMI):} \qquad \sum_{i=\kappa}^{\kappa+N_p-1} \left(\check{D}_{i|k}^\top \bar{D}_{i|k} + \bar{D}_{i|k}^\top \check{D}_{i|k} + \bar{D}_{i|k}^\top \bar{D}_{i|k}\right) \succeq \hat{\beta}_{\kappa|k} I$$

for $\kappa = -N_p + 1, \dots, 0, \dots, N$

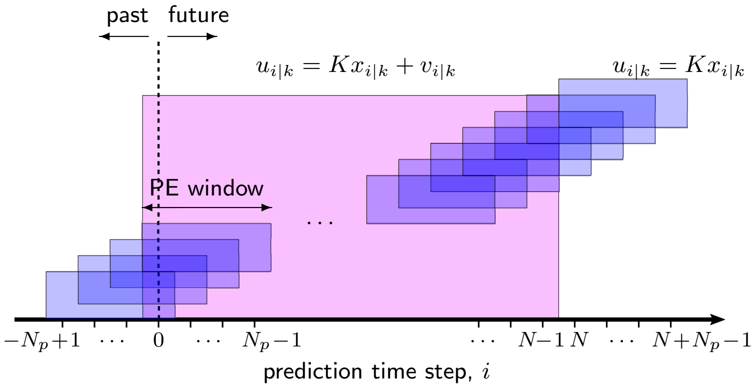

Robust adaptive MPC algorithm with PE

Offline: Choose Θ_0, $\mathcal{X}$, N_p, K, and compute P

Online, at each time $k = 1, 2, \ldots$:

1. Given x_k, update Θ_k, $\hat{\theta}_k$ and compute $\bar{x}_{i|k}, \bar{u}_{i|k}$, $i = 0, \ldots, N + N_p - 1$
2. Compute $\hat{\beta}_{\kappa|k} := \min_{x_\kappa \in \mathcal{X}_{\kappa|k-1}} \max_{\hat{\beta}} \hat{\beta}$ s.t. (PE-LMI) and $v_{i|k} = v^*_{i+1|k-1}\ \forall i$ for $\kappa = -N_p + 1, \ldots, 0, \ldots, N$
3. Compute the solution $(\mathbf{v}_k^*, \boldsymbol{\alpha}_k^*, \boldsymbol{\Lambda}_k^*, \boldsymbol{\beta}_k^*)$ of the semidefinite program

   $$\min_{\mathbf{v}_k, \boldsymbol{\alpha}_k, \boldsymbol{\Lambda}_k, \boldsymbol{\beta}_k} J(x_k, \hat{\theta}_k, \mathbf{v}_k)$$

   subject to $(\mathbf{v}_k, \boldsymbol{\alpha}_k, \boldsymbol{\Lambda}_k) \in \mathcal{F}(x_k, \Theta_k)$ and (PE-LMI)
4. Apply the control law $u_k^* = K x_k + v^*_{0|k}$

Robust adaptive MPC algorithm with PE

Closed loop PE condition

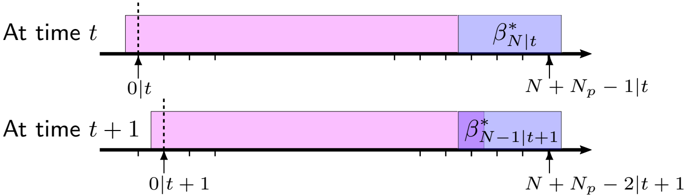

Robust adaptive MPC algorithm with PE

Closed loop PE condition

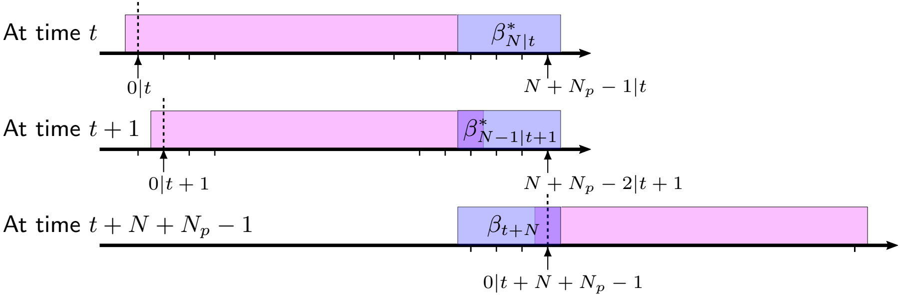

$$\beta_{t+N} = \beta^*_{-N_p+1|t+N+N_p-1} \geq \cdots \geq \beta^*_{N-1|t+1} \geq \beta^*_{N|t}$$

$$\Downarrow$$

$$\mathbb{E}\{\beta_{t+N}\} \geq \cdots \geq \mathbb{E}\{\beta^*_{N|t}\} \geq \epsilon_w \sigma_K^2$$

