

Preference-based MPC calibration

Mengjia Zhu, Alberto Bemporad, Dario Piga

Abstract—Automating the calibration of the parameters of a control policy by means of global optimization requires quantifying a closed-loop performance function. As this can be impractical in many situations, in this paper we suggest a semi-automated calibration approach that requires instead a human calibrator to express a *preference* on whether a certain control policy is “better” than another one, therefore eliminating the need of an explicit performance index. In particular, we focus our attention on semi-automated calibration of *Model Predictive Controllers* (MPCs), for which we attempt computing the set of best calibration parameters by employing the recently-developed active preference-based optimization algorithm GLISp. Based on the preferences expressed by the human operator, GLISp learns a surrogate of the underlying closed-loop performance index that the calibrator (unconsciously) uses and proposes, iteratively, a new set of calibration parameters to him or her for testing and for comparison against previous experimental results. The resulting semi-automated calibration procedure is tested on two case studies, showing the capabilities of the approach in achieving near-optimal performance within a limited number of experiments.

I. INTRODUCTION

The design of *Model Predictive Controllers* (MPC) typically requires the tuning of several parameters such as the prediction and control horizons, the weight matrices defining the cost function, various numerical solver tolerances, etc. A calibrator typically adjusts these knobs based on experience and trial-and-error, until the closed-loop system behaves as he or she desires. Such a tuning process thus requires skilled calibrators, domain knowledge, and it can be costly and time consuming.

To automate the tuning process, usually a figure of merit is defined to quantitatively assess closed-loop performance, and experiment-driven optimization algorithms are usually adopted to find near-optimal MPC parameters. In particular, black-box global optimization methods based on surrogate functions, like *Bayesian Optimization* (BO) [1] and the recently introduced GLIS (GLobal optimization based on Inverse distance weighting and radial basis function Surrogates) algorithm [2], have been recently applied for MPC parameter tuning [3], [4], choice of the MPC predictive model [5], [6] and also in other control engineering problems and applications such as PID and state-feedback control tuning [7], [8], position and force control in robot

manipulators [9], [10], and control of mobile robots and quadrotors [11], [12], just to cite a few.

The aforementioned optimization approaches iteratively suggest new parameters to be tested based on a surrogate function, which is estimated from closed-loop performance figures gathered from previous experiments. An exploration term takes care of sufficiently covering the set of parameters to search, thus avoiding the solver to be trapped in a local minimum.

However, in order to use such methods for MPC calibration, it is essential to have a well-defined performance index that captures the desired closed-loop behavior of the system. Unfortunately, in many practical control applications a performance is usually formulated based on multiple criteria, and thus it is difficult for a calibrator to formally define and quantify objectively the scoring function. On the other hand, it is usually easier for a calibrator to express a preference (such as “controller A is better than controller B”) between the outcome of two experiments.

A. Contribution

Motivated by the above consideration, this paper proposes a novel calibration approach to tune the MPC parameters based on pairwise preferences between experiment outcomes. The approach is semi-automated in that the control parameters to be tested are selected automatically by an algorithm while closed-loop performance is assessed manually by the calibrator, that therefore is no longer required to formulate a performance index upfront.

The proposed approach for preference-based MPC calibration relies on the derivative-free global optimization algorithm recently developed in [13], which iteratively proposes a new comparison to the calibrator to make, based on actively learning a surrogate of the latent objective function from past sampled decision vectors and pairwise preferences. Preference-based optimization method is becoming popular in the field of reinforcement learning (RL) [14] and a comprehensive review is presented in the survey paper [15]. Inverse (reinforcement) learning and semi-supervised machine learning approaches are also closely related to preference-based learning, which has been applied to parameter tuning, such as for autonomous driving [16], [17], [18]. The main concern of the inverse-learning approach is that it can not improve upon the demonstrations by the expert. The algorithm implemented in this paper [13], called GLISp (preference-based GLIS) proposed a different acquisition function form that balances the trade-off between exploitation and exploration. GLISp has been shown to be very efficient in terms of number of experiments and comparisons required to compute

M. Zhu and A. Bemporad are with IMT School for Advanced Studies Lucca, Lucca, Italy. mengjia.zhu@imtlucca.it; alberto.bemporad@imtlucca.it

D. Piga is with IDSIA Dalle Molle Institute for Artificial Intelligence, SUPSI-USI, Lugano, Switzerland. dario.piga@supsi.ch

This paper was partially supported by the Italian Ministry of University and Research under the PRIN’17 project “Data-driven learning of constrained control systems”, contract no. 2017J89ARP.

the global optimum, which is a major drawback of many preference-based RL methods [15].

We show the efficiency of the proposed preference-based MPC calibration in two case studies. The first one considers the control of a *Continuous Stirring Tank Reactor* (CSTR), while the second one is related to autonomous driving of a vehicle with obstacle avoidance. In both cases, overall satisfactory performance is achieved within a relatively small number of closed-loop experiments, without the hard and time-consuming need to specify a quantitative scoring function driving a fully-automated MPC calibration.

The rest of this paper is organized as follows. Section II describes the MPC calibration problem. The preference-based tuning approach based on GLISp is presented in Section III. The application of the proposed approach for two case studies are discussed in Section IV. Finally, conclusions and directions for future research are drawn in Section V.

II. PROBLEM DESCRIPTION

Let us consider the problem of controlling a nonlinear multi-input multi-output system described by the continuous-time state-space representation:

$$\begin{aligned}\dot{x} &= f(x, u) \\ y &= g(x, u),\end{aligned}\quad (1)$$

where $x \in \mathbb{R}^{n_x}$ and $\dot{x} \in \mathbb{R}^{n_x}$ are the state vector and its time derivative, respectively; $u \in \mathbb{R}^{n_u}$ is the control input; $y \in \mathbb{R}^{n_y}$ is the vector of controlled outputs; and $f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$ and $g : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_y}$ are the state and output mappings, respectively.

