



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Ira, AS;Manzie, C;Shames, I;Chin, R;Nešić, D;Nakada, H;Sano, T

Title:

Tuning of multivariable model predictive controllers through expert bandit feedback

Date:

2021-01-01

Citation:

Ira, A. S., Manzie, C., Shames, I., Chin, R., Nešić, D., Nakada, H. & Sano, T. (2021).
Tuning of multivariable model predictive controllers through expert bandit
feedback. *International Journal of Control*, 94 (10), pp.2650-2658. <https://doi.org/10.1080/00207179.2020.1727959>.

Persistent Link:

<https://hdl.handle.net/11343/251422>

Tuning of multivariable model predictive controllers through expert bandit feedback

Alex. S. Ira, Chris Manzie, Iman Shames, Robert Chin, Dragan Nešić, Hayato Nakada, Takeshi Sano

Abstract—For certain industrial control applications an explicit function capturing the nontrivial trade-off between competing objectives in closed loop performance is not available. In such scenarios it is common practice to use the human innate ability to implicitly learn such a relationship and manually tune the corresponding controller to achieve the desirable closed loop performance. This approach has its deficiencies because of individual variations due to experience levels and preferences in the absence of an explicit calibration metric. Moreover, as the complexity of the underlying system and/or the controller increase, in the effort to achieve better performance, so does the tuning time and the associated tuning cost. To reduce the overall tuning cost, a tuning framework is proposed herein, whereby a supervised machine learning is used to extract the human-learned cost function and an optimization algorithm that can efficiently deal with a large number of variables, is used for optimizing the extracted cost function. Given the interest in the implementation across many industrial domains and the associated high degree of freedom present in the corresponding tuning process, a Model Predictive Controller applied to air path control in a diesel engine is tuned for the purpose of demonstrating the potential of the framework.

I. INTRODUCTION

From a high level perspective, controller tuning involves modifying the gains to improve the closed loop performance of a system. For controllers with low degrees of freedom, such as PI loops that have been commonplace in industrial applications, this approach has typically involved online experiments and a human expert with sufficient experience to evaluate and then improve the quality of the response. However, with increasing degrees of freedom in modern controllers and the complexity in many systems, faster and more efficient methods of controller tuning are being sought to reduce the time burden. To utilize a digital twin of the system in an initial offline calibration, two linked problems can be addressed. The first of these relates to interpreting the goal of the tuning from the perspective of human experts. This is often not explicitly known, yet is instrumental in forming an optimization problem that can be tackled using an offline approach. The second problem involves posing an efficient solution to the resulting

optimization problem, which may involve high numbers of parameters, constraints and have non-unique solutions.

Because of the existence of human experts who, for a given data point (e.g., a pair of closed loop output trajectories) can provide a label (e.g., a quantitative or qualitative evaluation), supervised machine learning (SML) is proposed to address the first problem. Using human labeled data to learn a rank function (Xia, 2019) has been extensively used in addressing the information retrieval (IR) problem. Some applications are search engines, review and recommendation sites and services such as Netflix watch recommendation and Amazon buy recommendations. Inverse Reinforcement Learning (IRL) (Ng et al., 2000), can be thought of as an approach of extracting a (human) reward function for sequence of actions. Since the labelling consists of only one action, IRL is conceptually equivalent to SML.

Although similar challenges exist for any modern control architecture, in this work the focus is on the tuning of model predictive control (MPC) controllers. The reason for this choice is the growing prevalence of MPC in many industrial domains (Samad, 2017), along with the challenges it presents for traditional calibrators by having a large number of variables that are not directly linked to the closed loop system response.

Early approaches for the calibration of MPCs included tuning rules-of-thumb and general guidelines documented in (Garriga and Soroush, 2010; Morari and Lee, 1999; Qin and Badgwell, 2003; Rani and Unbehauen, 1997). However, these focused on only a few parameters of the MPC. Replacement of the existing controllers with MPC counterparts was considered in (Maciejowski, 2007). There, the guidelines on how to reverse-engineer the existing controller for the MPC design are provided. This served as a useful method for a rapid introduction of MPC, although its performance was limited to that achievable with the original controller.

More recently, metaheuristic optimizers, such as particle swarms in (Júnior et al., 2014), genetic algorithms in (Van der Lee et al., 2008) and a gradient-descent algorithm in (Bunin et al., 2012) have been considered. There are also metaheuristic off-the-shelf methods of goal attainment in (Exadaktylos and Taylor, 2010) and (Vega et al., 2008), and analytical approaches (employing appropriate simplifications) in (Bagheri and Khaki-Sedigh, 2014; Shah and Engell, 2010; Shridhar and Cooper, 1998). The potential shortcoming of the above approaches is they are not necessarily tailored for the optimization problem resulting from the proposed tuning framework. Thus, they may result in slower convergence properties or lower performance when included in an offline optimization routine involving a high-fidelity plant model. This becomes

Alex. S. Ira, Chris Manzie, Iman Shames, Robert Chin, Dragan Nešić, are with the Department of Electrical and Electronic Engineering at the University of Melbourne, Parkville, Australia. Robert Chin is also affiliated with the University of Birmingham, United Kingdom. E-mails: {alex.ira, manziec, ishames, dnesic}@unimelb.edu.au, chinr@student.unimelb.edu.au

Hayato Nakada and Takeshi Sano are with Advanced Unit Management System Development Division, Toyota Motor Corporation, Higashi-Fuji Technical Center, 1200, Mishuku, Susono-city, Shizuoka, 410-1193. E-mails: {hayato_nakada, takeshi_sano_aa}@mail.toyota.co.jp

particularly problematic as the controller complexity increases, and to date we are unaware of any of the approaches being adopted in an industrial context.

To address the potential shortcomings of existing MPC tuning algorithms not exploiting the full potential of the methodology and the issues with scalability for more generic approaches, in this paper we propose a framework to effectively tune advanced multivariable controllers such as MPC. In a similar vein to (Nguyen et al., 2017), a simulated human response to a closed loop output is generated, although here informed initially by real experts. The simulated responses can then be used within bespoke optimisers to efficiently tune the controller gains and deliver near-optimal closed loop responses. To demonstrate the proposed approach, the application considered here is a diesel engine air path control. The objective is to ensure multiple outputs track separate reference trajectories while satisfying the corresponding constraints. The experimental results are promising and provide incentive for further research and development.

