MODEL PREDICTIVE CONTROL FROM BASICS TO LEARNING-BASED DESIGN

Alberto Bemporad

imt.lu/ab







1st ELO-X Seasonal School - March 22, 2022

CONTENTS OF MY LECTURE

- Model predictive control (MPC): basic concepts
- Linear MPC and extensions to nonlinear MPC
- Embedded quadratic optimization
- Learning-based nonlinear MPC (feedforward and recurrent neural networks)
- Learning-based hybrid MPC (piecewise affine models)
- Active preference learning for MPC calibration

MODEL PREDICTIVE CONTROL: BASIC CONCEPTS

(extended slide set: http://cse.lab.imtlucca.it/~bemporad/mpc_course.html)

MODEL PREDICTIVE CONTROL (MPC)



simplified likely Use a dynamical model of the process to predict its future evolution and choose the "best" control action a good

MODEL PREDICTIVE CONTROL

• MPC problem: find the best control sequence over a future horizon of N steps

past

future

t+k

t+1 t+1+k

r(t) predicted outputs

manipulated inputs



numerical optimization problem

- **1** estimate current state x(t)
- **2** optimize wrt $\{u_0, \ldots, u_{N-1}\}$
- **3** only apply optimal u_0 as input u(t)

Repeat at all time steps t

t+N+1

t+N

DAILY-LIFE EXAMPLES OF MPC



• MPC is like playing chess !

• You use MPC too when you drive !



©2022 A. Bemporad - MPC: from basics to learning-based design

MPC IN INDUSTRY

• Conceived in the 60's (Rafal, Stevens, 1968) (Propoi, 1963)



- Used in the process industries since the 80's (Qin, Badgewell, 2003)
- Nowadays spreading to the automotive industry and other sectors
- MPC by General Motors and ODYS in high-volume production since 2018

(Bemporad, Bernardini, Long, Verdejo, 2018)



First known mass production of MPC in the automotive industry



www.odys.it

MPC IN INDUSTRY

Table 2

The percentage of survey respondents indicating whether a control technology had demonstrated ("Current Impact") or was likely to demonstrate over the next five years ("Future Impact") high impact in practice.

	Current Impact	Future Impact	
Control Technology	%High	%High	
PID control	91%	78%	
System Identification	65%	72%	
Estimation and filtering	64%	63%	
Model-predictive control	62%	85%	
Process data analytics	51%	70%	
Fault detection and identification	48%	78%	
Decentralized and/or coordinated control	29%	54%	
Robust control	26%	42%	
Intelligent control	24%	59%	
Discrete-event systems	24%	39%	
Nonlinear control	21%	42%	
Adaptive control	18%	44%	
Repetitive control	12%	17%	
Hybrid dynamical systems	11%	33%	
Other advanced control technology	11%	25%	
Game theory	5%	17%	

(Samad et al., 2020)

"As can be observed, MPC is clearly considered more impactful, and likely to be more impactful, vis-à-vis other control technologies, especially those that can be considered the "crown jewels" of control theory - robust control, adaptive control, and nonlinear control."

WORD TRENDS



(source: https://books.google.com/ngrams)

MODEL PREDICTIVE CONTROL - THE BASICS

LINEAR MPC

• Linear prediction model:
$$\begin{cases} x_{k+1} = Ax_k + Bu_k & x \in \mathbb{R}^n \\ y_k = Cx_k & u \in \mathbb{R}^m \\ y \in \mathbb{R}^p \end{cases}$$

• Constrained optimal control problem (quadratic performance index):

$$\min_{z} \quad x'_{N} P x_{N} + \sum_{k=0}^{N-1} x'_{k} Q x_{k} + u'_{k} R u_{k}$$

s.t.
$$u_{\min} \leq u_{k} \leq u_{\max}, \ k = 0, \dots, N-1$$
$$y_{\min} \leq y_{k} \leq y_{\max}, \ k = 1, \dots, N$$

$$\begin{array}{rcl} R & = & R' \succ 0 \\ Q & = & Q' \succeq 0 \\ P & = & P' \succeq 0 \end{array} & z = \left[\begin{array}{c} u_0 \\ u_1 \\ \vdots \\ u_{N-1} \end{array} \right]$$

©2022 A. Bemporad - MPC: from basics to learning-based design

LINEAR MPC

• Optimization problem (condensed form): $x_k = A^k x_0 + \sum_{i=1}^{k} A^i B u_{k-1-i}$

$$V(x_0) = \frac{1}{2}x'_0Yx_0 + \min_{z} \quad \frac{1}{2}z'Hz + x'_0F'z \quad \text{(quadratic objective)}$$

s.t. $Gz \le W + Sx_0 \quad \text{(linear constraints)}$

convex Quadratic Program (QP)

•
$$z = \begin{bmatrix} u_0 \\ u_1 \\ \vdots \\ u_{N-1} \end{bmatrix} \in \mathbb{R}^{Nm}$$
 is the optimization vector



- QP matrices depend on chosen weights, model, and constraints
- Alternative: keep also x_1, \ldots, x_N as optimization variables and the equality constraints $x_{k+1} = Ax_k + Bu_k$ (non-condensed form, which is sparse)

LINEAR MPC ALGORITHM

@ each sampling step t:

• Estimate the current state x(t)

• Get the solution
$$z^* = \begin{bmatrix} u_0^* \\ u_1^* \\ \vdots \\ u_{N-1}^* \end{bmatrix}$$
 of the QP
$$\begin{cases} t & t+k & \text{freedback} \\ \min & \frac{1}{2}z'Hz + x'(t)F'z \\ \text{s.t.} & Gz \leq W + S & x(t) \\ \text{freedback} & \text{freedback} \end{cases}$$

past

future

r(t)

predicted outputs

manipulated inputs

- Apply only $u(t) = u_0^*,$ discarding the remaining optimal inputs u_1^*, \ldots, u_{N-1}^*

• Unconstrained MPC: Hz + Fx(t) = 0 $u(t) = -[I \ 0 \ \dots \ 0]H^{-1}Fx(t)$ linear state feedback!

BASIC CONVERGENCE PROPERTIES

(Keerthi, Gilbert, 1988) (Bemporad, Chisci, Mosca, 1994)

• Theorem: Let the MPC law be based on

V

*
$$(x(t)) = \min$$

$$\sum_{k=0}^{N-1} x'_k Q x_k + u'_k R u_k$$

s.t.
$$x_{k+1} = A x_k + B u_k$$
$$u_{\min} \le u_k \le u_{\max}$$
$$y_{\min} \le C x_k \le y_{\max}$$
$$x_N = 0 \quad \leftarrow \text{"terminal constraint"}$$

with $R, Q \succ 0, u_{\min} < 0 < u_{\max}, y_{\min} < 0 < y_{\max}$. If the optimization problem is feasible at time t = 0 then

$$\lim_{t \to \infty} x(t) = 0, \quad \lim_{t \to \infty} u(t) = 0$$

and the constraints are satisfied at all time $t \ge 0$, for all $R, Q \succ 0$.