A popular strategy to achieve a reference-tracking objective for a system described by (1) under input and output constraints, is linear-time varying MPC, a.k.a. real-time iteration scheme [19], [20]. The MPC is designed based on the following predictive model obtained via linearization of (1) around a nominal trajectory \bar{x}_k , \bar{u}_k , \bar{y}_k and discretization with sampling time T_s , resulting in the prediction model

$$\begin{aligned}\tilde{x}_{k+1} &= A_k \tilde{x}_k + B_k \tilde{u}_k \\ \tilde{y}_k &= C_k \tilde{x}_k + D_k \tilde{u}_k,\end{aligned}\quad (2)$$

where $\tilde{x}_k = x_k - \bar{x}_k$, $\tilde{u}_k = u_k - \bar{u}_k$ and $\tilde{y}_k = y_k - \bar{y}_k$.

At each sampling time t , the MPC action $u_{t|t}$ to apply to the system is computed by solving the Quadratic Programming (QP) problem:

$$\begin{aligned}\min_{\{u_{t+k|t}\}_{k=0}^{N_p-1}, \varepsilon} & \sum_{k=0}^{N_p-1} \|y_{t+k|t} - y_{t+k}^{\text{ref}}\|_{Q_y}^2 \\ & + \sum_{k=0}^{N_p-1} \|u_{t+k|t} - u_{t+k}^{\text{ref}}\|_{Q_u}^2 \\ & + \sum_{k=0}^{N_p-1} \|\Delta u_{t+k|t}\|_{Q_{\Delta u}}^2 + Q_\varepsilon \|\varepsilon\|^2\end{aligned}\quad (3)$$

s.t. model equation (2) and the following constraints

$$\begin{aligned}y_{\min} - \varepsilon &\leq y_{t+k|t} \leq y_{\max} + \varepsilon, \quad k = 1, \dots, N_p \\ u_{\min} &\leq u_{t+k|t} \leq u_{\max}, \quad k = 1, \dots, N_p \\ \Delta u_{\min} &\leq \Delta u_{t+k|t}, \quad k = 1, \dots, N_p \\ \Delta u_{t+k|t} &\leq \Delta u_{\max}, \quad k = 1, \dots, N_p \\ u_{t+N_u+j|t} &= u_{t+N_u|t}, \quad j = 1, \dots, N_p - N_u,\end{aligned}\quad (4)$$

where $\|\nu\|_Q^2$ is the weighted squared norm, i.e., $\nu^\top Q \nu$; u_{ref} and y_{ref} are the input and output references, respectively; and $\Delta u_{t+k|t} = u_{t+k|t} - u_{t+k-1|t}$.

Several tuning parameters appear in the MPC problem (3)–(4) and must be tuned, such as:

- prediction horizon N_p and control horizon N_u ;
- positive-semidefinite weight matrices Q_y , Q_u , $Q_{\Delta u}$;
- positive constant Q_ε used to soften the constraints and thus to guarantee feasibility of the optimization problem (3);
- tolerances used in the QP algorithm solving (3)–(4).

In order to compact the notation, all the MPC knobs are collected in a parameter vector $\theta \in \Theta \subseteq \mathbb{R}^{n_\theta}$, where Θ is a given bounded set in which the optimal tuning for θ is sought.

In this paper, we present an experiment-driven approach to tune the MPC knobs θ in order to optimize the overall closed-loop performance based on pairwise preferences expressed by a controller calibrator. More precisely, we assume that we do not have a quantitative scoring function used by the calibrator to measure the closed-loop performance index and thus to be usable for auto-tuning θ by global optimization. The only assumption we make is that for a given pair of different MPC parameters $\theta_1, \theta_2 \in \Theta$, only a preference expressed by the calibrator is available in terms of “performance achieved with parameters θ_1 was better (or worse, or similar) than the one achieved with θ_2 ”.

The rationale of our problem formulation is that, in multi-criteria decision making such as in control system design, it is difficult (and sometimes impossible) to quantify an overall scoring function, but anyway it is easier for a calibrator to assess a *preference* between the outcomes of two experiments.

In order to compute the optimal MPC parameters θ , the active preference learning algorithm GLISp [13] is used, which proposes the experiments to be performed for tuning θ , through pairwise comparisons, that we review in the next section.

III. PREFERENCE-BASED TUNING

We adopt the active preference learning algorithm GLISp proposed in [13] to iteratively suggest a sequence of MPC parameters $\theta_1, \dots, \theta_N$ to be tested and compared such that θ_N approaches the “optimal” combination of parameters as N grows.

Formally, given two possible MPC knobs θ_1 and θ_2 , we define the *preference function* $\pi : \mathbb{R}^{n_\theta} \times \mathbb{R}^{n_\theta} \rightarrow \{-1, 0, 1\}$ as

$$\pi(\theta_1, \theta_2) = \begin{cases} -1 & \text{if } \theta_1 \text{ “better” than } \theta_2 \\ 0 & \text{if } \theta_1 \text{ “as good as” } \theta_2 \\ 1 & \text{if } \theta_2 \text{ “better” than } \theta_1. \end{cases}\quad (5)$$

The modeling assumption behind GLISp is that the calibrator assigns preferences based on an underlying closed-loop performance index J that he or she wants to minimize:

$$\pi(\theta_1, \theta_2) = \begin{cases} -1 & \text{if } J(\theta_1) < J(\theta_2) \\ 0 & \text{if } J(\theta_1) = J(\theta_2) \\ 1 & \text{if } J(\theta_1) > J(\theta_2), \end{cases} \quad (6)$$

where J is totally unknown (to the algorithm and to the calibrator as well). Only the evaluations $\pi(\theta_1, \theta_2)$ of the preference function are accessible.