The remainder of the paper is organized as follows. Section II presents the two main challenges associated with the proposed tuning framework. In Section III, the proposed framework is demonstrated by tuning an MPC controller which is used in a diesel engine air path control. The conclusions are provided in Section IV.

II. PROBLEM FORMULATION

Many industrial control applications achieve the desired *closed loop performance* via a tuning process. This involves many choices and tasks but essentially it is an optimization problem, solved by a human, see Fig. 1. This problem can be

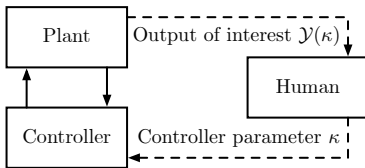


Fig. 1. Controller tuning framework with a human in the loop. Dashed lines indicate a different time scale to a closed loop time scale.

abstracted as the following optimization problem

$$\begin{aligned} & \text{minimize } \mathcal{C}_h(\mathcal{Y}(\kappa)) \\ & \text{subject to } \kappa \in K, \end{aligned} \quad (1)$$

where κ is a controller parameter, K is the set of admissible control parameters, $\mathcal{Y}(\kappa)$ is a closed loop output of interest with respect to a controller with the parameter κ , and $\mathcal{C}_h(\cdot)$ is an approximation of a cost function $\mathcal{C}(\cdot)$.

The cost function $\mathcal{C}(\cdot)$ captures a cost-benefit relationship between multiple, often competing, objectives. Unfortunately, more often than not, it is not available in an explicit form. Therefore, a common approach is to approximate it by exploiting the innate human ability to recognize and learn patterns, resulting in $\mathcal{C}_h(\cdot)$.

Even though humans are good at obtaining $\mathcal{C}_h(\cdot)$, they are not necessarily efficient at solving (1), particularly when there

is not an explicit relationship between the tuning variables and $\mathcal{C}_h(\cdot)$. A more efficient alternative might be to use a machine.

One approach to accomplish this is to first extract the human learned cost function $\mathcal{C}_h(\cdot)$, resulting in $\hat{\mathcal{C}}_h(\cdot)$. Given the availability of a human who can label¹ the corresponding closed loop performance, SML may be used (Abu-Mostafa et al., 2012).

Once $\hat{\mathcal{C}}_h(\cdot)$ is obtained, the tuning problem (1) is approximated by

$$\begin{aligned} & \text{minimize } \hat{\mathcal{C}}_h(\mathcal{F}(\mathcal{Y}(\kappa))) \\ & \text{subject to } \kappa \in K, \end{aligned} \quad (2)$$

where $\mathcal{F}(\cdot)$ is a feature extractor function which transforms an output of interest (time signal) into a vector consisting of relevant features, see Fig 2.

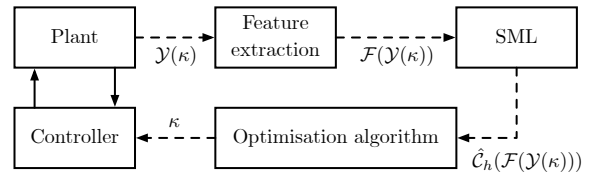


Fig. 2. Controller tuning framework with a machine in the loop. Dashed lines indicate a different time scale to a closed loop time scale.

There are two major steps in the transition from a tuning framework with a human-in-the-loop (as in Fig. 1) to the proposed automated tuning framework (as in Fig. 2): (i) estimation of a human learned cost function, $\mathcal{C}_h(\cdot)$, and (ii) selection of an optimization algorithm for solving (2).

A. Estimation of a human learned cost function

In this section SML is proposed to estimate the implicit cost function of an experienced calibrator using an appropriate regressor, for which there exist already well established tools (Abu-Mostafa et al., 2012).

One potentially time consuming aspect of the estimation of a human cost function is to collect the corresponding data used for training, development (validation) and testing. In this work the data set is defined as $\mathcal{D} := \{(\mathcal{Y}^{(i)}(\kappa), \mathcal{L}^{(i)}) = (\mathcal{Y}, \mathcal{L})^{(i)} | i \in \{1, \dots, N\}, N \in \mathbb{N}, \mathcal{Y} : \mathcal{I} \rightarrow \mathbb{R}^{p \times T}, \mathcal{L} \in \mathbb{R}_{\geq 0}\}$, where superscript $^{(i)}$ refers to the i -th data point, N is the number of data points, \mathcal{I} is the relevant time interval of the output of interest, $p \in \mathbb{N}$ and \mathcal{L}^i is the corresponding continuous label (or grading) assigned by the human for the given \mathcal{Y}^i .

One of the challenges in obtaining \mathcal{D} is related to the cost associated with labelling. Namely, the high cost may lead to small N which further may result in a poor estimate of the human learned cost function. In some cases, one might be able to increase N by using data synthesis together with active learning (Settles, 2012).

Another challenge is related to the extraction of the features of the output of interest \mathcal{Y} as it is not always clear which features to use. Depending on the control application and the

¹Quantitative and/or qualitative evaluation.

experience of the human expert, s/he may be able to provide an insight into which ones to use. Otherwise, one might use an appropriate statistical approach, such as principal component analysis, for selecting features.

Let $\mathcal{F} : \mathbb{R}^{p^x} \rightarrow \mathbb{R}^l$, $l \in \mathbb{N}$ be a feature extractor function resulting from human expert providing an insight or some appropriate statistical analysis. Then, the raw data set \mathcal{D} can be transformed into

$$\mathcal{D}_{\mathcal{F}} := \{(\mathcal{F}(\mathcal{Y}), \mathcal{L})^{(i)} | i \in \{1, \dots, N\}\} \quad (3)$$

which is used for training, development and testing.

Given a regressor the challenge is to obtain the corresponding data set $\mathcal{D}_{\mathcal{F}}$, which is dealt with on a case by case basis, as will be illustrated in Section III.