Robust adaptive MPC algorithm with PE

Closed loop properties:

If $\theta^\star \in \Theta_0$ and $\mathcal{F}(x_0, \Theta_0) \neq \emptyset$, then for all $k > 0$:

1. $\theta^\star \in \Theta_k$
2. $\mathcal{F}(x_k, \Theta_k) \neq \emptyset$
3. $Fx_k + Gu_k \leq \mathbf{1}$

4. The system $x_{k+1} = A(\theta^\star)x_k + B(\theta^\star)u_k^* + w_k$ is ISS
5. $\displaystyle\sum_{i=0}^{N_p-1} \mathbb{E}\big\{D(x_{k+i}, u_{k+i})^\top D(x_{k+i}, u_{k+i})\big\} \succeq \epsilon_w \sigma_K^2 I$

Robust adaptive MPC algorithm with PE

Example with $N = 25$, $N_p = 3$ [Marafioti, Bitmead, Hovd, 2014]

Mean and range of β_t for 30 disturbance sequences

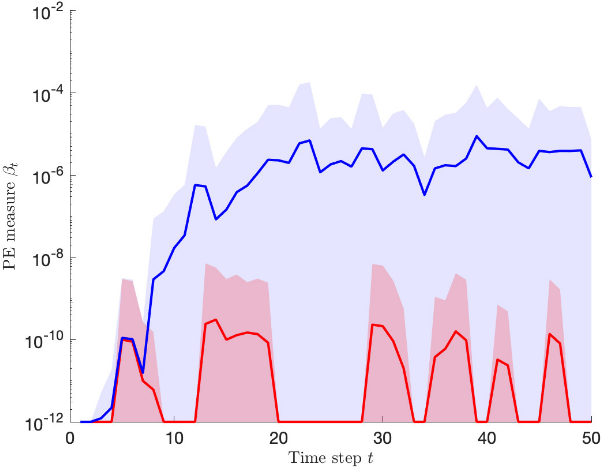

blue: with PE constraints
red: without PE constraints

Robust adaptive MPC algorithm with PE

Example with $N = 25$, $N_p = 3$ [Marafioti, Bitmead, Hovd, 2014]

Mean and range of β_t for 30 disturbance sequences

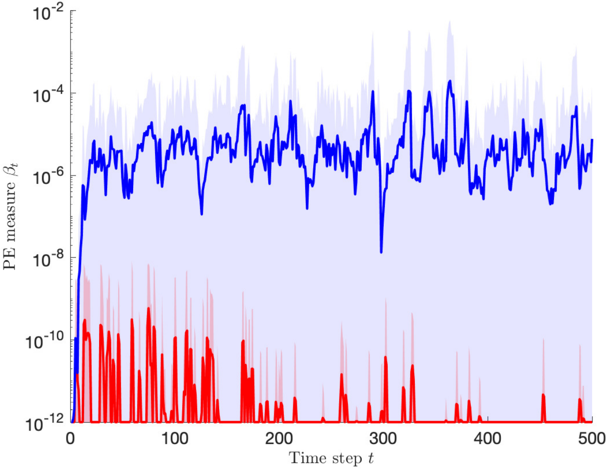

blue: with PE constraints

red: without PE constraints

Robust adaptive MPC algorithm with PE

Convergence of parameter set estimate, $\mathrm{vol}(\Theta_t)$

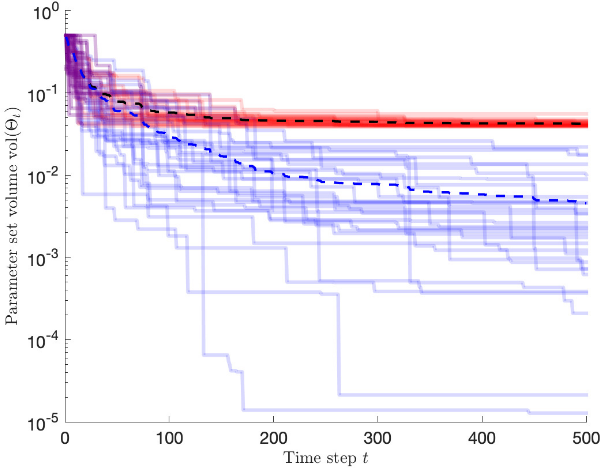

blue: with PE constraints
red: without PE constraints

Robust adaptive MPC algorithm with PE

Convergence and computation for $N = 25$, $N_p = 3$

	Volume %			Mean β_t	CPU time	
	Θ_{10}	Θ_{100}	Θ_{500}		step 2	step 3
with PE	25.4	2.67	0.26	4.9×10^{-5}	0.958	0.073
without PE	25.3	5.77	4.22	9.0×10^{-10}	–	0.052