The goal of the preference-based MPC calibration approach discussed in this paper is to find the optimal MPC parameters $\theta^* \in \Theta$ such that θ^* is “better” (or “no worse”) than any other parameter θ according to the preference function π , namely

$$\theta^* \text{ such that } \pi(\theta^*, \theta) \leq 0, \quad \forall \theta \in \Theta. \quad (7)$$

If such a vector θ^* is found, as (7) implies that $J(\theta^*) \leq J(\theta)$, $\forall \theta \in \Theta$, clearly we have also found a global minimizer θ^* of J on Θ .

A. Training a surrogate function from preferences

Assume that we have generated $N \geq 2$ samples $\{\theta_1 \dots \theta_N\}$ of the decision vector, with $\theta_i, \theta_j \in \mathbb{R}^{n_\theta}$ such that $\theta_i \neq \theta_j$, $\forall i \neq j$, $i, j = 1, \dots, N$. For each of these parameters, a closed-loop experiment is performed and the calibrator has provided a *preference vector* $B = [b_1 \dots b_M]^T \in \{-1, 0, 1\}^M$ with

$$b_h = \pi(\theta_{i(h)}, \theta_{j(h)}), \quad (8)$$

where M is the number of expressed preferences, $1 \leq M \leq \binom{N}{2}$, $h \in \{1, \dots, M\}$, $i(h), j(h) \in \{1, \dots, N\}$, $i(h) \neq j(h)$. Note that the element b_h of vector B represents the preference expressed by the calibrator between the closed-loop performance achieved with parameters $\theta_{i(h)}$ and $\theta_{j(h)}$.

The observed preferences are then used to learn a surrogate function $\hat{J} : \mathbb{R}^{n_\theta} \rightarrow \mathbb{R}$ of the underlying performance index J . The surrogate \hat{J} is constructed by imposing the constraints

$$\hat{\pi}(\theta_{i(h)}, \theta_{j(h)}) = \pi(\theta_{i(h)}, \theta_{j(h)}), \quad \forall h = 1, \dots, M, \quad (9)$$

on \hat{J} , where $\hat{\pi}$ is defined from \hat{J} as in (6).

The function \hat{J} is parametrized as the following linear combination of Radial Basis Functions (RBFs) [21], [22]:

$$\hat{J}(\theta) = \sum_{k=1}^N \beta_k \phi(\epsilon d(\theta, \theta_i)), \quad (10)$$

where $d : \mathbb{R}^{n_\theta} \times \mathbb{R}^{n_\theta} \rightarrow \mathbb{R}$ is the squared Euclidean distance

$$d(\theta_1, \theta_2) = \|\theta_1 - \theta_2\|_2^2, \quad (11)$$

$\epsilon > 0$ is a scalar parameter, $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is an RBF, and $\beta = [\beta_1 \dots \beta_N]^T$ are the unknown coefficients to be computed based on the preferences imposed in (9). Examples of RBFs are $\phi(\epsilon d) = \frac{1}{1+(\epsilon d)^2}$ (*inverse quadratic*), $\phi(\epsilon d) = e^{-(\epsilon d)^2}$ (*Gaussian*), $\phi(\epsilon d) = (\epsilon d)^2 \log(\epsilon d)$ (*thin plate spline*), see more examples in [21], [2].

According to (9) and the preference relation (6), the following constraints are imposed on \hat{J} :

$$\begin{aligned} \hat{J}(\theta_{i(h)}) &\leq \hat{J}(\theta_{j(h)}) - \sigma + \epsilon_h & \text{if } \pi(\theta_{i(h)}, \theta_{j(h)}) = -1 \\ \hat{J}(\theta_{i(h)}) &\geq \hat{J}(\theta_{j(h)}) + \sigma - \epsilon_h & \text{if } \pi(\theta_{i(h)}, \theta_{j(h)}) = 1 \\ |\hat{J}(\theta_{i(h)}) - \hat{J}(\theta_{j(h)})| &\leq \sigma + \epsilon_h & \text{if } \pi(\theta_{i(h)}, \theta_{j(h)}) = 0 \end{aligned} \quad (12)$$

for all $h = 1, \dots, M$, where $\sigma > 0$ is a given tolerance and ϵ_h are positive slack variables.

Accordingly, and similarly to Support Vector Machines (SVMs) [23], the coefficient vector β describing the surrogate \hat{J} is obtained by solving the convex QP problem

$$\begin{aligned} \min_{\beta, \epsilon} \quad & \sum_{h=1}^M c_h \epsilon_h + \frac{\lambda}{2} \sum_{k=1}^N \beta_k^2 \\ \text{s.t.} \quad & \sum_{k=1}^N (\phi(\epsilon d(\theta_{i(h)}, \theta_k)) - \phi(\epsilon d(\theta_{j(h)}, \theta_k))) \beta_k \\ & \leq -\sigma + \epsilon_h, \quad \forall h : b_h = -1 \\ & \sum_{k=1}^N (\phi(\epsilon d(\theta_{i(h)}, \theta_k)) - \phi(\epsilon d(\theta_{j(h)}, \theta_k))) \beta_k \\ & \geq \sigma - \epsilon_h, \quad \forall h : b_h = 1 \\ & \left| \sum_{k=1}^N (\phi(\epsilon d(\theta_{i(h)}, \theta_k)) - \phi(\epsilon d(\theta_{j(h)}, \theta_k))) \beta_k \right| \\ & \leq \sigma + \epsilon_h, \quad \forall h : b_h = 0 \\ & h = 1, \dots, M \end{aligned} \quad (13)$$

where c_h are positive weights, for example $c_h = 1$, $\forall h = 1, \dots, M$. The slack variables ϵ_h in (13) are used to relax the constraints imposed by the preference vector B in (9). Constraint infeasibility might be due to an inappropriate selection of the RBF (namely, poor flexibility in the parametric description of the surrogate \hat{J}) and/or to outliers in the acquired preferences, for instance due to inconsistent assessments done by the calibrator. The scalar λ in the cost function (13) is a regularization parameter. When $\lambda > 0$, problem (13) is a QP problem that admits a unique solution because $c_h > 0$ for all $h = 1, \dots, M$. If $\lambda = 0$, problem (13) becomes a Linear Program (LP), whose solution may not be unique.