B. Gradient-free optimization

The final step in the proposed tuning framework of Fig. 2 is to select or develop a bespoke optimization algorithm.

As mentioned in Section I, one objective of the proposed framework is to efficiently deal with a large number of controller parameters (e.g., optimization variables). To that end, optimization approaches requiring no gradients or higher order derivatives, such as in (Kolda et al., 2003; Torczon, 1997), are required.

Further, in some cases (as demonstrated in the case study below) an approach which can efficiently accommodate constraints is necessary. This may require enforcement of certain restrictions on a problem so it becomes amenable to a certain family of algorithms.

The following case study is used to illustrate one possible implementation of the high level framework and demonstrate the potential benefits of this approach relative to the existing tuning methodologies.

III. CASE STUDY

In what follows, the proposed tuning framework is applied to the problem of air path control in a diesel engine using MPC. This is a particularly relevant example, as the potential for advanced control implementations in the diesel air path to improve efficiency have been widely reported in the research literature, yet the industry uptake of these controllers has been slow due to the time and cost of using existing tuning methods on the new controllers. This is coupled with the issue of the non-explicit relationship between tuning parameters in algorithms like MPC and their time domain outputs, along with the limited exposure of many engine calibrators to these control methods.

A. Engine air path model

An illustration of a simplified diesel engine air path is given in Fig. 3. Modern (automotive) diesel engines improve the fuel economy by using the variable geometry turbine (VGT), while they decrease the nitrogen oxide (NO_x) emissions by using the exhaust gas recirculation (EGR) (Huang et al., 2016). The dynamics of the air path is manipulated via throttle, EGR valve and VGT vane. In particular, the compressor, followed by the

intercooler, increases the density of the fresh air entering the air path. Then, the EGR system is used to cool and recirculate part of the burnt gas in the exhaust manifold. Finally, the part of the burnt gas is used to drive VGT which spins the compressor.

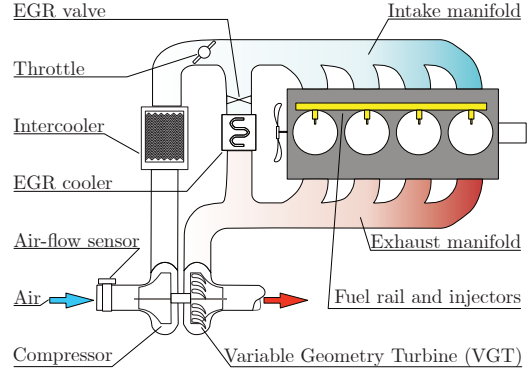


Fig. 3. Illustration of a diesel engine air path.

A diesel engine air path is typically modeled as a continuous-time system². Although slight variations on the model structure exist in (Broomhead et al., 2015; Wahlström and Eriksson, 2010, 2011), they can all be captured in the following general form

$$\dot{x} = f(x, u, \theta), \quad (4a)$$

$$y = h(x, u, \theta), \quad (4b)$$

where $x \in \mathcal{X} \subseteq \mathbb{R}^n$ is the state, $u \in \mathcal{U} \subseteq \mathbb{R}^m$ is the input, $\theta \in \Theta \subseteq \mathbb{R}^b$ is the parameter, $y \in \mathcal{Y} \subseteq \mathbb{R}^p$ is the output, mappings $f : \mathcal{X} \times \mathcal{U} \times \Theta \rightarrow \mathcal{X}$ and $h : \mathcal{X} \times \mathcal{U} \times \Theta \rightarrow \mathcal{Y}$ are nonlinear, and $n, m, b, p \in \mathbb{N}$.

A four-state model is adopted here (Huang et al., 2016; Sankar et al., 2019). The model consists of

$$x = (p_{in}, p_{ex}, W_{comp}, F_{EGR}), \quad (5a)$$

$$\theta = (w, w_{fuel}), \quad (5b)$$

$$u = (u_{thr}, u_{EGR}, u_{VGT}), \quad (5c)$$

$$y = (p_{in}, F_{EGR}), \quad (5d)$$

where p_{in} is the pressure in the intake manifold; p_{ex} is the pressure in the exhaust manifold; W_{comp} is the compressor mass flow rate; F_{EGR} is EGR rate; w is the engine speed; w_{fuel} is the volumetric fuel-rate; u_{thr} is the throttle position; u_{EGR} is the EGR valve position; u_{VGT} is the setting of VGT. Note that full state feedback is available either via direct sensor measurements or it can be reliably estimated within the engine control unit (ECU). It is possible to incorporate more states in order to improve the accuracy of the model predictions but the cost for doing that is the increased online computational load and the requirement for a state estimation.

One approach to handling the nonlinearities in the diesel air path is to use a switched linear-time-invariant (LTI) model to represent the engine. In such an approach, the engine operating range Θ is divided into N_r subregions, with one

²Note that this is a shorthand notation, for relevant details see Appendix A.

possible division of 12 regions illustrated in Fig. 4, where $\theta^{\#l} \in \Theta^{\#} := \{\theta \in \Theta \mid \forall w \in \{w^L, w^M, w^H\}, \forall w_{fuel} \in \{w_{fuel}^L, w_{fuel}^{ML}, w_{fuel}^{MH}, w_{fuel}^H\}\}, \forall l \in \{1, 2, \dots, 12\}$.

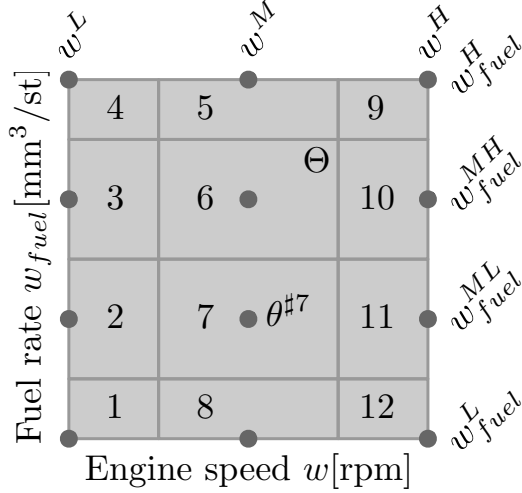


Fig. 4. Illustration of a parameter space Θ divided into 12 regions with respect to 12 evenly spaced operating points.