Many more convergence and stability results exist (Mayne, 2014)

LINEAR MPC - TRACKING

- Objective: make the output y(t) track a reference signal r(t)
- Let us parameterize the problem using the input increments

$$\Delta u(t) = u(t) - u(t-1)$$

- As $u(t)=u(t-1)+\Delta u(t)$ we need to extend the system with a new state $x_u(t)=u(t-1)$

$$\begin{cases} x(t+1) = Ax(t) + Bu(t-1) + B\Delta u(t) \\ x_u(t+1) = x_u(t) + \Delta u(t) \end{cases}$$

$$\begin{cases} \begin{bmatrix} x(t+1) \\ x_u(t+1) \end{bmatrix} = \begin{bmatrix} A & B \\ 0 & I \end{bmatrix} \begin{bmatrix} x(t) \\ x_u(t) \end{bmatrix} + \begin{bmatrix} B \\ I \end{bmatrix} \Delta u(t) \\ y(t) = \begin{bmatrix} C & 0 \end{bmatrix} \begin{bmatrix} x(t) \\ x_u(t) \end{bmatrix} \end{cases}$$

- Again a linear system with states $x(t), x_u(t)$ and input $\Delta u(t)$

LINEAR MPC - TRACKING

• Optimal control problem (quadratic performance index):

$$\begin{vmatrix} \min_{z} & \sum_{k=0}^{N-1} \|W^{y}(y_{k+1} - r(t))\|_{2}^{2} + \|W^{\Delta u} \Delta u_{k}\|_{2}^{2} \\ & [\Delta u_{k} \triangleq u_{k} - u_{k-1}], u_{-1} = u(t-1) \\ \text{s.t.} & u_{\min} \le u_{k} \le u_{\max}, k = 0, \dots, N-1 \\ & y_{\min} \le y_{k} \le y_{\max}, k = 1, \dots, N \\ & \Delta u_{\min} \le \Delta u_{k} \le \Delta u_{\max}, k = 0, \dots, N-1 \end{vmatrix}$$

$$z = \begin{bmatrix} \Delta u_{0} \\ \Delta u_{1} \\ \vdots \\ \Delta u_{N-1} \end{bmatrix} \text{ or } z = \begin{bmatrix} u_{0} \\ u_{1} \\ \vdots \\ u_{N-1} \end{bmatrix}$$

weight $W^{(\cdot)}$ = diagonal matrix

$$\min_{z} \quad J(z, x(t)) = \frac{1}{2}z'Hz + [x'(t) r'(t) u'(t-1)]F'z$$

s.t. $Gz \le W + S \begin{bmatrix} x(t) \\ r(t) \\ u(t-1) \end{bmatrix}$

convex Quadratic Program

- Add the extra penalty $\|W^u(u_k u_{ref}(t))\|_2^2$ to track input references
- Constraints may depend on r(t), such as $e_{\min} \leq y_k r(t) \leq e_{\max}$

ANTICIPATIVE ACTION (A.K.A. "PREVIEW")

$$\min_{\Delta U} \sum_{k=0}^{N-1} \|W^{y}(y_{k+1} - r(t+k))\|_{2}^{2} + \|W^{\Delta u} \Delta u(k)\|_{2}^{2}$$

• Reference not known in advance (causal):



• Future refs (partially) known in advance (anticipative action):



• Same for previewing measured disturbances $x_{k+1} = Ax_k + Bu_k + B_v v(t+k)$

OUTPUT INTEGRATORS AND OFFSET-FREE TRACKING

• Add constant unknown disturbances on measured outputs:

$$\begin{cases} x_{k+1} = Ax_k + Bu_k \\ d_{k+1} = d_k \\ y_k = Cx_k + d_k \end{cases}$$

- Use the extended model to design a state observer (e.g., Kalman filter) that estimates both the state $\hat{x}(t)$ and disturbance $\hat{d}(t)$ from y(t)
- Why we get offset-free tracking in steady-state (intuitively):
 - the observer makes $C\hat{x}(t) + \hat{d}(t) \rightarrow y(t)$ (estimation error)
 - the MPC controller makes $C\hat{x}(t) + \hat{d}(t) \rightarrow r(t)$
 - the combination of the two makes $y(t) \rightarrow r(t)$

(predicted tracking error) (actual tracking error)

- In steady state, the term $\hat{d}(t)$ compensates for model mismatch
- See more on survey paper (Pannocchia, Gabiccini, Artoni, 2015)

EMBEDDED QUADRATIC OPTIMIZATION FOR MPC

EMBEDDED LINEAR MPC AND QUADRATIC PROGRAMMING

MPC based on linear models requires solving a Quadratic Program (QP)

$$\begin{array}{ll} \min_{z} & \frac{1}{2}z'Qz + x'(t)F'z + \frac{1}{2}x'(t)Yx(t)\\ \text{s.t.} & Gz \leq W + Sx(t) \end{array} \qquad z = \begin{bmatrix} u_{0}\\ u_{1}\\ \vdots\\ u_{N-1} \end{bmatrix} \end{array}$$

$$\begin{array}{l} \text{MINIMEZING A CONVEX FUNCTION SUBJECT TO LINEAR INEQUALITIES}\\ \text{By E. M. I. BALE}\\ \text{Advisition of a convex function of variable subject to linear inequalities is discrimination of the last two sections a form of linear programming the strenged to yield finite algorithms for minimizing either a donce quadratic function of the sum of the last two sections a form of linear programming to of a convex function. \\ \end{array}$$

$$\begin{array}{l} \text{Beale, 1955} \end{array}$$

$$\begin{array}{l} \text{Minimum Section of good QP algorithms is available todays } \end{array}$$

Not all QP algorithms are suitable for industrial embedded control

 $Gz \le W + Sx(t)$

s.t.