We remark that computing the surrogate function \hat{J} requires one to choose the hyper-parameter ϵ defining the shape of the RBFs ϕ in (10). This parameter can be chosen through K -fold cross-validation [24], by testing the capabilities of \hat{J} in reconstructing the preferences in slices of the dataset not used to estimate \hat{J} .

B. Acquisition function

Once a surrogate \hat{J} is estimated, this function can be in principle minimized in order to find the optimal MPC parameter vector θ . More specifically, the following steps can be followed: (i) generate a new sample by pure minimization of the estimated surrogate function \hat{J} defined in (10), *i.e.*,

$$\theta_{N+1} = \arg \min \hat{J}(\theta) \text{ s.t. } \theta \in \Theta;$$

(ii) ask the calibrator to evaluate the preference $\pi(\theta_{N+1}, \theta_N^*)$, where $\theta_N^* \in \mathbb{R}^{n_\theta}$ is the best vector of

MPC parameters found so far, corresponding to the smallest index i^* such that

$$\pi(\theta_{i^*}, \theta_i) \leq 0, \forall i = 1, \dots, N; \quad (14)$$

(iii) update the estimate of \hat{J} through (13); and (iv) iterate over N .

Such a procedure, which only *exploits* the current available observations in finding the optimal MPC parameter vector θ , may easily miss the global minimum of (7). Therefore, looking only at the surrogate function \hat{J} is not enough to search for a new sample θ_{N+1} . A term promoting the *exploration* of the parameter space should thus be considered.

In the GLISp algorithm, an acquisition function is employed to balance exploitation vs. exploration when generating the new sample θ_{N+1} . As proposed in [2], the exploration function is constructed by using the inverse distance weighting (IDW) function $z : \mathbb{R}^{n_\theta} \rightarrow \mathbb{R}$ defined by

$$z(\theta) = \begin{cases} 0 & \text{if } \theta \in \{\theta_1, \dots, \theta_N\} \\ \tan^{-1} \left(\frac{1}{\sum_{i=1}^N w_i(\theta)} \right) & \text{otherwise} \end{cases} \quad (15)$$

where $w_i(\theta) = \frac{1}{d^2(\theta, \theta_i)}$. Clearly $z(\theta) = 0$ for all parameters already tested, and $z(\theta) > 0$ in $\mathbb{R}^{n_\theta} \setminus \{\theta_1, \dots, \theta_N\}$. The arc tangent function in (15) avoids that $z(\theta)$ gets excessively large far away from all sampled points.

Then, given an exploration parameter $\delta \geq 0$, the *acquisition function* $a : \mathbb{R}^{n_\theta} \rightarrow \mathbb{R}$ is constructed as

$$a(\theta) = \frac{\hat{J}(\theta)}{\Delta \hat{J}} - \delta z(\theta), \quad (16)$$

where

$$\Delta \hat{J} = \max_i \{\hat{J}(\theta_i)\} - \min_i \{\hat{J}(\theta_i)\}$$

is the range of the surrogate function on the samples in $\{\theta_1, \dots, \theta_N\}$ and is used in (16) as a normalization factor to simplify the choice of the exploration parameter δ . Clearly $\Delta \hat{J} \geq \sigma$ if at least one comparison $b_h = \pi(\theta_{i(h)}, \theta_{j(h)}) \neq 0$.

As discussed below, given a set $\{\theta_1, \dots, \theta_N\}$ of samples and a vector B of preferences defined by (8), the next MPC parameter θ_{N+1} to test is computed as the solution of the (non-convex) optimization problem

$$\theta_{N+1} = \arg \min_{\theta \in \Theta} a(\theta). \quad (17)$$

The preference-based global optimization algorithm based on the RBF interpolants (10) and the acquisition function (16) is detailed in [13] and will be referred to as the GLISp algorithm since now on. We will use GLISp to compute the optimal MPC parameter vector θ^* by preferences. We use the Particle Swarm Optimization (PSO) algorithm of [25] to solve problem (17). Note that the construction of the acquisition function a is rather heuristic, therefore finding highly accurate solutions of (17) is not required.

In the acquisition function (16), the exploration parameter δ promotes sampling the space in Θ in areas that have not been explored yet. While, as observed earlier, $\delta = 0$ can make the GLISp algorithm rely only on the surrogate

function \hat{J} and miss the global optimum, setting $\delta \gg 1$ makes the GLISp algorithm exploring the entire feasible region regardless of the results of the comparisons. For a sensitivity analysis example of GLISp with respect to δ , the reader is referred to [13, Section 7.5]. Regarding the other main hyper-parameter ϵ of GLISp defining the RBF in (10), during the active learning phase K -fold cross-validation is executed repeatedly to automatically choose and possibly adapt ϵ .

We finally remark that, in executing the GLISp, some values of the MPC knobs may lead to an unstable closed-loop behaviour. In this case, the experiment should be interrupted (e.g., for safety reasons) and the calibrator may still express a preference. It is also possible that in the initialization phase of GLISp a comparison should be performed over two experiments both exhibiting closed-loop instability. In this case, one experiment can be preferred to another one, for instance, based on the time it could run before being interrupted.

IV. CASES STUDIES

We apply the proposed preference-based MPC calibration method on two case studies: control of a *Continuous Stirring Tank Reactor* (CSTR) and autonomous driving with obstacle avoidance. The role of the calibrator is played by the first author of the paper, which compares the observed closed-loop performance manually according to a qualitative scoring function constructed in her mind based on engineering insights.