Each region is approximated with a linear model around the corresponding operating point. In particular, consider $\theta^{\#l} = \theta^{\#} \in \Theta^{\#}$, $l \in \{1, 2, \dots, N_r\}$. Then, the corresponding discrete-time linear approximation is denoted as

$$\tilde{x}(k+1) = A_l \tilde{x}(k) + B_l \tilde{u}(k), \quad (6a)$$

$$\tilde{y}(k) = C_l \tilde{x}(k) + D_l \tilde{u}(k), \quad (6b)$$

where \tilde{x} , \tilde{u} and \tilde{y} are the corresponding perturbed quantities (see (13)) and A_l , B_l , C_l and D_l are the corresponding matrices for the active region; see Appendix A for the details.

Diesel engine control is challenging due to the presence of multiple competing objectives and constraints. The overarching problem requires attaining good drivability and fuel economy whilst ensuring that the legislated emissions constraint is met (along with other constraints on states and inputs). This requires close tracking of boost pressure and EGR.

For a given engine operating condition θ , the corresponding reference $r(\theta) \in \mathbb{R}^2$ (see (12)) corresponds to the “optimal” driver demand responsiveness and fuel efficiency, and the satisfaction of the emission regulations. These optimal trajectories are often unable to be tracked exactly, and the calibration problem involves a tradeoff when attempting to track both sets of references. An experienced calibrator is well-versed in managing this tradeoff at different engine operating points.

B. Control architecture

In the switched-LTI approach, there are N_r MPC controllers with the active controller dependent on the value of θ . The MPC cost function is quadratic and it is denoted as

$$J(\tilde{x}(0), \{\tilde{u}\}) = \|\tilde{x}(N)\|_P + \sum_{i=0}^{N-1} (\|\tilde{x}(i)\|_Q + \|\tilde{u}(i)\|_R), \quad (7)$$

where the tuning parameters of interest consist of $P \in \mathbb{R}^{4 \times 4}$, $Q \in \mathbb{R}^{4 \times 4}$ and $R \in \mathbb{R}^{3 \times 3}$, while $\|v\|_M := v^T M v$, $M \in$

$\{P, Q, R\}$. Note that although the calibration problem considers the closed loop response of only two states on a high fidelity model, the MPC (for stability reasons) must contain a positive definite Q and so considers four states in an open loop optimization on a reduced order model.

The corresponding MPC optimization problem is defined as

$$\begin{aligned} & \text{minimize } J(\tilde{x}(0), \{\tilde{u}\}) \\ & \text{subject to } \begin{cases} (6a), \\ \tilde{x}(0) = x(kT_s) - x_r^*(\theta^{\#}), \\ \tilde{x} \in \mathcal{X}, \\ \tilde{u} \in \mathcal{U}, \end{cases} \end{aligned} \quad (8)$$

where $x(kT_s)$ is the sampled state of system (4) at time $t = kT_s$, $x_r^*(\theta^{\#})$ is “optimal” steady state value with respect to a given output reference $r(\theta^{\#})$, and \mathcal{X} and \mathcal{U} represent the relevant state and input constraints, respectively. The solution of the optimization problem (8) is a sequence of optimally predicted control values over the horizon N , that is

$$\{\tilde{u}^*\} := \{\tilde{u}^*(0), \tilde{u}^*(1), \dots, \tilde{u}^*(N-1)\}. \quad (9)$$

Only the first element of the sequence is applied to the engine at each time step.

For more details on the MPC problem formulation refer to Appendix A and Appendix B, and for details on the parameters used in the case study problem here (including a discussion on state and input constraints and prediction and control horizon lengths) interested readers are referred to (Shekhar et al., 2017).

C. Automated tuning framework

As discussed above, the first step towards the transition from a tuning framework with a human in the loop to an automated tuning framework is to identify the human learned cost function C_h . For the purposes of this paper, the following assumption is made.

Assumption 1 (Human learned cost function). *The human learned cost function C_h is invariant to $\theta \in \Theta$.* \square

This assumption is limiting in the context of industrial application, in that different priorities may arise at different engine operating points. For instance, the calibrator may preference greater torque in certain subregions, however this is dealt with simply by increasing the training dataset.

1) *Data collection and training stage:* The output trajectory of interest, $\mathcal{Y}(\kappa)$, is (6b) over the time interval \mathcal{I} and in this case study consists of the pressure in the intake manifold p_{in} (or boost pressure) and the EGR rate F_{EGR} .

The human expert uses plots like in Fig. 5 to evaluate the behavior of the corresponding outputs of interest, which can be considered to represent the performance of the chosen controller parameters. In order to extract the corresponding cost function the first step is necessary to generate a set of output trajectory pairs and query for its labels. This step can be achieved either by using closed loop responses, or by artificially generating trajectories. The advantage of the latter

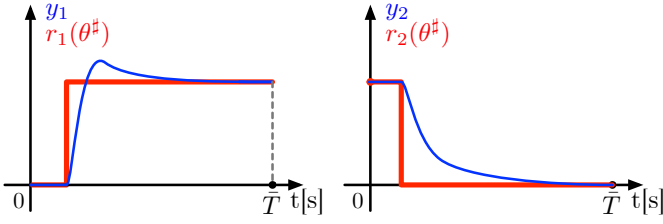


Fig. 5. A typical pair of normalized boost pressure (y_1) and normalised EGR rate (y_2) responses with respect to the step reference $r_1(\theta^\#)$ and $r_2(\theta^\#)$, respectively; T is the duration of a reference signal.

approach is the trajectories can explore a more specific part of the operational space at significantly lower computational effort, but not all trajectories may be achievable. This is not necessarily a significant issue as the key information being extracted is the labels of the trajectory pairs.