 $\frac{1}{z}z'Qz + x(t)'F'z = \text{constant}$

MPC IN A PRODUCTION ENVIRONMENT

Key requirements for deploying MPC in production:

- speed (throughput)
 - worst-case execution time less than sampling interval
 - also fast on average (to free the processor to execute other tasks)
- 2. limited memory and CPU power (e.g., 150 MHz / 50 kB)
- 3. numerical robustness (single precision arithmetic)
- 4. certification of worst-case execution time
- code simple enough to be validated/verified/certified (library-free C code, easy to check by production engineers)









vſil



©2022 A. Bemporad - MPC: from basics to learning-based design

EMBEDDED SOLVERS IN INDUSTRIAL PRODUCTION

- Multivariable MPC controller
- Sampling frequency = 40 Hz (= 1 QP solved every 25 ms)
- Vehicle operating \approx 1 hr/day for \approx 360 days/year on average
- Controller running on 10 million vehicles



DUAL GRADIENT PROJECTION FOR QP

(Goldstein, 1964) (Levitin, Poljak, 1965) (Combettes, Waijs, 2005)

• Consider the strictly convex QP and its dual

$$\min_{\substack{1 \\ \text{s.t.}}} \frac{1}{2}z'Qz + x'F'z \\ \text{s.t.} \quad Gz \le W + Sx$$

$$\min_{\substack{1 \\ \text{s.t.}}} \frac{1}{2}y'Hy + (Dx + W)'y \\ \text{s.t.} \quad y \ge 0$$

with $H = GQ^{-1}G'$, $D = S + GQ^{-1}F$. Take $L \ge \lambda_{\max}(H)$

• Apply proximal gradient method to dual QP:

$$y^{k+1} = \max\{y^k - \frac{1}{L}(Hy^k + Dx + W), 0\}$$
 $y_0 = 0$

• The primal solution is related to the dual solution by

$$z^k = -Q^{-1}(Fx + G'y^k)$$

- Convergence is slow: the initial error $f(z^0) - f(z^*)$ reduces as 1/k

FAST GRADIENT PROJECTION FOR (DUAL) QP

(Nesterov, 1983) (Beck, Teboulle, 2008) (Patrinos, Bemporad, 2014)

• The fast gradient method is applied to solve the dual QP problem

$ \min_{z} s.t. $		$\frac{1}{2}z'Qz + x'F'z$ $Gz \le W + Sx$	w^k z^k	=	$y^k + \beta_k (y^k - y^{k-1})$ $-Kw^k - Jx$	<pre>while k<maxiter beta=max(k<1)/(k+2),0); w=y+beta*(y-y0); z=-(iM6*w+iMc); s=GL*z-bL; y0=y;</maxiter </pre>
K J L	= = >	$Q^{-1}G'$ $Q^{-1}F$ $\lambda_{\max}(GQ^{-1}G')$	s^k y^{k+1}	=	$\frac{1}{L}Gz^k - \frac{1}{L}(W + Sx)$ $\max\left\{w^k + s^k, 0\right\}$	<pre>% Termination if all(s<=eps(L) gapL===w*s; if gapL<=epsVL return end</pre>
β_k		$\max\{\frac{k-1}{k+2}, 0\}$				y=w+s; k=k+1; end

• Very simple to code

1

FAST GRADIENT PROJECTION FOR (DUAL) QP

• Termination criteria: when the following two conditions are met

 $\begin{array}{rcl} s_i^k & \leq & \frac{1}{L}\epsilon_G, \, i=1,\ldots,m \\ -(w^k)'s^k & \leq & \frac{1}{L}\epsilon_f \end{array} \qquad \qquad \mbox{primal feasibility} \\ \end{array}$

the solution $z^k = -Kw^k - Jx$ satisfies $G_i z^k - W_i - S_i x \leq \epsilon_G$ and, if $w^k \geq 0$,



- Can be useful to warm-start active-set methods (Bemporad, Paggi, 2015)
- Extended to mixed-integer quadratic programming (MIQP) (Naik, Bemporad, 2017)



(Gabay, Mercier, 1976) (Glowinski, Marrocco, 1975) (Douglas, Rachford, 1956) (Boyd et al., 2010)

• Alternating Directions Method of Multipliers for QP

$$\begin{aligned} & \overset{k+1}{=} & -(Q + \rho A'A)^{-1}(\rho A'(v^k - s^k) + c) \\ & \overset{k+1}{=} & \min\{\max\{Az^{k+1} + v^k, \ell\}, u\} \\ & \overset{k+1}{=} & v^k + Az^{k+1} - s^{k+1} \end{aligned}$$

in
$$\frac{1}{2}z'Qz + c'z$$

.t. $\ell \le Az \le u$

m

s

while k<maxiter
 k=k+1;
 z=-iM*(c+A'*(rho*(v-s)));
 Az=A*z;
 s=max(min(Az+v,u),ell);
 v=v+Az-s;
end</pre>

(7 lines EML code) (\approx 40 lines of C code)

 ρv = dual vector

 z^{k}

 s^{k}

 v^k

- Matrix $(Q + \rho A'A)$ must be nonsingular
- The factorization of matrix $(Q + \rho A'A)$ can be done at start and cached
- Very simple to code. Sensitive to matrix scaling (as gradient projection)
- Used in many applications (control, signal processing, machine learning)

REGULARIZED ADMM FOR QUADRATIC PROGRAMMING

(Stellato, Banjac, Goulart, Bemporad, Boyd, 2020)

• Robust "regularized" ADMM iterations:

$$\begin{aligned} z^{k+1} &= -(Q + \rho A^T A + \epsilon I)^{-1} (c - \epsilon z_k + \rho A^T (v^k - z^k)) \\ s^{k+1} &= \min\{\max\{Az^{k+1} + v^k, \ell\}, u\} \\ v^{k+1} &= v^k + Az^{k+1} - s^{k+1} \end{aligned}$$

- Works for any $Q \succeq 0, A$, and choice of $\epsilon > 0$
- Simple to code, fast, and robust

• Only needs to factorize
$$\begin{bmatrix} Q + \epsilon I & A' \\ A & -\frac{1}{\rho}I \end{bmatrix}$$
 once

 Implemented in free osQP solver (Python interface: ≈ 1,700,000 downloads) http://osqp.org

• Extended to solve mixed-integer quadratic programming problems

(Stellato, Naik, Bemporad, Goulart, Boyd, 2018) ©2022 A. Bemporad - MPC: from basics to learning-based design

ODYS QP SOLVER

• General purpose QP solver designed for industrial production

$$\min_{z} \qquad \frac{1}{2}z'Qz + c'z$$

s.t.
$$b_{\ell} \le Az \le b_{u}$$
$$\ell \le z \le u$$
$$Ez = f$$



odys.it/qp

- Implements a proprietary state-of-the-art method for QP
- Completely written in ANSI-C and MISRA-C 2012 compliant
- Fast, robust (also in single precision), low-memory requirements
- optimized version for MPC available (\approx 50% faster)
- Licensed to several automotive OEMs and Tier-1 suppliers
- Certifiable execution time

PRIMAL-DUAL INTERIOR-POINT METHOD FOR QP

(Nocedal, Wright, 2006) (Gondzio, Terlaki, 1994)