In both case studies we set the maximum number of function evaluations $N_{\max} = 50$, the exploration parameter $\delta = 0.3$, and the tolerance parameter $\sigma = 10^{-6}$. The GLISp algorithm is initialized with $N_{\text{init}} = 10$ random samples generated by Latin hypercube sampling [26] using the *lhsdesign* function of the *Statistics and Machine Learning Toolbox* of MATLAB [27]. The hyper-parameter ϵ defining the RBF functions (10) is initialized as 1 and updated at iterations 10, 20, 30, 40 via K -fold cross-validation with $K = 3$. The tests are run on an x64 with Intel i7-8550U 1.8 GHz CPU and 8GB of RAM. The *Model Predictive Control Toolbox* of MATLAB [28] is used for MPC design and simulation.

A. CSTR optimal steady-state switching policy

System description

The first case study considers the control of a CSTR, extensively described in [29]. The CSTR system consists of a jacketed non-adiabatic tank where an exothermic reaction occurs. The tank is assumed to be perfectly mixed with constant inlet and outlet rate. The chemical laws describing the CSTR reaction can be derived based on the energy and material balance, and are given by:

$$\begin{aligned} \frac{dT(t)}{dt} &= \frac{F}{V}(T_f(t) - T(t)) - \frac{H}{C_p \rho} r(t) - \frac{US}{C_p \rho V}(T(t) - T_c(t)) \\ \frac{dC_A(t)}{dt} &= \frac{F}{V}(C_{Af}(t) - C_A(t)) - r(t), \end{aligned} \quad (18)$$

TABLE I
CSTR PARAMETERS [29]

Parameter	Value	unit
F/V	1.0	hr^{-1}
k_0	$3.49e7$	hr^{-1}
H	5960	kcal/kgmol
E	11843	kcal/kgmol
ρC_p	500	$\text{kcal}/(\text{m}^3\text{K})$
US/V	150	$\text{kcal}/(\text{m}^3\text{K hr})$
T_{c0}	298	K
R	1.987	$\text{Kcal}/(\text{kgmol K})$

where V [m^3] is the volume of the reactor, and F [m^3/hr] is the rate of reactant A feeding the tank, which is equal to the rate of the product stream that exists from the reactor. Both V and F are assumed to be constant. $C_A(t)$ and $C_{A_f}(t)$ [kgmol/m^3] represent the concentration at time t of reactant A in the tank and in the inlet feed stream, respectively. $T(t)$, $T_f(t)$ and $T_c(t)$ [K] are the temperature of the reactor, of the inlet feed stream and of the coolant stream, respectively. Constant heat of reaction H [kcal/kgmol], fluid heat capacity C_p [$\text{kcal}/(\text{kg K})$] and density ρ [kg/m^3] are assumed. U [$\text{kcal}/(\text{m K hr})$] and S [m^2] are the overall heat transfer coefficient and heat exchange area, respectively. $r(t)$ [$\text{kgmol}/(\text{m}^3\text{hr})$] is the reaction rate per unit volume, which can be calculated through Arrhenius rate law:

$$r(t) = k_0 \exp\left(\frac{-E}{RT(t)}\right) C_A(t) \quad (19)$$

where k_0 [hr^{-1}] is the pre-exponential factor; E [kcal/kgmol] is the activation energy for the reaction; and R [$\text{Kcal}/(\text{kgmol K})$] is the gas constant. The CSTR process parameters are taken from [29] and listed in Table I.

Control objectives

Initially, the plant is operating at steady state with a reactant concentration $C_A = 8.56 \text{ kgmol}/\text{m}^3$ and low conversion. The objective is to design an MPC controller to achieve a steady-state concentration $C_A = 2 \text{ kgmol}/\text{m}^3$, thus increasing the conversion rate. The feed stream concentration $C_{A_f}(t)$ and temperature $T_f(t)$ are treated as measured disturbances, and the coolant temperature $T_c(t)$ is the control input. In this test, the condition of the feed stream is kept constant, with $T_f = 298.15 \text{ K}$ and $C_{A_f} = 10 \text{ kgmol}/\text{m}^3$.

At each time step, the nonlinear model of the CSTR system in (18) is linearized around its operating point, and a linear MPC problem is solved. Although the output variables of interest are both the reactor temperature $T(t)$ and the concentration $C_A(t)$ in the product stream, only the latter is tracked and compensated by MPC.

Different competitive objectives should be taken into account in the MPC calibration, such as fast steady-state transition, reasonable final target achievement, and low energy consumption. The following guidelines are used to assist the calibration. First, the steady-state switching is expected to be completed within two days. Second, the final achieved steady state should be within $\pm 3\%$ of the desired value $C_A = 2 \text{ kgmol}/\text{m}^3$. Third, in order to take into account energy consumption due to the cooling process, the temperature of




Fig. 1. CSTR query window for one iteration of GLISp. Left experiment is preferred to the right one because of a faster transient time.

the coolant stream $T_c(t)$ is restricted to be in the range of [284 310] K, with a maximum change at each time step ($T_{c_{max}}$) set to 10 K. The first two requirements just reflect the desired performance, and thus are not treated as *hard* constraints during calibration.

The following three design parameters are tuned: the sampling time T_s used to close the loop and to discretize the continuous-time model (18); the prediction horizon N_p and the weight $Q_{\Delta u}$ penalizing the input change. These tuning parameters T_s , N_p and $\log(Q_{\Delta u})$ are restricted to the ranges [0.25 1.5] hr, [4 40] and [-5 3], respectively. The control horizon N_u is set equal to $N_p/3$ and rounded to the closet integer. Other MPC knobs are fixed, with Q_y and Q_u equal to $\begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$ and 0, respectively.

Calibration process

At each iteration of GLISp, a set of new MPC parameters is suggested and the closed-loop experiment is performed in simulation. The experiment is interrupted either when the steady-state condition is reached or after 48 hrs (namely, two simulated days). The experiment is also interrupted if an un-safe behavior is observed. The calibrator is then asked to compare the performance of the experiment with the best performance achieved until that time and choose the one she prefers based on the control objectives and the aforementioned guidelines.