For the specific case study used here, the undershoot, overshoot, settling time and steady state error, for both boost pressure and EGR rate represent possible features of interest and are used to populate the \mathcal{Y}^i data. Each collected step response is a finite array of numerical values. Extracting the features of interest (undershoot, overshoot, settling time and steady state error) is done with a bespoke algorithm which outputs these values from the collected step response using the standard definitions (e.g. the step response has settled to within 1% of the steady state value). To allow for noisy measurements, a small tolerance band is applied to all the features (e.g. noise on the step response causing a temporary excursion from the 1% boundary does not invalidate the calculation). The data set, \mathcal{D} in (3), can then be populated from a number of trajectories with the extracted features as detailed above and their associated grading by the calibrator.

For this specific implementation of the framework, a multilayer feedforward neural network is used to approximate the relationship between the features and the labels provided by the human expert, \hat{C}_h . The developed network has two hidden layers, with 24 and 8 nodes in the first and second layers respectively. This choice of regressor typically requires at least 10 times as many data pairs as there are weights in the network (Abu-Mostafa et al., 2012), and so relevant trajectory generation is crucial in order to not overburden the human expert. Active learning approaches (Settles, 2012) may be useful in maximising the new information content in the trajectories presented to the human for labelling, and can be tailored to only present features in a relevant region of operation.

To avoid issues with the optimisation problem later, the function \hat{C}_h should have good coverage of the entire feature set, however this is at odds with attempts to minimise the trajectories considered by the human to those that are likely to be acceptable. One method of dealing with this issue is to synthesise labels for features that are clearly unacceptable and use this additional data to supplement the rated responses. We found that *modelling* a human expert who monotonically penalizes the trajectories as the features become more undesirable as a way of synthesising the required data (e.g. as the

overshoot increases, the synthesised label decreases).

2) *Tuning stage*: As discussed earlier, the engine operating range Θ is divided into N_r regions, each approximated with a separate linear model (6) and each governed with the corresponding MPC controller (7) and (8).

For notational simplicity, in the subsequent analysis we replace $\theta^{\#l}$ with $\theta^\#$ and intend that the relevant index l is always considered. Further, by Assumption 1, the same \hat{C}_h is used in all N_r optimization problems. The tuning problem then amounts to:

$$\begin{aligned} & \text{minimize } \hat{C}_h(f^y(P_{\theta^\#}, Q_{\theta^\#}, R_{\theta^\#})) \\ & \text{subject to } \begin{cases} P_{\theta^\#} = P_{\theta^\#}^\top, P_{\theta^\#} \succeq 0, \\ Q_{\theta^\#} = Q_{\theta^\#}^\top, Q_{\theta^\#} \succeq 0, \\ R_{\theta^\#} = R_{\theta^\#}^\top, R_{\theta^\#} \succ 0, \end{cases} \end{aligned} \quad (10)$$

where symmetric constraints³ lower the number of parameters from 46 to 26, while definiteness constraints are due to stability reasons, (Rawlings and Mayne, 2009). The tuning stage for region $l \in \{1, 2, \dots, N_r\}$ is performed in a simulation setting using the digital twin as illustrated in Fig. 6.

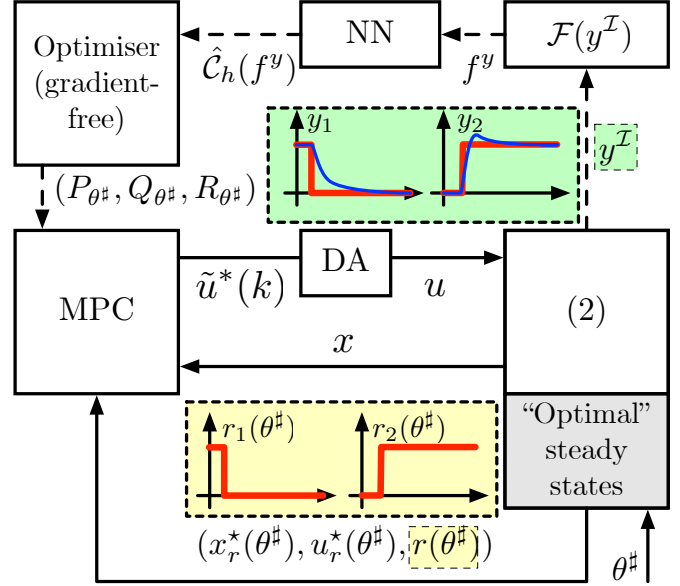


Fig. 6. Block diagram illustration of tuning MPC controller with respect to region $l \in \{1, 2, \dots, 12\}$. DA stands for digital-to-analog.

The chosen optimizer utilizes an algorithm for random search for nonsmooth and stochastic optimization, (Nesterov and Spokoiny, 2011, Section 4). This algorithm considers constrained optimization problems, i.e., it does not need gradient or any higher derivative information. The former is important because of constraints in (10) while the latter is important for practical purposes. The tuning parameters are the elements of $P \in \mathbb{R}^{4 \times 4}$, $Q \in \mathbb{R}^{4 \times 4}$ and $R \in \mathbb{R}^{3 \times 3}$ matrices. Finding the gradient or any other higher derivative information would result in a considerable computational cost. Details regarding the optimization algorithm used to solve (10) are provided in Appendix C.

³In addition to lowering the number of parameters, they are imposed due to implementation reasons, (Ira et al., 2018).

The performance of the tuned switched-MPC architecture is then ready to be experimentally tested over the extra-urban driving cycle (EUDC).

Remark 1. We note that although the implemented controller is MPC, other advanced controllers with tuning gains replacing $(P_{\theta^\sharp}, Q_{\theta^\sharp}, R_{\theta^\sharp})$ can be substituted into (10). However, the constraints on the optimisation variables must be altered from requiring positive semi-definite matrices as in (10) to appropriate conditions to guarantee stability under the chosen controller architecture.

D. Experimental results

The experimental testing of the tuning efficacy of the proposed framework was performed at Toyota’s Higashi-Fuji Technical Center in Susono, Japan. The test bench is equipped with a diesel engine and a transient dynamometer. A dSPACE DS1006 real-time processor board is used to implement the switched MPC architecture. A block diagram of the experimental configuration is shown in Fig. 7. The ECU logs sensor data from the engine and transmits the current state information to the controller. Note that ECU directly controls all engine subsystems. However, its commands for the three actuators, i.e., throttle, EGR valve and VGT, can be overridden with the MPC commands by toggling the switch (from the ControlDesk interface), see Fig. 7. Finally, depending on the current engine speed and fueling rate, i.e., θ , the corresponding MPC controller is selected based on the switched MPC architecture.