Differentiable MPC

- MPC law: $u_N(x_k, \hat{\theta}_k, \Theta_k)$ is the solution of a multiparametric programming problem
- Differentiable MPC uses the gradient $\nabla_{\hat{\theta}} u_N(\cdot)$ to train a neural network (NN) with weights $\hat{\theta}_k$ via back-propagation
- Update $\hat{\theta}_k$ with MPC optimization embedded in a NN layer; retain parameter set estimate Θ_k for safe constraint handling

Differentiable MPC: learning model parameters

Linearly parameterised system model:

$$x_{k+1} = f(x_k, u_k, \theta^*) + w_k \qquad \begin{cases} D_k = D(x_k, u_k) \\ d_k = d(x_k, u_k) \end{cases}$$
$$f(x_k, u_k, \theta) = D_k \theta + d_k$$

Parameter set estimate:

$$\Theta_{k+1} \supseteq \Theta_k \cap \Delta_{k+1}$$
$$\Delta_{k+1} = \{\theta : x_{k+1} - D_k\theta - d_k \in \mathcal{W}\}$$

Imitation learning problem: identify θ^* by observing an expert controller

Train $\hat{\theta}_k$ to minimize a loss function

$$\frac{1}{T} \sum_{t=k-T+1}^{k} \left(\|\mathbf{u}_t - \mathbf{u}_N(x_t, \theta_k, \Theta_k)\|^2 + \sigma \|\hat{w}_t\|^2 \right)$$

where

$\mathbf{u}_t = \{u_t, \ldots, u_{t+N-1}\}$ = observed expert control sequence

$\mathbf{u}_N(x_t, \hat{\theta}_k, \Theta_k) = \{u_{0|t}, \ldots, u_{N-1|t}\}$ = MPC law for an initial state x_t

$\hat{w}_t = x_{t+1} - f(x_k, u_k, \hat{\theta}_k)$ = 1-step ahead error

Differentiable MPC: learning model parameters

Linearly parameterised system model:

$$x_{k+1} = f(x_k, u_k, \theta^*) + w_k \qquad \begin{cases} D_k = D(x_k, u_k) \\ d_k = d(x_k, u_k) \end{cases}$$

$$f(x_k, u_k, \theta) = D_k \theta + d_k$$

Parameter set estimate:

$$\Theta_{k+1} \supseteq \Theta_k \cap \Delta_{k+1}$$

$$\Delta_{k+1} = \{\theta : x_{k+1} - D_k\theta - d_k \in \mathcal{W}\}$$

Imitation learning problem: identify θ^* by observing an expert controller

Train $\hat{\theta}_k$ to minimize a loss function

$$\frac{1}{T} \sum_{t=k-T+1}^{k} \Big(\|\mathbf{u}_t - \mathbf{u}_N(x_t, \theta_k, \Theta_k)\|^2 + \sigma \|\hat{w}_t\|^2 \Big)$$

where

$\mathbf{u}_t = \{u_t, \ldots, u_{t+N-1}\}$ = observed expert control sequence

$\mathbf{u}_N(x_t, \hat{\theta}_k, \Theta_k) = \{u_{0|t}, \ldots, u_{N-1|t}\}$ = MPC law for an initial state x_t

$\hat{w}_t = x_{t+1} - f(x_k, u_k, \hat{\theta}_k)$ = 1-step ahead error

Differentiable MPC: learning model parameters

Problem: Regulate $(y, \dot{y})$ subject to bounds on y
Prior assumptions: 2nd order LTI model

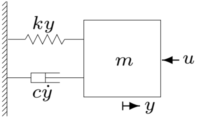

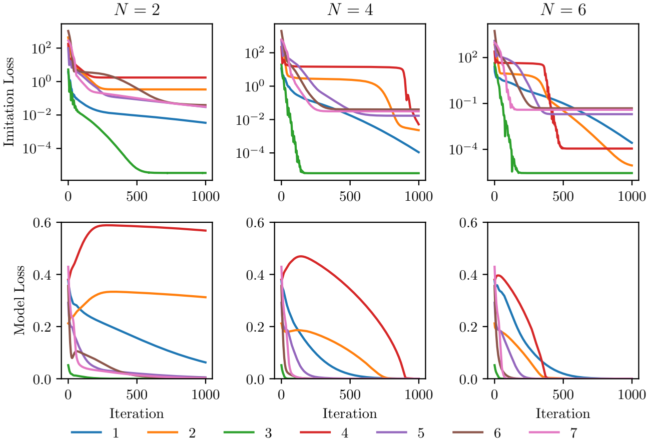

Training results for varying damping c & horizon N

Differentiable MPC: learning model parameters

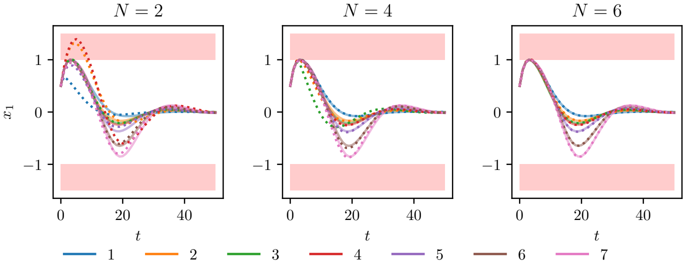

Closed-loop responses: expert (solid lines), learned controllers (dots)