As an example, Fig. 1 shows one iteration of the semi-automated calibration process. At the top of the figure, the MPC design parameters, the achieved values of $C_{A_{end}}$ and duration of the switching process t_f are displayed. Both the MPC designs achieved the desired $C_{A_{end}}$ within 48 hrs. However, the transient and settling times of the left-hand-side experiment are shorter than the right-hand-side one. Moreover, the variations of the input signal T_c are smaller for the experiment on the left. Therefore, in this iteration, the left-hand-side MPC design is preferred.

Results

The GLISp algorithm terminates after $N_{max} = 50$ iterations (namely, closed-loop experiments) and hence 49

comparisons. The best MPC design parameters T_s , N_p and $\log(Q_{\Delta u})$ are found to be 0.31 hr, 26 and -1.79 , respectively. Fig. 2 (left panels) shows the corresponding trajectory of reactant concentration $C_A(t)$ and of the manipulated variable $T_c(t)$.

Fully-automated calibration

For the sake of comparison, a fully-automated calibration is performed. This requires to define a multi-objective quantitative scoring function describing the expected closed-loop performance. This step can be very hard and time-consuming, requiring proper scaling and balancing of the competitive control objectives.

The following closed-loop performance index function to be minimized is constructed after several trial-and-error iterations based on the preferences of the calibrator:

$$\frac{t_f}{t_{f_{max}}} + \frac{\sum_{k=1}^{N_T} (T_{c_k} - T_{c_{k-1}})^2}{T_{c_{max}} N_T} + \frac{|C_{A_{end}} - C_{A_{ref}}|}{AR\% C_{A_{ref}}} \quad (20)$$

where $t_{f_{max}} = 48$ hrs is the maximum duration of the switching process; $T_{c_{max}} = 10$ K is the maximum allowed temperature change of the coolant fluid between each time step; T_{c_k} and $T_{c_{k-1}}$ are the temperature of the coolant fluid at time step k and $k-1$, respectively; N_T is the total number of time steps in the steady-state transition; $AR\% = 3\%$ is the acceptable range of the final steady state concentration; $C_{A_{end}}$ and $C_{A_{ref}}$ [kgmol/m³] are the achieved and desired final steady-state concentration of reactant A .

The GLIS algorithm (without preference) of [2] is used for experiment-driven optimization of the cost function in (20) with respect to the MPC design parameters. For a fair comparison with the proposed preference-based algorithm, 50 function evaluations are performed. This value actually does not account for the number of trials (namely, experiments) needed to construct the scoring function in (20).

The obtained optimal MPC parameters T_s , N_p and $\log(Q_{\Delta u})$ are 0.25 hr, 26 and -0.91 . Fig. 2 (right panels) shows the corresponding closed-loop trajectory of $C_A(t)$ and of the manipulated variable $T_c(t)$. It can be observed that the preference- and the non-preference- based approaches achieve similar closed-loop performance.

Overall, this case study demonstrates the capability of the semi-automated approach for solving calibration tasks with multiple competitive objectives. For such calibration tasks, when using a fully-automated approach, significant efforts need to be devoted to construct a proper performance scoring function as in (20). Therefore, the semi-automated approach can greatly reduce calibration time and efforts by eliminating this step.

B. Autonomous driving vehicle

System description

As a second case study, we consider the problem of lane-keeping (LK) and obstacle-avoidance (OA) in autonomous driving. A simplified two degree-of-freedom bicycle model is used to describe the vehicle dynamics, with the front wheel as the reference point. The model involves the following




Fig. 2. CSTR closed-loop performance obtained by calibrating MPC parameters through the proposed semi-automated preference-based approach (left panels) and through a fully-automated approach minimizing the scoring function (20) (right panels).

three state variables: longitudinal x_f and lateral y_f [m] position of the front wheel, and yaw angle θ [rad]. The control inputs are the reference for the vehicle velocity v [m/s] and the steering angle δ_s [rad]. The system behaviour is thus described by the continuous-time kinematic model

$$\begin{aligned} \dot{x}_f &= v \cos(\theta + \delta_s) \\ \dot{y}_f &= v \sin(\theta + \delta_s) \\ \dot{\theta} &= \frac{v \sin(\delta_s)}{L} \end{aligned} \quad (21)$$

where L [m] is the vehicle length. This model is linearized around its operating point at each time step and used to design a linear MPC.

Control objectives

In tuning the MPC parameters, the objective is to keep the vehicle at the same horizontal lane with constant speed and to overtake other moving vehicles in an optimal way if they are within safety distance. However, similar to the CSTR case, it is difficult to define a proper quantitative scoring function for this multi-objective calibration task (for example, due to the ambiguity of transferring “optimal obstacle avoidance” into a mathematical formula). Therefore, to achieve good MPC performance using a fully-automated approach not based on preferences, either the calibrator needs to find a proper scoring function via trial-and-error (like in the previous CSTR case) or an advanced path planner model is required to provide a well-defined reference path to compare with. Both procedures can be time-consuming and computationally heavy. On the other hand, when using the semi-automated approach discussed in this paper, none of the pre-mentioned steps is required.

The test scenario is the following. The vehicle starts at position $(x_f, y_f) = (0, 0)$ m with θ equal to 0 rad and there is another vehicle (obstacle) at position (30, 0) m which is moving at constant speed (40 km/hr), with a constant yaw angle of 0 rad. Both vehicles are assumed to have a simplified rectangular shape with length equal to 4.5 m and width of

1.8 m. At nominal LK condition, the vehicle moves at 50 km/hr with $y_f = 0$ m. When the obstacle is within safety distance (10 m in this case), the vehicle needs to overtake it and the minimum lateral distance between two vehicles is set to 3 m. During LK and OA periods, the vehicle is allowed to vary its velocity in the range of [40 70] and [50 70] km/hr, respectively, with its reference velocity set to 50 and 60 km/hr, respectively. The maximum and minimum of θ in both LK and OA periods are $\pm \frac{\pi}{4}$ rad with the maximum and minimum rate of change between each time step set to ± 0.0873 rad/s.