• The Karush-Kuhn-Tucker(KKT) optimality conditions for the convex QP

are

- $\begin{array}{rcl} r_Q &=& Qx + c + E'y + A'z &=& 0 & x = \text{primal vars} \\ r_E &=& Ex f &=& 0 & y = \text{dual vars (eq. constr.)} \\ r_A &=& Ax + s b &=& 0 & s = \text{slacks (ineq. constr.)} \\ r_S &=& [z_1s_1 \dots z_ms_m]' &=& 0 & z = \text{dual vars (ineq. constr.)} \\ z,s &\geq& 0 \end{array}$
- In a nutshell, **interior-point** methods use Newton's method with line search to solve the above nonlinear system of equations
- The complementary slackness constraint is replaced by $z_i s_i = \mu$ and $\mu \rightarrow 0$

PRIMAL-DUAL INTERIOR-POINT METHOD FOR QP

(Nocedal, Wright, 2006) (Gondzio, Terlaki, 1994)

• Each interior-point iteration requires solving a linear system of the form

- In MPC the structure $x_{k+1} = Ax_k + Bu_k$ can be heavily exploited to factorize/solve the linear systems efficiently (Rao, Wright, Rawlings, 1998) (Wright, 2018)
- IP provides good solutions within 10-15 IP iterations (usually ...).
- Linear systems tends to become ill-conditioned at convergence
- IP usually faster for sparse and large QPs (say >500 vars & constraints)

MPC WITHOUT ON-LINE QP





• Can we implement constrained linear MPC without an on-line QP solver?

EXPLICIT MODEL PREDICTIVE CONTROL

Continuous & piecewise affine solution of strictly convex multiparametric QP

$$z^*(x) = \arg \min_z \quad \frac{1}{2}z'Qz + x'F'z$$

s.t. $Gz \le W + Sx$

(Bemporad, Morari, Dua, Pistikopoulos, 2002)



Corollary: linear MPC is continuous & piecewise affine !

$$z^* = \begin{bmatrix} \mathbf{u}_0 \\ u_1 \\ \vdots \\ u_{N-1}^* \end{bmatrix} \qquad \qquad u_0^*(x) = \begin{cases} F_1 x + g_1 & \text{if} \quad H_1 x \le K_1 \\ \vdots & \vdots \\ F_M x + g_M & \text{if} \quad H_M x \le K_M \end{cases}$$

 New mpQP solver based on NNLS available (Bemporad, 2015) and included in MPC Toolbox since R2014b (Bemporad, Morari, Ricker, 1998-today)

Is explicit MPC better than on-line QP (=implicit MPC)?

COMPLEXITY CERTIFICATION FOR ACTIVE-SET QP SOLVERS

• **Result**: The **number of iterations** to solve the QP via a dual active-set method is a **piecewise constant function** of the parameter *x*



(Cimini, Bemporad, 2017)

We can **exactly** quantify how many iterations (flops) the QP solver takes in the worst-case !

	inverted pendulum	DC motor	nonlinear demo	AFTI F16
Explicit MPC				
max flops	3382	1689	9184	16434
max memory (kB)	55	30	297	430
Implicit MPC				
max flops	3809	2082	7747	7807
sqrt	27	9	37	33
max memory (kB)	15	13	20	16

• QP certification algorithm currently used in industrial production projects

FROM LINEAR TO NONLINEAR MPC

LINEAR TIME-VARYING MODELS

• Linear Time-Varying (LTV) model

$$\begin{cases} x_{k+1} = A_k(t)x_k + B_k(t)u_k \\ y_k = C_k(t)x_k \end{cases}$$

- At each time t the model can also change over the prediction horizon k
- Possible measured disturbances are embedded in the model
- On-line optimization is still a QP

$$\min_{z} \qquad \frac{1}{2}z'H(t)z + \begin{bmatrix} x(t) \\ r(t) \\ u(t-1) \end{bmatrix}' F(t)'z$$

s.t.
$$G(t)z \le W(t) + S(t) \begin{bmatrix} x(t) \\ r(t) \\ u(t-1) \end{bmatrix}$$

• The QP matrices cannot be constructed offline

LINEARIZING A NONLINEAR MODEL

• LTV models can be obtained by linearizing a nonlinear model

$$\begin{cases} \frac{dx_c(t)}{dt} &= f(x_c(t), u_c(t)) \\ y_c(t) &= g(x_c(t)) \end{cases}$$

• At time t, consider the **nominal trajectory**

$$U = \{ \bar{u}_c(t), \bar{u}_c(t+T_s), \dots, \bar{u}_c(t+(N-1)T_s) \}$$

For example U = shifted previous sequence optimized by MPC @t-1

• Integrate the model from $\bar{x}_c(t)$ and get nominal state/output trajectories

$$X = \{ \bar{x}_c(t), \bar{x}_c(t+T_s), \dots, \bar{x}_c(t+(N-1)T_s) \}$$

$$Y = \{ \bar{y}_c(t), \bar{y}_c(t+T_s), \dots, \bar{y}_c(t+(N-1)T_s) \}$$

For example $\bar{x}_c(t) = \text{current state}$
LINEARIZING A NONLINEAR MODEL

• Linearize the nonlinear model around the nominal states and inputs:

• Define $x \triangleq x_c - \bar{x}_c$, $u \triangleq u_c - \bar{u}_c$, $y \triangleq y_c - \bar{y}_c$ and get the linear system

$$\frac{dx}{dt} = A_c x + B_c u \qquad \qquad y = C x$$

- Convert linear model to discrete-time and get matrices (A_k, B_k, C_k)
- Alternative: compute (A_k, B_k, C_k) (a.k.a. sensitivities) during integration

FROM LTV-MPC TO NONLINEAR MPC

- How to use the LTV-MPC machinery to handle nonlinear MPC?
- Key idea: Solve a sequence of LTV-MPC problems at each time t

For h = 0 to $h_{\max} - 1$ do:

- 1. Simulate from x(t) with inputs U_h and get state trajectory X_h
- 2. Linearize around (X_h, U_h) and discretize in time
- 3. Get $U_{h+1}^* = \mathbf{QP}$ solution of corresponding LTV-MPC problem
- 4. Line search: find optimal step size $\alpha_h \in (0, 1]$;
- 5. Set $U_{h+1} = (1 \alpha_h)U_h + \alpha_h U_{h+1}^*$;

Return solution $U_{h_{\max}}$

• Special case: just solve one iteration with $\alpha = 1$ (a.k.a. Real-Time Iteration)