Implementation issues:

- ⋆ nonconvexity and non-uniqueness of optimal parameters
- ⋆ gradient information can vanish or explode across a prediction horizon
- ⋆ expert controller may not be persistently exciting

Differentiable MPC: learning the MPC performance index

Platoon problem:

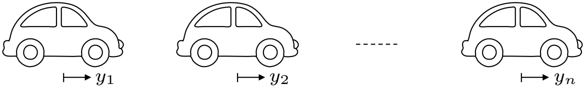

regulate $y_1, \dots y_n$ a so that $\dot{y}_{i+1} - \dot{y}_i \to 0$ subject to $y_{i+1} - y_i \geq \underline{y}$

$$a \leq \ddot{y}_i \leq b$$

Prior assumptions:

- System model ($\ddot{y}_i = u_i$) is known
- $\underline{y}$, a, b are known

Unknown MPC cost to be learnt from observations of an expert controller

Differentiable MPC: learning the MPC performance index

- $n = 10$ vehicles $\implies$ $x \in \mathbb{R}^{18}$, $u \in \mathbb{R}^{10}$
- Cost weights Q, R initialized as random diagonal matrices
- 500 training iterations:

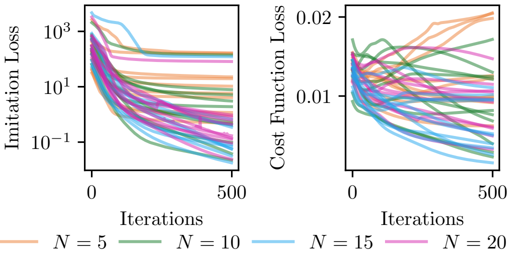

Differentiable MPC: learning the MPC performance index

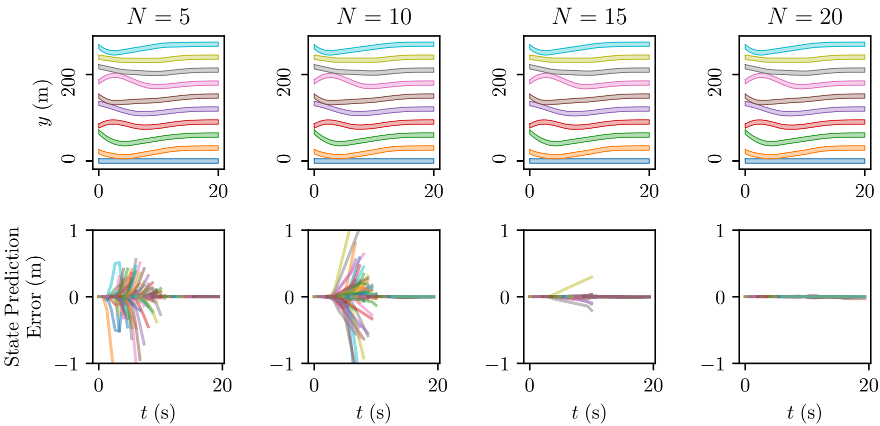

Performance of learnt MPC with horizons $N = 5, 10, 15, 20$:

- constraints $y_{i+1} - y_i \geq \underline{y} = 30$ m are satisfied
- approximately constrained LQ-optimal for $N = 20$

Conclusions

- Adaptive robust MPC is computationally tractable
- Set-membership parameter estimation and robust tube MPC
- Closed loop stability (ISS) and parameter convergence (PE)

Future work

- Can we relax the assumption of bounded disturbances?
- How to combine PE conditions and RNN model adaptation?

References:

– M. Lorenzen, M. Cannon, F. Allgöwer, Robust MPC with recursive model update. *Automatica*, 2019

– S. East, M. Gallieri, J. Masci, J. Koutnik, M. Cannon, Infinite-horizon differentiable Model Predictive Control. *ICLR*, 2020

– X. Lu, M. Cannon, D. Koksal-Rivet, Robust adaptive model predictive control: performance and parameter estimation. *Int. J. Robust & Nonlinear Control*, 2020

Thanks to: Xiaonan Lu, Sebastian East, Matthias Lorenzen

Questions?