Five MPC design parameters are tuned. The sampling time T_s is allowed to vary in the range [0.085 0.5] s. The prediction horizon N_p is restricted to [10 30] and the control horizon N_u is taken as a fraction ϵ_c of N_p rounded to the closest integer. Here, ϵ_c can take values in the range [0.1 1]. The weight matrix of manipulated variables ($Q_{\Delta u}$) is set to be diagonal with diagonal entries q_{u11} and q_{u22} such that $\log(q_{u11})$ and $\log(q_{u22})$ vary in the range [-5 3]. The other MPC design parameters are fixed, with Q_y and Q_u set to $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$ and $\begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}$, respectively.

Calibration process

The calibrator selects the preferred controller based on the following observations: (i) during both LK and OA periods, the worst-case computational time (t_{comp}) required for solving the QP problem (3) at each time step needs to be smaller than T_s , so that the MPC can be implemented in real time; (ii) during the LK phase, the vehicle should move at constant speed with y_f and θ close to 0 m and 0° ; (iii) during the OA phase, the vehicle should keep reasonable safety distance away from the obstacle and should guarantee passengers' comfort (i.e., aggressive lateral movements during overtaking should be avoided); (iv) the velocity in both LK and OA period should be close to the reference value with its variations kept to minimum; (v) variations of steering angles should not be aggressive; (vi) when there is a conflict combination among aforementioned criteria, criterion (i) has the highest priority and if the conflict is among criteria (ii)–(v), preference is given to the one leading to safer driving practice based on calibrator's experience.

The closed-loop test is simulated for 15 seconds. Fig. 3 shows the query window for one iteration of the calibration process. The MPC design parameters and t_{comp} are displayed at the top of the figure. In both of the cases illustrated in the figure, the QP problem (3) is solved within the chosen sampling time T_s . The performance of the right-hand-side experiment is preferred since in the left-hand-side experiment large lateral movements are present. Indeed, these movements can be much more dangerous (as the car may cross to the other lane) comparing to slightly more aggressive δ_s variations.

Results

The GLISp algorithm terminates after $N_{max} = 50$ function evaluations and 49 comparisons. The best MPC design parameters T_s , ϵ_c , N_p , $\log(q_{u11})$ and $\log(q_{u22})$ are found to

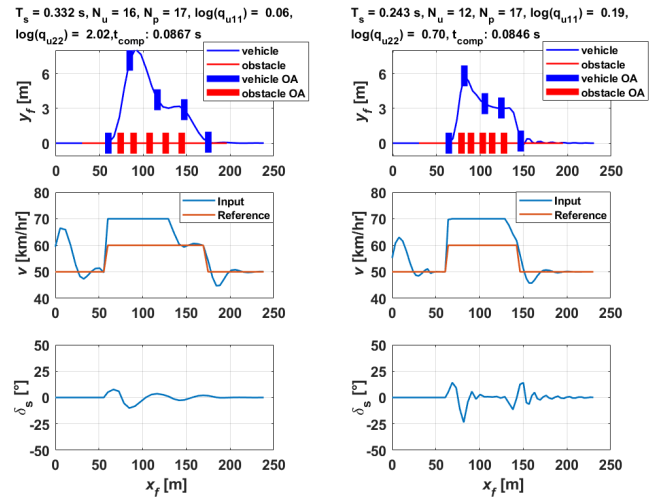


Fig. 3. Vehicle control query window. The top subplots show the location trajectories of the vehicle and the obstacle, in which the “vehicle OA” and “obstacle OA” bars show five relative positions of the vehicle and obstacle during the obstacle-avoidance phase. The middle subplots show the actual and reference velocity v at different longitudinal positions. The steering angle δ_s over the longitudinal position is depicted in the bottom subplots. For ease of assessment, the unit of v and δ_s in the figure is converted to km/hr and degree ($^\circ$), respectively. The results on the right panels are preferred.




Fig. 4. Vehicle control final performance obtained by the designed MPC controller.

be equal to 0.085 s, 0.310, 16, 0.261 and 0.918, respectively with a worst-case computational time t_{comp} needed to solve the QP problem (3) equal to 0.0808 s.

The final closed-loop results achieved by the designed MPC are depicted in Fig. 4, which demonstrates that solely based on calibrator's preferences, after only 50 experiments, the proposed algorithm is able to tune the MPC parameters with satisfactory performance. It is also important to remark that, for the same problem, the authors were not able to find, via several trial-and-error tests, a proper scoring function of the closed-loop performance to be used for a fully-automated calibration.

V. CONCLUSION

In this paper, a novel semi-automated MPC calibration approach was presented which turns out to be efficient in terms of number of experiments required for calibration. The key feature of the proposed methodology is that it allows calibration only based on pairwise preferences between the outcomes of the experiments, and thus it is very useful for calibration tasks with qualitative, subjective, or hard-to-quantify performance index functions, and with calibrators having limited MPC design knowledge. The same preference-based approach can be also used for calibration of other type of controllers, such as PIDs.

Future research activities are devoted to add new features to the calibration algorithm in order to further improve its performance. For instance, when the outcome of the experiment can be easily labelled as “very bad”, one can try to exploit this information instead of specifying a preference w.r.t. other experiments. A straightforward solution could be to train a classifier to separate the feasible space into good/bad regions, and thus reduce the search space.