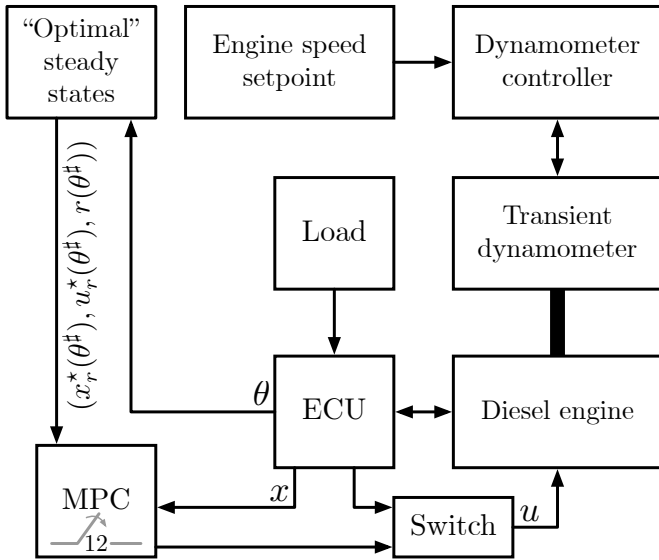


Fig. 7. Experimental setup.

For real time implementation of the controller on the dSPACE platform, it is desirable to choose a quadratic program (QP) solver that runs on embedded hardware. To this end, the MPC is formulated using the condensed formulation (Jerez et al., 2011) (which reduces the number of decision variables in the QP) while the choice of the solver is QPKWIK algorithm (Schmid and Biegler, 1994). This solver is implemented

in MATLAB’s MPC toolbox and it is used for both tuning (in simulation time) and testing (in real time).

The tracking performance of the tuned switched MPC architecture is tested over the EUDC. The results are displayed in Fig. 8.

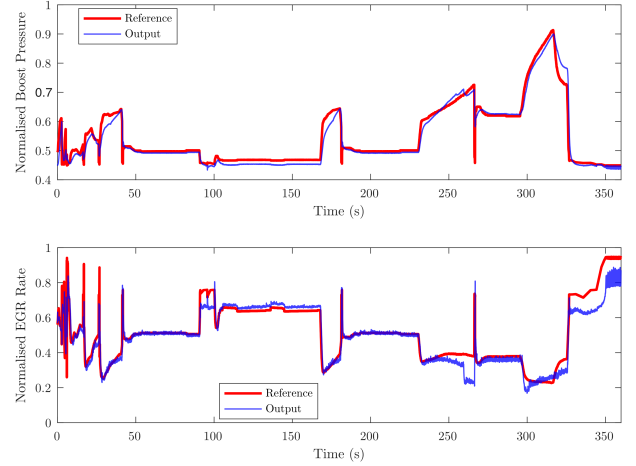


Fig. 8. Tuned switched MPC architecture performance on EUDC.

The performance in Fig. 8 is promising but there are tracking errors on both trajectories. In what follows, two possible error origins are presented and potential solutions are discussed.

The first possible origin is insufficient fidelity of the plant model (4). Since the proposed tuning framework is implemented in the offline setting of Fig. 6, and so the corresponding results depend on the accuracy of the approximation (4). Two possible solutions are either to i) use a higher-fidelity model which would increase the computational cost; or ii) implement the tuning structure online which would amount to replacing model (4) in Fig. 6 with a real engine. Indeed, the latter may appear appealing from an accuracy perspective but would incur an implementation cost as the time per optimisation step is increased and requires human intervention. One might address this through mixing the approaches by initialising the online tuning with an offline optimised solution to reduce the number of optimisation iterations requiring the plant in-the-loop. This has no strict guarantees on improvements but practically may be beneficial.

The second potential origin for tracking errors is the amount of data collected or the number of features considered to approximate the human cost function. Increasing the amount of data collected (as would be required through removing Assumption 1 for example) would address this problem but also comes at a cost.

Both of these issues are of course fundamental to data science and therefore not unexpected.

IV. CONCLUSIONS

A tuning framework, using machine learning to extract the human learned cost function, is implemented and experimentally tested. The results are promising but there are also areas for improvement. The issues related to the approximation of

the human learned cost function can be addressed by collecting more data. However, this increases the overall cost. The issues related to the plant model can be addressed by using a higher-fidelity models or implementing the tuning framework in an online setting.

Even though better results might necessitate more data which would lead to the implementation cost increase, it is important to note that extracting the human learned cost function might be one-off investment. Namely, once the cost function is extracted, the tuning reduces to solving an offline optimization problem that may be applicable across multiple engine models.

Finally, although the case studies presented here focus on MPC-based architectures, the high level framework is equally deployable to other existing and proposed algorithmic solutions to closed loop control. We also note that other combinations of optimiser and machine learning tools may provide superior performance than that depicted in the case study, although the same core elements as the algorithms used herein should be retained.

APPENDIX

A. Approximation of system dynamics

Consider the following continuous-time nonlinear system

$$\frac{dx(t)}{dt} = f(x(t), u(t), \theta(t)), \quad (11a)$$

$$y(t) = h(x(t), u(t), \theta(t)), \quad (11b)$$

where $x \in \mathcal{X} \subseteq \mathbb{R}^n$ is the state, $u \in \mathcal{U} \subseteq \mathbb{R}^m$ is the input, $\theta \in \Theta \subseteq \mathbb{R}^b$ is the parameter and $y \in \mathcal{Y} \subseteq \mathbb{R}^p$ is the output of the system (11), $t \in \mathbb{R}_{\geq 0} \subset \mathcal{T}$ is the continuous-time, $n = 4, m = 3, b = 2$ and $p = 2$ (see (5)).

Now, consider a scenario where system dynamics (4) is such that practically it is expensive to consider it for each $\theta \in \Theta$. In such scenarios, one can resort to a *discretization* of the parameter space Θ , and an approximation of (4) with respect to the resulting discrete values.