(Diehl, Bock, Schloder, Findeisen, Nagy, Allgower, 2002) = LTV-MPC

• Example



OUTPUT FEEDBACK - EXTENDED KALMAN FILTER

• For state estimation, an Extended Kalman Filter (EKF) can be used based on the same nonlinear model (with additional noise)

$$x(k+1) = f(x(k), u(k), \xi(k))$$

 $y(k) = g(x(k)) + \zeta(k)$

• measurement update:

• time update:

$$\begin{split} \hat{x}(k+1|k) &= f(\hat{x}(k|k), u(k)) \\ A(k) &= \frac{\partial f}{\partial x}(\hat{x}_{k|k}, u(k), E[\xi(k)]), \ G(k) &= \frac{\partial f}{\partial \xi}(\hat{x}_{k|k}, u(k), E[\xi(k)]) \\ P(k+1|k) &= A(k)P(k|k)A(k)' + G(k)Q(k)G(k)' \end{split}$$

ODYS EMBEDDED MPC TOOLSET

• **ODYS Embedded MPC** is a software toolchain for design and deployment of MPC solutions in industrial production



- Support for linear & nonlinear MPC and extended Kalman filtering
- Extremely flexible, all MPC parameters can be changed at runtime
- Integrated with MPC-optimized version of ODYS QP Solver
- Library-free C code, MISRA-C 2012 compliant, supports also single precision
- Currently used worldwide by several automotive OEMs in R&D and production
- MPC Toolbox Plugin to easily import NL-MPC projects from MPC Toolbox
- ODYS Deep Learning supports neural networks as prediction models

odys.it/embedded-mpc

LEARNING-BASED NONLINEAR MPC

• Good mathematical foundations from artificial intelligence, statistics, optimization

- Works very well in practice (despite training is most often a nonconvex optimization problem ...)
- Used in myriads of very diverse application domains

• Availability of excellent open-source software tools also explains success scikit-learn, TensorFlow/Keras, PyTorch, JAX, Flux.jl,... ቅ python julia

CONTROL-ORIENTED NONLINEAR MODELS

• Black-box modeling: purely data-driven. Use training data to fit a prediction model that can explain them



• **Physics-based modeling**: use physical principles to create a prediction model (e.g.: weather forecast, chemical reaction, mechanical laws, ...)



• Gray-box modeling is a mix of the two. It can be quite effective

"All models are wrong, but some are useful."

NONLINEAR SYS-ID BASED ON NEURAL NETWORKS

• Neural networks proposed for nonlinear system identification since the '90s

(Hunt et al., 1992) (Suykens, Vandewalle, De Moor, 1996)

- NNARX models: use a feedforward neural network to approximate the nonlinear difference equation $y_t \approx \mathcal{N}(y_{t-1}, \dots, y_{t-n_a}, u_{t-1}, \dots, u_{t-n_b})$
- Neural state-space models:
 - w/ state data: fit a neural network model $x_{t+1} \approx \mathcal{N}_x(x_t, u_t), \ y_t \approx \mathcal{N}_y(x_t)$
 - I/O data only: set xt = value of an inner layer of the network (Prasad, Bequette, 2003)
- Alternative for MPC: learn entire prediction (Masti, Smarra, D'Innocenzo, Bemporad, 2020)

$$y_{t+k} = h_k(x_t, \boldsymbol{u_t}, \dots, \boldsymbol{u_{t+k-1}}), \, k = 1, \dots, N$$



• **Recurrent neural networks** are more appropriate for accurate open-loop predictions, but more difficult to train (see later ...)

NONLINEAR STATE-SPACE MODELS VIA AUTOENCODERS

• Idea: use autoencoders and artificial neural networks to learn a nonlinear state-space model of desired order from input/output data



ANN with hourglass structure (Hinton, Salakhutdinov, 2006)

• Quasi-LPV structure for MPC: set $(A_{ij}, B_{ij}, C_{ij} = \text{feedforward NNs})$



$$\begin{array}{rcl} x_{k+1} & = & A(x_k, u_k) \left[\begin{smallmatrix} x_k \\ 1 \end{smallmatrix}\right] + B(x_k, u_k) u_k \\ y_k & = & C(x_k, u_k) \left[\begin{smallmatrix} x_k \\ 1 \end{smallmatrix}\right] \end{array}$$

LEARNING NEURAL NETWORK MODELS FOR CONTROL

TRAINING FEEDFORWARD NEURAL NETWORKS

• Feedforward neural network model:

$$y_{k} = f_{y}(x_{k}, \theta) = \begin{cases} v_{1k} = A_{1}x_{k} + b_{1} \\ v_{2k} = A_{2}f_{1}(v_{1k}) + b_{2} \\ \vdots & \vdots \\ v_{Lk} = A_{Ly}f_{L-1}(v_{(L-1)k}) + b_{L} \\ \hat{y}_{k} = f_{L}(v_{Lk}) \end{cases}$$



Examples: x_k = measured state, or $x_k = (y_{k-1}, \dots, y_{k-n_a}, u_{k-1}, \dots, u_{k-n_b})$

• Training problem: given a dataset $\{x_0, y_0, \ldots, x_{N-1}, y_{N-1}\}$ solve

$$\min_{\theta} r(\theta) + \sum_{k=0}^{N-1} \ell(y_k, f(x_k, \theta))$$



• It is a nonconvex, unconstrained, nonlinear programming problem that can be solved by **stochastic gradient descent**, **quasi-Newton** methods, ... and **EKF** !

©2022 A. Bemporad - MPC: from basics to learning-based design

TRAINING FEEDFORWARD NEURAL NETWORKS BY EKF

(Singhal, Wu, 1989) (Puskorius, Feldkamp, 1994)

• Key idea: treat parameter vector *θ* of the feedforward neural network as a constant state

$$\begin{array}{rcl} \theta_{k+1} &=& \theta_k + \eta_k \\ y_k &=& f(x_k, \theta_k) + \zeta_k \end{array}$$

and use EKF to estimate θ_k on line from a streaming dataset $\{x_k, y_k\}$

- Ratio $\operatorname{Var}[\eta_k]/\operatorname{Var}[\zeta_k]$ is related to the learning-rate
- Initial matrix $(P_{0|-1})^{-1}$ is related to quadratic regularization on θ
- Implemented in ODYS Deep Learning library
- Extended to rather arbitrary convex loss functions/regularization terms (Bemporad, 2021 https://arxiv.org/abs/2111.02673)

RECURRENT NEURAL NETWORKS

• Recurrent Neural Network (RNN) model:

 $x_{k+1} = f_x(x_k, u_k, \theta_x) \left| f_x, f_y \neq \text{feedforward neural network} \right. \ y_k = f_y(x_k, \theta_y)$