REFERENCES

- [1] E. Brochu, V. M. Cora, and N. De Freitas, “A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning,” *arXiv preprint arXiv:1012.2599*, 2010.
- [2] A. Bemporad, “Global optimization via inverse distance weighting and radial basis functions,” *Computational Optimization and Applications*, vol. 77, pp. 571–595, 2020, code available at <http://cse.lab.imtlucca.it/~bemporad/glis>.
- [3] M. Forgione, D. Piga, and A. Bemporad, “Efficient calibration of embedded MPC,” in *Proc. of the 21st IFAC World Congress*, Berlin, Germany, 2020.
- [4] A. Lucchini, S. Formentin, M. Corno, D. Piga, and S. M. Savaresi, “Torque vectoring for high-performance electric vehicles: an efficient MPC calibration,” *IEEE Control Systems Letters*, 2020.
- [5] D. Piga, M. Forgione, S. Formentin, and A. Bemporad, “Performance-oriented model learning for data-driven MPC design,” *IEEE Control Systems Letters*, vol. 3, no. 3, pp. 577–582, 2019.
- [6] S. Bansal, R. Calandra, T. Xiao, S. Levine, and C. J. Tomiini, “Goal-driven dynamics learning via bayesian optimization,” in *Proc. of the IEEE 56th Annual Conference on Decision and Control*, 2017, pp. 5168–5173.
- [7] M. Fiducioso, S. Curi, B. Schumacher, M. Gwerder, and A. Krause, “Safe contextual bayesian optimization for sustainable room temperature PID control tuning,” *CoRR*, vol. abs/1906.12086, 2019. [Online]. Available: <http://arxiv.org/abs/1906.12086>
- [8] A. Marco, P. Hennig, J. Bohg, S. Schaal, and S. Trimpe, “Automatic LQR tuning based on gaussian process global optimization,” in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 270–277.
- [9] L. Roveda, M. Forgione, and D. Piga, “Two-stage robot controller auto-tuning methodology for trajectory tracking applications,” in *Proc. of the 21st IFAC World Congress*, Berlin, Germany, 2020.
- [10] D. Drieß, P. Englert, and M. Toussaint, “Constrained Bayesian optimization of combined interaction force/task space controllers for manipulations,” in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 902–907.
- [11] R. Calandra, N. Gopalan, A. Seyfarth, J. Peters, and M. P. Deisenroth, “Bayesian gait optimization for bipedal locomotion,” in *International Conference on Learning and Intelligent Optimization*. Springer, 2014, pp. 274–290.
- [12] F. Berkenkamp, A. P. Schoellig, and A. Krause, “Safe controller optimization for quadrotors with gaussian processes,” in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 491–496.
- [13] A. Bemporad and D. Piga, “Global optimization based on active preference learning with radial basis functions,” *Machine Learning*, vol. 110, pp. 417–448, 2020.
- [14] P. F. Christiano, J. Leike, T. Brown, M. Martic, S. Legg, and D. Amodei, “Deep reinforcement learning from human preferences,” in *Advances in Neural Information Processing Systems*, 2017, pp. 4299–4307.
- [15] C. Wirth, R. Akrou, G. Neumann, and J. Fürnkranz, “A survey of preference-based reinforcement learning methods,” *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 4945–4990, 2017.
- [16] M. Menner, K. Berntorp, M. N. Zeilinger, and S. Di Cairano, “Inverse learning for human-adaptive motion planning,” in *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 809–815.
- [17] S. Rosbach, V. James, S. Großjohann, S. Homoceanu, and S. Roth, “Driving with style: Inverse reinforcement learning in general-purpose planning for automated driving,” *arXiv preprint arXiv:1905.00229*, 2019.
- [18] W. Wang, J. Xi, A. Chong, and L. Li, “Driving style classification using a semisupervised support vector machine,” *IEEE Transactions on Human-Machine Systems*, vol. 47, no. 5, pp. 650–660, 2017.
- [19] S. Gros, M. Zanon, R. Quirynen, A. Bemporad, and M. Diehl, “From linear to nonlinear MPC: bridging the gap via the real-time iteration,” *International Journal of Control*, vol. 93, no. 1, pp. 62–80, 2020.
- [20] M. Diehl, H. G. Bock, and J. P. Schlöder, “A real-time iteration scheme for nonlinear optimization in optimal feedback control,” *SIAM Journal on Control and Optimization*, vol. 43, no. 5, pp. 1714–1736, 2005.
- [21] H. Gutmann, “A radial basis function method for global optimization,” *Journal of Global Optimization*, vol. 19, pp. 201–2227, 2001.
- [22] D. McDonald, W. Grantham, W. Tabor, and M. Murphy, “Global and local optimization using radial basis function response surface models,” *Applied Mathematical Modelling*, vol. 31, no. 10, pp. 2095–2110, 2007.
- [23] A. Smola and B. Schölkopf, “A tutorial on support vector regression,” *Statistics and Computing*, vol. 14, pp. 199–222, 2004.
- [24] M. Stone, “Cross-validated choice and assessment of statistical predictions,” *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 36, no. 2, pp. 111–133, 1974.
- [25] A. Vaz and L. Vicente, “PSwarm: A hybrid solver for linearly constrained global derivative-free optimization,” *Optimization Methods and Software*, vol. 24, pp. 669–685, 2009, <http://www.norg.uminho.pt/aivaz/pswarm/>.
- [26] M. McKay, R. Beckman, and W. Conover, “Comparison of three methods for selecting values of input variables in the analysis of output from a computer code,” *Technometrics*, vol. 21, no. 2, pp. 239–245, 1979.
- [27] MathWorks, *Statistics and Machine Learning Toolbox*, Natick, Massachusetts, United State, 2019. [Online]. Available: <https://it.mathworks.com/products/statistics.html>
- [28] —, *Model Predictive Control Toolbox*, Natick, Massachusetts, United State, 2019. [Online]. Available: <https://it.mathworks.com/products/mpc.html>
- [29] B. W. Bequette, *Process dynamics: modeling, analysis, and simulation*. Prentice hall PTR New Jersey, 1998.