Discretized version of the parameter space $\Theta \subseteq \mathbb{R}^2$ is defined as $\Theta^\# := \{\theta = (\theta_1, \theta_2) \in \Theta | \forall \theta_1 \in \Theta_1, \forall \theta_2 \in \Theta_2\}$, where $\theta_1 = w$, $\theta_2 = w_{fuel}$, $\Theta_1 = \{w^L, w^M, w^H\}$ and $\Theta_2 = \{w_{fuel}^L, w_{fuel}^{ML}, w_{fuel}^{MH}, w_{fuel}^H\}$.

Now, consider a $\theta^\# \in \Theta^\#$ and system (11), and let a given reference value be denoted with $r(\theta^\#) \in \mathcal{Y}$, while the corresponding state and input with $x_r(\theta^\#)$ and $u_r(\theta^\#)$, respectively. Determining the pair $(x_r(\theta^\#), u_r(\theta^\#))$ consists of solving the following system of equations subject to operational and input constraints:

$$0 = f(x_r(\theta^\#), u_r(\theta^\#), \theta^\#), \quad (12a)$$

$$r(\theta^\#) = h(x_r(\theta^\#), u_r(\theta^\#), \theta^\#). \quad (12b)$$

yields a feedforward mapping of input-reference pairs. $u_r^*(\theta^\#)$ and $x_r^*(\theta^\#)$.

To improve the transient response, in this case study a MPC controller is used in conjunction with the feedforward control. At each time instant, MPC solves an optimization problem, it is desirable to reduce the complexity of the considered plant

dynamics (if possible) and to transform it from continuous-time to discrete-time. To that end, consider again a $\theta^\# \in \Theta^\#$, system (11), and let $r(\theta^\#) \in \mathcal{Y}$ be the given reference value. Next, define the following perturbed quantities

$$\tilde{x}(t, \theta^\#) := x(t) - x_r^*(\theta^\#), \quad (13a)$$

$$\tilde{u}(t, \theta^\#) := u(t) - u_r^*(\theta^\#), \quad (13b)$$

$$\tilde{y}(t, \theta^\#) := y(t) - r(\theta^\#). \quad (13c)$$

(Note that in the sequel, when appropriate, for notational convenience, time t and parameter $\theta^\#$ are dropped out in (13).) After linearisation and discretisation of (11), one arrives at:

$$\tilde{x}(k+1) = A\tilde{x}(k) + B\tilde{u}(k), \quad (14a)$$

$$\tilde{y}(k) = C\tilde{x}(k) + D\tilde{u}(k). \quad (14b)$$

B. MPC cost function

At each time instant k , the MPC controller samples the state of the system (11). This sampled state x_k is used to initialize the plant model (14). The cost of using a sequence of control values $\{\tilde{u}\} := \{\tilde{u}(k), \tilde{u}(k+1), \dots, \tilde{u}(k+N-1)\}$ on the plant model (14), over a horizon $N \in \mathbb{N}$, is defined as

$$J(\tilde{x}(0), \{\tilde{u}\}) := \|\tilde{x}(N)\|_P + \sum_{i=0}^{N-1} (\|\tilde{x}(i)\|_Q + \|\tilde{u}(i)\|_R),$$

where matrix $P \in \mathbb{R}^{n \times n}$, $Q \in \mathbb{R}^{n \times n}$ and $R \in \mathbb{R}^{m \times m}$, $\|v\|_M := v^T M v$, $M \in \{P, Q, R\}$ and v is the relevant vector.

C. Optimization algorithm

To obtain a desirable behaviour of the closed loop system using the controller of the preceding section, one needs to tune parameters P, Q and R. In this section we define the concept of optimisation using a random oracle as defined below for the MPC tuning problem:

Definition 1 (Random Oracle). *Let M_j^P, M_j^Q and M_j^R be random symmetric matrices of appropriate dimensions and definitenesses. Then,*

$$g_\mu(P_j) := \left((\hat{C}_h(f^y(P_j + \mu M_j^P, Q_j + \mu M_j^Q, R_j + \mu M_j^R)) - \hat{C}_h(f^y(P_j, Q_j, R_j))) M_j^P \right) / \mu$$

is a random oracle for P_j, Q_j , and R_j . \square

Recall that due to the stability requirements, matrices P, Q and R need to have a particular structure. In order to generate a random matrix with such a structure, a random unitary matrix, constructed using a technique described in (Ozols, 2009), is used. The algorithm based on (Nesterov and Spokoiny, 2011, Section 4), applied to the present problem, consists of the following steps,

Choose: P_0, Q_0, R_0 ,

While: $j \leq N, N \in \mathbb{N}_0$,

Compute: $P_{j+1} \leftarrow \pi_P(P_j - h_j B^{-1} g_\mu(P_j))$,
 $Q_{j+1} \leftarrow \pi_Q(Q_j - h_j B^{-1} g_\mu(Q_j))$,
 $R_{j+1} \leftarrow \pi_R(R_j - h_j B^{-1} g_\mu(R_j))$,
 $j \leftarrow j + 1$,

where $h_j > 0$, while π_P and π_Q are projections on a positive semidefinite cone and π_R is a projection on a positive definite cone⁴.

Note that the approximated cost function \hat{C}_h might consist of many local extrema. Thus, the performance of the used algorithm depends on the initial point (P_0, Q_0, R_0) . Therefore, a generation of a batch of initial points is performed and a promising candidate is chosen. The promising candidate is defined as the initial point (P_0, Q_0, R_0) , such that $\hat{C}_h(f^y(P_0, Q_0, R_0))$ achieves the smallest value with respect to the corresponding batch. In this work, μ is chosen to be very small and h_j is defined according to (42) in (Nesterov and Spokoiny, 2011). Specifically $\mu = 10^{-9}$ and $h_j = 10^{-6}/\sqrt{j+1}$, $j \in [1, 2, \dots, 50]$, are used for all controllers in 12 regions introduced in Fig. 4.