• Training problem: given a dataset $\{u_0, y_0, \dots, u_{N-1}, y_{N-1}\}$ solve

$$\min_{\substack{\theta_x, \theta_y \\ x_0, x_1, \dots, x_{N-1}}} r(x_0, \theta_x, \theta_y) + \sum_{k=0}^{N-1} \ell(y_k, f_y(x_k, \theta_y))$$

s.t. $x_{k+1} = f_x(x_k, u_k, \theta_x)$

• Main issue: x_k are hidden states, i.e., are unknowns of the problem

TRAINING RNNS ONLINE BY EKF

• Estimate both hidden states x_k and parameters θ_x, θ_y by EKF based on

$$\begin{cases} x_{k+1} &= f_x(x_k, u_k, \theta_{xk}) + \xi_k \\ \begin{bmatrix} \theta_{x(k+1)} \\ \theta_{y(k+1)} \end{bmatrix} &= \begin{bmatrix} \theta_{xk} \\ \theta_{yk} \end{bmatrix} + \eta_k \\ y_k &= f_y(x_k, \theta_{yk}) + \zeta_k \end{cases}$$

- RNN and its hidden state x_k can be estimated on line from a streaming dataset $\{u_k, y_k\}$, and/or offline by processing multiple epochs of a given dataset
- Can handle general smooth strictly convex loss functions/regularization terms
- Can add ℓ_1 -penalty $\lambda \left\| \begin{bmatrix} \theta_x \\ \theta_y \end{bmatrix} \right\|_1$ to sparsify θ_x, θ_y by changing EKF update into

$$\begin{bmatrix} \hat{x}(k|k) \\ \theta_x(k|k) \\ \theta_y(k|k) \end{bmatrix} = \begin{bmatrix} \hat{x}(k|k-1) \\ \theta_x(k|k-1) \\ \theta_y(k|k-1) \end{bmatrix} + M(k)e(k) - \lambda P(k|k-1) \begin{bmatrix} 0 \\ \operatorname{sign}(\theta_x(k|k-1)) \\ \operatorname{sign}(\theta_y(k|k-1)) \end{bmatrix}$$

©2022 A. Bemporad - MPC: from basics to learning-based design

TRAINING RNNS BY EKF - EXAMPLES

- Dataset: 3499 I/O data of magneto-rheological fluid damper (Wang et al., 2009)
- N=2000 data used for training, 1499 for testing the model
- Same data used in NNARX modeling demo of SYS-ID Toolbox for MATLAB
- RNN model: 4 hidden states shallow state-update and output functions 6 neurons each, leaky-ReLU activation
- Compare with gradient descent (AMSGrad)



• Training time measured on MATLAB+CasADi implementation of EKF/AMSGrad

TRAINING RNNS BY EKF - EXAMPLES

• Compare NRMSE¹ wrt NNARX model (SYS-ID TBX):

EKF = **91.97**, AMSGrad = **85.58**, NNARX(6,2) = **88.18** (training) EKF = **90.54**, AMSGrad = **80.95**, NNARX(6,2) = **85.15** (test)

• Repeat training with ℓ_1 -penalty $\lambda \left\| \begin{bmatrix} \theta_x \\ \theta_y \end{bmatrix} \right\|_1$



¹normalized root-mean-square error



TRAINING RNNS BY EKF - EXAMPLES

Dataset: 2000 I/O data of linear system with binary outputs

$$\begin{aligned} x(k+1) &= \begin{bmatrix} \frac{.8}{0} & \frac{.2}{.1} & -.1\\ \frac{.0}{.1} & -.1 & \frac{.7}{.1} \end{bmatrix} x(k) + \begin{bmatrix} -1\\ \frac{.5}{.1} \end{bmatrix} u(k) + \xi(k) \qquad \text{Var}[\xi_i(k)] = \sigma^2 \\ y(k) &= \begin{cases} 1 & \text{if } [-2 & 1.5 & 0.5] x(k) - 2 + \zeta(k) \ge 0\\ 0 & \text{otherwise} \end{cases} \qquad \text{Var}[\zeta(k)] = \sigma^2 \end{aligned}$$

- + N=1000 data used for training, 1000 for testing the model
- Train linear state-space model with 3 states and sigmoidal output function
 σ σ</l

$$f_1^y(y) = 1/(1 + e^{-A_1^y[x'(k) \ u(k)]' - b_1^y})$$

• Training loss: (modified) cross-entropy loss $\ell_{CE\epsilon}(y(k), \hat{y}) = \sum_{i=1}^{n_y} -y_i(k) \log(\epsilon + \hat{y}_i) - (1 - y_i(k)) \log(1 + \epsilon - \hat{y}_i)$

	σ	training	test	
	0.000	99.20	98.90	
	0.001	99.30	98.90	
	0.010	99.20	98.70	
	0.100	96.50	97.00	
	0.200	93.00	93.80	

TRAINING RNNS BY SEQUENTIAL LEAST-SQUARES

(Bemporad, 2021 - http://arxiv.org/abs/2112.15348)

• RNN training problem = **optimal control** problem:

$$\min_{\theta_x, \theta_y, x_0, x_1, \dots, x_{N-1}} \quad r(x_0, \theta_x, \theta_y) + \sum_{k=0}^{N-1} \ell(y_k, \hat{y}_k)$$

s.t.
$$x_{k+1} = f_x(x_k, u_k, \theta_x)$$
$$\hat{y}_k = f_y(x_k, \theta_y)$$

- θ_x, θ_y, x_0 = manipulated variables, \hat{y}_k = output, y_k = reference signal
- $r(x_0, \theta_x, \theta_y)$ = input penalty, $\ell(y_k, \hat{y}_k)$ = output penalty
- N = prediction horizon, control horizon = 1
- Linearized model:

$$\begin{aligned} \Delta x_{k+1} &= (\nabla_x f_x)' \Delta x_k + (\nabla_{\theta_x} f_x)' \Delta \theta_x \\ \Delta y_k &= (\nabla_{x_k} f_y)' \Delta x_k + (\nabla_{\theta_y} f_y)' \Delta \theta_y \end{aligned}$$

Idea: take 2nd-order expansions of the loss *l* and regularization term *r* and use sequential least-squares + line search to minimize wrt x₀, θ_x, θ_y

TRAINING RNNS BY SEQUENTIAL LS AND ADMM

(Bemporad, 2021 - http://arxiv.org/abs/2112.15348)

- 10 FKF FKF AMSGrad AMSGrad MSE Sea. LS MSE Seq. LS 10 10 50 100 150 200 250 0 8 10 12 epochs training time (s)
- Fluid-damper example:

• We want to also handle non-smooth (and non-convex) regularization terms

$$\min_{\theta_x, \theta_y, x_0} \quad r(x_0, \theta_x, \theta_y) + \sum_{k=0}^{N-1} \ell(y_k, f_y(x_k, \theta_y)) + g(\theta_x, \theta_y)$$

s.t.
$$x_{k+1} = f_x(x_k, u_k, \theta_x)$$

• Idea: use alternating direction method of multipliers (ADMM) by splitting

$$\min_{\theta_x, \theta_y, x_0, \nu_x, \nu_y} \quad r(x_0, \theta_x, \theta_y) + \sum_{k=0}^{N-1} \ell(y_k, f_y(x_k, \theta_y)) + g(\nu_x, \nu_y)$$
s.t.
$$x_{k+1} = f_x(x_k, u_k, \theta_x)$$

$$\begin{bmatrix} \nu_x \\ \nu_y \end{bmatrix} = \begin{bmatrix} \theta_x \\ \theta_y \end{bmatrix}$$

©2022 A. Bemporad - MPC: from basics to learning-based design

TRAINING RNNS BY SEQUENTIAL LS AND ADMM

(Bemporad, 2021 - http://arxiv.org/abs/2112.15348)

ADMM + Seq. LS = NAILS algorithm (Nonconvex ADMM Iterations and Sequential LS)

$$\begin{bmatrix} x_0^{t+1} \\ \theta_x^{t+1} \\ \theta_y^{t+1} \\ \theta_y^{t+1} \end{bmatrix} = \arg \min_{x_0, \theta_x, \theta_y} V(x_0, \theta_x, \theta_y) + \frac{\rho}{2} \left\| \begin{bmatrix} \theta_x - \nu_x^t + w_x^t \\ \theta_y - \nu_y^t + w_y^t \end{bmatrix} \right\|_2^2 \quad \text{(sequential) LS}$$

$$\begin{bmatrix} \nu_x^{t+1} \\ \nu_y^{t+1} \end{bmatrix} = \operatorname{prox}_{\frac{1}{\rho}g}(\theta_x^{t+1} + w_x^t, \theta_y^{t+1} + w_y^t) \quad \text{proximal step}$$

$$\begin{bmatrix} w_x^{t+1} \\ w_y^{t+1} \end{bmatrix} = \begin{bmatrix} w_x^h + \theta_x^{t+1} - \nu_x^{t+1} \\ w_y^h + \theta_y^{t+1} - \nu_y^{t+1} \end{bmatrix} \quad \text{update dual varse}$$

• Fluid-damper example: group-Lasso regularization $g(\nu_i^g) = \tau \sum_{i=1}^{n_x} \|\nu_i^g\|_2$ to zero entire rows and columns and reduce state-dimension automatically



©2022 A. Bemporad - MPC: from basics to learning-based design

TRAINING RNNS BY SEQUENTIAL LS AND ADMM

(Bemporad, 2021 - http://arxiv.org/abs/2112.15348)

• Fluid-damper example: quantization of θ_x , θ_y for simplifying model arithmetic +ReLU activation function

$$g(\theta_i) = \begin{cases} 0 & \text{if } \theta_i \in \mathcal{Q} \\ +\infty & \text{otherwise} \end{cases} \qquad \qquad \mathcal{Q} = \text{multiples of } 0.1 \text{ between } -0.5 \text{ and } 0.5 \end{cases}$$

- NRMSE = 83.10 (training), 80.51 (test)
- NRMSE = 8.83 (training), 2.69 (test) ← no ADMM, just quantize after training
- Training time: \approx 5 s

.

- Note: no convergence to a global minimum is guaranteed
- NAILS = very flexible & efficient learning algorithm for control-oriented RNNs

LEARNING HYBRID PREDICTION MODELS

LEARNING HYBRID MODELS

• Switching dynamics are better captured by piecewise affine (PWA) models and handled by hybrid MPC techniques

$$v(k) = \begin{cases} F_1 z(k) + g_1 & \text{if } H_1 z(k) \le K_1 \\ \vdots \\ F_s z(k) + g_s & \text{if } H_s z(k) \le K_s \end{cases}$$
$$v(k) = \begin{bmatrix} x(k+1) \\ y(k) \end{bmatrix}, \quad z(k) = \begin{bmatrix} x(k) \\ u(k) \end{bmatrix}$$



• **PWA regression**: learn **both** the $\{F_i, g_i\}$ and the partition $\{H_i, K_i\}$

(Ferrari-Trecate, Muselli, Liberati, Morari, 2003) (Roll, Bemporad, Ljung, 2004) (Juloski, Wieland, Heemels, 2004) (Bemporad, Garulli, Paoletti, Vicino, 2005) (Pillonetto, 2016) (Breschi, Piga, Bemporad, 2016) (...)

• Any ML technique can be applied that leads to PWA models, such as (leaky)ReLU-NNs, decision trees, softmax regression, KNN, ...

PARC - PIECEWISE AFFINE REGRESSION AND CLASSIFICATION

(Bemporad, 2021, arXiv)

- New Piecewise Affine Regression and Classification (PARC) algorithm
- Training dataset:
 - feature vector $z \in \mathbb{R}^n$ (categorical features one-hot encoded in $\{0, 1\}$)
 - target vector $v_c \in \mathbb{R}^{m_c}$ (numeric), $v_{di} \in \{w_{di}^1, \dots, w_{di}^{m_i}\}$ (categorical)
- PARC iteratively clusters training data in K sets and fit linear predictors
 - 1. fit $v_c = a_j z + b_j$ by ridge regression (= ℓ_2 -regularized least squares)
 - 2. fit $v_{di} = w_{di}^{h_*}$, $h_* = \arg \max\{a_{dih}^h z + b_{di}^h\}$ by softmax regression
 - 3. fit a convex PWL separation function by softmax regression

$$\Phi(z) = \omega^{j(z)} z + \gamma^{j(z)}, \qquad j(z) = \min\left\{\arg\max_{j=1,\dots,K} \{\omega^j z + \gamma^j\}\right\}$$

- Data reassigned to clusters based on weighted fit/PWL separation criterion
- PARC is a block-coordinate descent algorithm ⇒ (local) convergence ensured
 ©2022 A. Bemporad MPC: from basics to learning-based design 53/66

PARC - PIECEWISE AFFINE REGRESSION AND CLASSIFICATION

- Simple PWA regression example:
 - 1000 samples of $y = \sin(4x_1 5(x_2 0.5)^2) + 2x_2$ (use 80% for training)
 - Look for PWA approximation over K = 10 polyhedral regions



Code download: 🔁

http://cse.lab.imtlucca.it/~bemporad/parc/

DATA-DRIVEN HYBRID-MPC EXAMPLE

• Example: moving cart and bumpers + heat transfer during bumps.