REFERENCES

- Abu-Mostafa, Y. S., Magdon-Ismail, M., and Lin, H.-T. (2012). *Learning from data*, volume 4. AMLBook New York, NY, USA.
- Bagheri, P. and Khaki-Sedigh, A. (2014). An analytical tuning approach to multivariable model predictive controllers. *Journal of Process Control*, 24(12):41–54.
- Broomhead, T., Manzie, C., Brear, M., and Hield, P. (2015). Model reduction of diesel mean value engine models. Technical report, SAE Technical Paper.
- Bunin, G. A., Tirado, F. F., François, G., and Bonvin, D. (2012). Run-to-run mpc tuning via gradient descent. In *Computer Aided Chemical Engineering*, volume 30, pages 927–931. Elsevier.
- Exadaktylos, V. and Taylor, C. J. (2010). Multi-objective performance optimisation for model predictive control by goal attainment. *International Journal of Control*, 83(7):1374–1386.
- Garriga, J. L. and Soroush, M. (2010). Model predictive control tuning methods: A review. *Industrial & Engineering Chemistry Research*, 49(8):3505–3515.
- Huang, M., Zaseck, K., Butts, K., and Kolmanovsky, I. (2016). Rate-based model predictive controller for diesel engine air path: Design and experimental evaluation. *IEEE Transactions on Control Systems Technology*, 24(6):1922–1935.
- Ira, A., Shames, I., Manzie, C., Chin, R., Nešić, D., Hayato, N., and Sano, T. (2018). A machine learning approach for tuning model predictive controllers. In *Proceedings of the 15th International Conference on Control, Automation, Robotics and Vision*.
- Jerez, J. L., Kerrigan, E. C., and Constantinides, G. A. (2011). A condensed and sparse qp formulation for predictive control. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, pages 5217–5222. IEEE.
- Júnior, G. A. N., Martins, M. A., and Kalid, R. (2014). A psobased optimal tuning strategy for constrained multivariable predictive controllers with model uncertainty. *ISA transactions*, 53(2):560–567.
- Kolda, T. G., Lewis, R. M., and Torczon, V. (2003). Optimization by direct search: New perspectives on some classical and modern methods. *SIAM review*, 45(3):385–482.
- Maciejowski, J. (2007). Reverse-engineering existing controllers for mpc design. *IFAC Proceedings Volumes*, 40(20):436–441.
- Morari, M. and Lee, J. H. (1999). Model predictive control: past, present and future. *Computers & Chemical Engineering*, 23(4-5):667–682.
- Nesterov, Y. and Spokoiny, V. (2011). Random gradient-free minimization of convex functions. Technical report, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE).
- Ng, A. Y., Russell, S. J., et al. (2000). Algorithms for inverse reinforcement learning. In *Icml*, pages 663–670.
- Nguyen, K., Daumé III, H., and Boyd-Graber, J. (2017). Reinforcement learning for bandit neural machine translation with simulated human feedback. *arXiv preprint arXiv:1707.07402*.
- Ozols, M. (2009). How to generate a random unitary matrix.
- Qin, S. J. and Badgwell, T. A. (2003). A survey of industrial model predictive control technology. *Control engineering practice*, 11(7):733–764.
- Rani, K. Y. and Unbehauen, H. (1997). Study of predictive controller tuning methods. *Automatica*, 33(12):2243–2248.
- Rawlings, J. B. and Mayne, D. Q. (2009). *Model predictive control: Theory and design*. Nob Hill Pub.
- Samad, T. (2017). A survey on industry impact and challenges thereof [technical activities]. *IEEE Control Systems*, 37(1):17–18.
- Sankar, G. S., Shekhar, R. C., Manzie, C., Sano, T., and Nakada, H. (2019). Fast calibration of a robust model predictive controller for diesel engine airpath. *IEEE Transactions on Control Systems Technology*.
- Schmid, C. and Biegler, L. T. (1994). Quadratic programming methods for reduced hessian sqp. *Computers & chemical engineering*, 18(9):817–832.
- Settles, B. (2012). Active learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 6(1):1–114.
- Shah, G. and Engell, S. (2010). Tuning mpc for desired closed-loop performance for siso systems. In *Control & Automation (MED), 18th Mediterranean Conference*, pages 628–633.
- Shekhar, R. C., Sankar, G. S., Manzie, C., and Nakada, H. (2017). Efficient calibration of real-time model-based controllers for diesel engines part i: Approach and drive cycle results. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 843–848.
- Shridhar, R. and Cooper, D. J. (1998). A tuning strategy for unconstrained multivariable model predictive control. *Industrial & engineering chemistry research*, 37(10):4003–4016.
- Torczon, V. (1997). On the convergence of pattern search algorithms. *SIAM Journal on optimization*, 7(1):1–25.
- Van der Lee, J., Svrcek, W., and Young, B. (2008). A tuning

⁴Recall that because $P = P^\top$, $Q = Q^\top$ and $R = R^\top$ one can first find the corresponding eigenvalue decompositions. Then, for the projection on a positive semidefinite cone (in the case of P and Q) instead of using the diagonal matrix, say Λ (which contains the eigenvalues), take $\Lambda_+ := \max(\Lambda, 0)$; which has non-negative diagonal values. On the other hand, for a projection on a positive definite cone (in the case of R) use $\Lambda_{++} := \max(\Lambda, d)$, $d \in \mathbb{R}_{>0}$; which has positive diagonal elements (in simulations, $d = 10^{-15}$).

- algorithm for model predictive controllers based on genetic algorithms and fuzzy decision making. *ISA transactions*, 47(1):53–59.
- Vega, P., Francisco, M., and Tadeo, F. (2008). Multiobjective optimization for automatic tuning of robust model based predictive controllers. *IFAC Proceedings Volumes*, 41(2):6980–6985.
- Wahlström, J. and Eriksson, L. (2010). Modeling of a diesel engine with intake throttle, vgt, and egr. *Department of Electrical Engineering, Linköpings Universitet, SE-581*, 83.
- Wahlström, J. and Eriksson, L. (2011). Modelling diesel engines with a variable-geometry turbocharger and exhaust gas recirculation by optimization of model parameters for capturing non-linear system dynamics. *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, 225(7):960–986.
- Xia, L. (2019). *Learning and Decision-Making from Rank Data*. Morgan & Claypool.