Spring and viscous forces are nonlinear.

- Categorical input $F \in \{-\bar{F}, 0, \bar{F}\}$ and categorical output $c \in \{green, yellow, red\}$
- Continuous-time system simulated for 2,000 s, sample time = 0.5 s (=4000 training samples)
- Feature vector $z_k = [y_k, \dot{y}_k, T_k, F_k]$
- Target vector $v_k = [y_{k+1}, \dot{y}_{k+1}, T_{k+1}, c_k]$
- Hybrid model learned by PARC (K = 5 regions)





DATA-DRIVEN HYBRID-MPC EXAMPLE

• **Open-loop** simulation on 500 s **test** data:





continuous-time system

discrete-time PWA model

• Model fit is good enough for MPC design purposes

DATA-DRIVEN HYBRID-MPC EXAMPLE

• MPC problem with prediction horizon N = 9:

$$\begin{array}{ll} \min_{F_0,\dots,F_{N-1}} & \displaystyle\sum_{k=0}^{N-1} |c_k - \mathbf{1}| + 0.25 |F_k| \\ \text{s.t.} & F_k \in \{-\bar{F}, 0, \bar{F}\} \\ & \mathsf{PWA} \text{ model equations} \end{array}$$





- Data-driven hybrid MPC controller can keep temperature in yellow zone
- Approximate explicit MPC: fit a decision tree on 10,000 samples (accuracy: 99.9%). CPU time = 52÷67 μs. Closed-loop trajectories very similar.

LEARNING OPTIMAL MPC CALIBRATION

MPC CALIBRATION PROBLEM

- The design depends on a vector x of MPC parameters
- MPC parameters are intuitive to set (e.g., weights)
- Still, can we auto-calibrate them?



• Define a **performance index** *f* over a closed-loop simulation or real experiment. For example:



 Auto-tuning = find the best combination of parameters by solving the global optimization problem

$$\min_{x} f(x)$$

AUTO-TUNING - GLOBAL OPTIMIZATION ALGORITHMS

- Several derivative-free global optimization algorithms exist: (Rios, Sahidinis, 2013)
 - Lipschitzian-based partitioning techniques:
 - DIRECT (Divide in RECTangles) (Jones, 2001)
 - Multilevel Coordinate Search (MCS) (Huyer, Neumaier, 1999)
 - Response surface methods
 - Kriging (Matheron, 1967), DACE (Sacks et al., 1989)
 - Efficient global optimization (EGO) (Jones, Schonlau, Welch, 1998)
 - Bayesian optimization (Brochu, Cora, De Freitas, 2010)
 - Genetic algorithms (GA) (Holland, 1975)
 - Particle swarm optimization (PSO) (Kennedy, 2010)
 - ...
- New method: radial basis function surrogates + inverse distance weighting

(GLIS) (Bemporad, 2020)

cse.lab.imtlucca.it/~bemporad/glis

- Pros:
 - Selection of calibration parameters x to test is fully automatic
 - Applicable to any calibration parameter (weights, horizons, solver tolerances, ...)
 - **a** Rather arbitrary performance index f(x) (tracking performance, response time, worst-case number of flops, ...)
- Cons:
 - **•** Need to **quantify** an objective function f(x)
 - No room for qualitative assessments of closed-loop performance
 - Often have multiple objectives, not clear how to blend them in a single one

ACTIVE PREFERENCE LEARNING

- Objective function f(x) is not available (latent function)
- We can only express a **preference** between two choices:

$$\pi(x_1, x_2) = \begin{cases} -1 & \text{if } x_1 \text{ "better" than } x_2 & [f(x_1) < f(x_2)] \\ 0 & \text{if } x_1 \text{ "as good as" } x_2 & [f(x_1) = f(x_2)] \\ 1 & \text{if } x_2 \text{ "better" than } x_1 & [f(x_1) > f(x_2)] \end{cases}$$

• We want to find a global optimum x^* (="better" than any other x)

find x^* such that $\pi(x^*, x) \leq 0, \forall x \in \mathcal{X}, \ell \leq x \leq u$

- Active preference learning: iteratively propose a new sample to compare
- Key idea: learn a surrogate of the (latent) objective function from preferences

SEMI-AUTOMATIC TUNING BY PREFERENCE-BASED LEARNING

- Use preference-based optimization (GLISp) algorithm for semi-automatic tuning of MPC (Zhu, Bemporad, Piga, 2021)
- Latent function = calibrator's (unconscious) score of closed-loop MPC performance
- GLISp proposes a new combination x_{N+1} of MPC parameters to test
- By observing test results, the calibrator expresses a **preference**, telling if x_{N+1} is "**better**", "**similar**", or "**worse**" than current best combination
- Preference learning algorithm: update the surrogate $\hat{f}(x)$ of the latent function, optimize the acquisition function, ask preference, and iterate



PREFERENCE-BASED TUNING: MPC EXAMPLE

• Example: calibration of a simple MPC for lane-keeping (2 inputs, 3 outputs)

$$\begin{cases} \dot{x} = v\cos(\theta + \delta) \\ \dot{y} = v\sin(\theta + \delta) \\ \dot{\theta} = \frac{1}{L}v\sin(\delta) \end{cases}$$



• Multiple control objectives:

"optimal obstacle avoidance", "pleasant drive", "keep CPU time small", ...

not easy to quantify in a single function

- 5 MPC parameters to tune:
 - sampling time
 - prediction and control horizons
 - weights on input increments $\Delta v, \Delta \delta$
PREFERENCE-BASED TUNING: MPC EXAMPLE

• Preference query window:



PREFERENCE-BASED TUNING: MPC EXAMPLE

• Convergence after 50 GLISp iterations (=49 queries):



Optimal MPC parameters:

- sample time = 85 ms (CPU time = 80.8 ms)
- prediction horizon = 16
- control horizon = 5
- weight on Δv = 1.82
- weight on $\Delta\delta$ = 8.28



- Note: no need to define a closed-loop performance index explicitly!
- Extended to handle also unknown constraints (Zhu, Piga, Bemporad, 2021)

CONCLUSIONS

- Learning-based MPC is a formidable combination for advanced control:
 - MPC / on-line optimization is an extremely powerful control methodology
 - ML extremely useful to get control-oriented models and control laws from data
- Ignoring ML tools would be a mistake (a lot to "learn" from machine learning)
- ML cannot replace control engineering:
 - Black-box modeling can be a failure. Better use gray-box models when possible
 - Approximating the control law can be a failure. Don't abandon on-line optimization
 - Pure AI-based reinforcement learning methods can be also a failure
- A wide spectrum of research opportunities and new practices is open !

